



ELSEVIER

Contents lists available at SciVerse ScienceDirect

Ultramicroscopy

journal homepage: www.elsevier.com/locate/ultramic

Accurate segmentation of dense nanoparticles by partially discrete electron tomography

T. Roelandts^{a,*}, K.J. Batenburg^{a,b}, E. Biermans^c, C. Kübel^d, S. Bals^c, J. Sijbers^a

^a IBBT-Vision Lab University of Antwerp, Universiteitsplein 1, 2610 Wilrijk, Belgium

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

^c EMAT, University of Antwerp, Groenenborgerlaan 171, 2020 Antwerp, Belgium

^d Institute of Nanotechnology, Karlsruhe Institute of Technology, Hermann-von-Helmholtz-Platz 1, 76344 Eggenstein-Leopoldshafen, Germany

ARTICLE INFO

Article history:

Received 17 June 2011

Received in revised form

12 October 2011

Accepted 22 December 2011

Available online 4 January 2012

Keywords:

Electron tomography

Segmentation

Discrete tomography

ABSTRACT

Accurate segmentation of nanoparticles within various matrix materials is a difficult problem in electron tomography. Due to artifacts related to image series acquisition and reconstruction, global thresholding of reconstructions computed by established algorithms, such as weighted backprojection or SIRT, may result in unreliable and subjective segmentations. In this paper, we introduce the *Partially Discrete Algebraic Reconstruction Technique* (PDART) for computing accurate segmentations of dense nanoparticles of constant composition. The particles are segmented directly by the reconstruction algorithm, while the surrounding regions are reconstructed using continuously varying gray levels. As no properties are assumed for the other compositions of the sample, the technique can be applied to any sample where dense nanoparticles must be segmented, regardless of the surrounding compositions. For both experimental and simulated data, it is shown that PDART yields significantly more accurate segmentations than those obtained by optimal global thresholding of the SIRT reconstruction.

© 2012 Elsevier B.V. All rights reserved.

1. Introduction

Electron tomography deals with the reconstruction of a three-dimensional (3D) representation of a microscopy sample from a tilt series of two-dimensional (2D) images. This technique has been applied successfully in materials science since the late 1980s [1]. Several imaging modes have been used for acquiring the projection images, in particular, bright-field TEM [2,3], annular dark field TEM [4], high-angle annular dark-field scanning TEM (HAADF STEM) [5–9], and energy-filtered TEM (EFTEM) [10–13].

Quantitative interpretation of the reconstructed 3D volume is often hampered by the presence of *artifacts*: structured distortions that do not correspond with the actual sample. In particular, limits on the number of projection images imposed by sample contamination or beam damage give rise to such artifacts. Furthermore, the limited spacing for specimen holders in between the pole pieces of the objective lens often restricts the range of tilt angles to about $\pm 70^\circ$, leading to a missing wedge in the collected data. As a consequence, features perpendicular to the electron beam are better resolved than features parallel to the beam, resulting in anisotropic resolution and distortions of the structure.

* Corresponding author. Tel.: +32 32652840; fax: +32 32652245.
E-mail address: tom.roelandts@ua.ac.be (T. Roelandts).

For many imaging tasks in materials science, the goal is to obtain an accurate *segmentation* of particular structures (i.e., particles, pores, tubules, etc.). Of particular importance is the problem of segmenting nanoparticles within various matrix materials [5–7,14]. Due to artifacts related to image acquisition and reconstruction, segmenting these structures from gray level volumes computed by established algorithms, such as weighted backprojection (WBP) or SIRT [15], may result in unreliable and subjective segmentations. In practice, reconstructions are often segmented using a global threshold. Since the threshold is estimated visually, this approach is highly subjective. Moreover, it does not account for the effect that the intensity of the features in the reconstruction strongly depends on their size [16]. Fully manual segmentation may avoid this effect, but remains a time consuming and subjective approach.

Recently, discrete tomography algorithms have demonstrated the ability to overcome some of these limitations by exploiting prior knowledge. Discrete tomography is based on the assumption that the sample consists of only a few different compositions. Two rather different variants of discrete tomography have been applied to electron tomography. The first variant was recently applied to the reconstruction of crystalline nanoparticles at atomic resolution [17,18]. For this variant, it is assumed that the crystal contains only a few atomic species, and that the atoms lie on a regular grid. Together, these assumptions allow to create a reconstruction from as few as two or three projections. For the second variant, which can be applied at lower resolutions, it is

only assumed that the sample consists of a few different compositions, each corresponding to a particular gray level in the reconstructed image. The discrete tomography algorithms that appear in this paper are of the second variant.

Major advantages of discrete tomography algorithms are that they require fewer projection images compared to alternative methods such as SIRT, and that missing wedge artifacts are strongly reduced [19]. Moreover, as the final result of the reconstruction process is a segmented image, a separate segmentation step is no longer required. The *Discrete Algebraic Reconstruction Technique* (DART) for discrete tomography has been successfully applied to a broad range of material samples [20–23]. The main restriction for using discrete tomography is that the entire sample must satisfy the discreteness requirement. If the sample contains a mixture of compositions, the results of discrete tomography cannot be relied upon, as the key assumptions are violated.

In this paper, we introduce the *Partially Discrete Algebraic Reconstruction Technique* (PDART) for computing accurate segmentations of dense nanoparticles of constant composition, regardless of the compositions in the remaining part of the sample. Embedded nanoparticles such as catalyst particles are often dense structures compared to their surroundings (e.g., porous materials), resulting in a high gray level in the reconstructed image. PDART is based on the assumption that the densest composition occurs in homogeneous regions that have a constant gray level. These dense regions are segmented discretely, while the surrounding regions are reconstructed using continuously varying gray levels. If the assumption of a homogeneous densest composition holds, the imaging mode that is used to record the tilt series – HAADF STEM for both samples in this paper – is not a restriction on the applicability of PDART, as long as the selected imaging mode is compatible with tomography. PDART imposes no restrictions on the nature of the sample (except that the densest composition must be homogeneous), which means that the application of the algorithm is not restricted to any specific type of samples.

This paper is structured as follows. In Section 2, the problem of segmenting dense particles is introduced, and the PDART algorithm is defined. Section 2 also introduces the figure of merit that is used for quantitative evaluation of the results. It concludes by describing how the parameters of the algorithm can be optimized automatically. In Section 3, the capabilities of PDART are assessed using two different experimental datasets and a number of simulation experiments. The results are discussed in Section 4 and conclusions are drawn in Section 5.

2. Algorithm

Before describing the PDART algorithm, we start by giving an example of its applicability. Fig. 1 illustrates the problem of nanoparticle segmentation. Fig. 1a shows a *phantom* (i.e., a simulated image), representing a microscopy sample that contains

nanoparticles of only a few pixels each, embedded in a cylinder of varying composition. From this phantom, a synthetic dataset was created by calculating 28 evenly spaced projections in the range of $\pm 70^\circ$. Fig. 1 also shows WBP (Fig. 1b), SIRT (Fig. 1c), DART (Fig. 1d), and PDART (Fig. 1e) reconstructions of this dataset.

The gray level reconstructions computed by WBP and SIRT have limited visual quality, as a result of the small number of projection angles and their limited angular range. When thresholding these images to determine the size and shape of the particles, it is not clear how the threshold should be chosen in an optimal way. The DART reconstruction, shown in Fig. 1d, is already segmented, yet the segmentation is not accurate at all when compared to the original phantom. The varying composition of the disk surrounding the nanoparticles violates the key discreteness assumption imposed by the DART algorithm. The PDART reconstruction, shown in Fig. 1e, seems much more accurate than the other reconstructions.

2.1. Algorithm description

The PDART algorithm has been designed to allow for accurate particle segmentation in cases where neither continuous methods nor fully discrete tomography leads to good results. The algorithm is based on the assumption that the particles have a constant composition, and that this composition represents the highest gray level in the reconstructed volume. PDART combines an iterative reconstruction algorithm, such as SIRT, with intermediate segmentation steps. Once pixels have been identified as “particle”, they are directly segmented (i.e., their value is set to the constant gray level for the particles) and kept fixed at this value in subsequent SIRT iterations. Note that, throughout this paper, we use the additive variant of SIRT, as described in [24].

Fig. 2 shows a flowchart of the PDART algorithm. Besides having the projection data as input, the algorithm has two parameters: a threshold τ and a gray level $\rho > \tau$, which corresponds to the gray level of the particles. Optimal values for both parameters can be determined automatically, as is outlined in Section 2.3.

Initially, the set F of *fixed pixels* is empty. In an iterative loop, the algorithm starts by performing one or more SIRT iterations on the entire image volume. Whenever one or more pixels are assigned a higher gray level than the threshold τ , it is decided that these pixels belong to a particle. Such pixels are added to F : their gray level is set to ρ and is kept fixed at this value during all subsequent SIRT iterations. In this way, the set F gradually expands as pixels are added, until some termination condition is satisfied. Typically, one aims for terminating the algorithm when no new pixels have been added to F for a sufficiently large number of iterations.

In its original form, the SIRT algorithm computes a weighted least square solution of the system $\mathbf{W}\mathbf{x} = \mathbf{p}$, where \mathbf{x} denotes the unknown image, \mathbf{p} denotes the projection data, and \mathbf{W} denotes the




Fig. 1. A simulation phantom and several reconstructions. The phantom represents a cylindrical sample that contains nanoparticles of only a few pixels each, embedded in a material of varying composition. (a) Phantom, (b) WBP, (c) SIRT, (d) DART and (e) PDART.

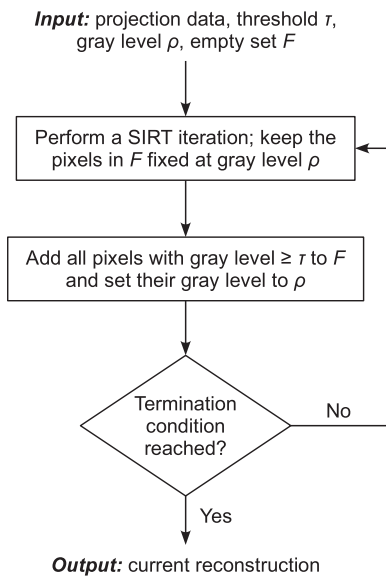


Fig. 2. Flowchart of the PDART algorithm.




Fig. 3. Illustration of several stages of the PDART algorithm. (a) Phantom. (b) Iteration 8 of PDART, before the application of the threshold. (c) Pixels that are added to F at iteration 8. (d) Iteration 8, after the application of the threshold. (e) PDART at iteration 57. (f) PDART at iteration 150.

projection operator. As pixels are added to F within the PDART algorithm, the number of unknowns in this system is reduced, while the number of equations remains the same, thereby improving the reconstruction for pixels that do not (yet) belong to F .

As an illustration, Fig. 3 shows several intermediate steps of the PDART algorithm, using the phantom from Fig. 1a. In this particular example, the first pixels cross the threshold τ at the eighth iteration (Fig. 3b). At this point, the result is still exactly the same as for SIRT, since no pixels have been fixed yet. The pixels that crossed the threshold are then added to the set F of fixed pixels (Fig. 3c) and fixed for the remainder of the reconstruction process (Fig. 3d). During the following iterations, the process of discovering and fixing more and more particle pixels continues until all of them are found at, in this example, iteration 57 (Fig. 3e). No new pixels are found during the following iterations, although the reconstruction quality of the background keeps improving somewhat. The result after 150 iterations is shown in Fig. 3f.

2.2. Figure of merit

After computing a reconstruction, the projections of the reconstructed image can be computed, and subsequently compared with the measured projections. The difference between the computed projections and the measured dataset is calculated by taking the sum of squares of the differences for all projection pixels, resulting in a number that indicates how well the reconstruction adheres to the measured projection data. Mathematically, this is known as the *projection distance*, defined as $d_{pr}(\mathbf{x}) = \|\mathbf{W}\mathbf{x} - \mathbf{p}\|_2$. It is this projection distance that we use as a figure of merit. For phantoms, we can also calculate the *phantom distance*, which is defined as $d_{ph}(\mathbf{x}) = \|\mathbf{x} - \mathbf{h}\|_2$, where \mathbf{h} denotes the phantom image. The phantom distance directly measures the difference between the reconstruction and the phantom.

2.3. Parameter optimization

In the flowchart of Fig. 2, the threshold τ and the gray level ρ are assumed to be known in advance. Their values may be set manually. A good value for ρ can be determined by calculating the average value within one or more particles in a SIRT reconstruction. For τ , a value that is somewhat higher than the highest gray level in the background material generally leads to accurate reconstructions. However, although this manual procedure may lead to satisfactory results, it is also subjective.

A more objective way to determine optimal values for τ and ρ is to search for those values that result in a reconstruction that corresponds maximally to the measured tilt series. After computing the PDART reconstruction for particular values (τ, ρ) , the projection distance can be computed, resulting in a number that indicates the quality of a particular pair (τ, ρ) . By optimizing the reconstruction quality over the space of possible values for τ and ρ , their optimal values can be determined. In the case studies that follow, this optimization was performed by applying an unconstrained nonlinear optimization, using the derivation-free Nelder–Mead simplex algorithm [25]. This procedure needs in the order of 100 reconstructions to reach a precision of three significant digits for the threshold and gray level.

3. Experiments

In this section, we report on a series of experiments that were performed to assess the capabilities of the PDART algorithm. The experiments were conducted based on phantom objects, as well as experimental electron tomography datasets of two different samples.

The phantom study in Section 3.1 illustrates the basic differences between PDART and DART, and establishes that PDART can be an alternative for DART in cases where that algorithm is not applicable.

In Sections 3.2 and 3.3, PDART is applied to two different experimental datasets, to further investigate the properties of the algorithm. The reconstructions from the experimental datasets also show that PDART is applicable in practice.

The first experimental sample is a heterogeneous catalyst, consisting of metal nanoparticles on a mesoporous silica support, acquired using an angular range of -70° to $+72^\circ$. The second sample consists of Pb nano-inclusions in a crystalline Si matrix, acquired using an On-Axis Rotation Tomography Holder, allowing for image acquisition over the full angular range.

To validate the reconstruction results for both datasets, simulation phantoms that resemble the experimental sample were designed. For these phantoms, reconstructions with known ground

truth were computed, thereby allowing a direct comparison with the original object.

Throughout this section, the results for PDART are compared with SIRT and, for the phantom study, with DART. Since both PDART and DART result in reconstructions that are already segmented, the results from those algorithms are compared with segmented SIRT reconstructions. The segmentation was performed by taking a pair (τ, ρ) , and setting all pixels with a value that exceeds τ to ρ . The optimal pair (τ, ρ) was determined by minimizing the projection distance, using the optimization procedure from Section 2.3.

3.1. Phantom study

For the phantom study, two phantoms containing dense elliptical particles were designed. For the first phantom (Fig. 4a), the particles are embedded in a homogeneous material; for the second one (Fig. 4h), in a continuously varying material. The gray value of the homogeneous background from Fig. 4a is the mean of the gray values of the background from Fig. 4h. The dimensions of both phantoms are 512×512 pixels.

From the discrete phantom (Fig. 4a), a synthetic dataset using 90 projections at evenly spaced 2° intervals was created. The resulting reconstructions are shown in the top row of Fig. 4. Visually, the results for DART (Fig. 4c) and PDART (Fig. 4d) seem to be of comparable quality. In the SIRT reconstruction (Fig. 4b), the size of the particles seems to be underestimated, and reconstruction artifacts are visible in the background material. Table 1 (rows 1–4) shows the numerical results for this experiment. In Table 1, the values for τ and ρ that yield the minimum projection distance d_{pr} are shown, together with the phantom distance d_{ph} for the same pair (τ, ρ) . The numerical

results from Table 1 (rows 1–4) confirm the visual assessment, but indicate that DART adheres better to the projection data than PDART for this fully discrete dataset (the value for d_{pr} is lower), even though this is not clear from the visual appearance of the reconstruction. This is confirmed by the phantom distance d_{ph} .

From the same discrete phantom of Fig. 4a, a second synthetic dataset using 36 projections at evenly spaced 5° intervals was created. The resulting reconstructions are shown in the middle row of Fig. 4. The DART reconstruction (Fig. 4f) is virtually identical to that of Fig. 4c, even though the number of projections was reduced from 90 to 36. The quality of the PDART reconstruction (Fig. 4g) has decreased somewhat, mainly in the background material. This effect is more obvious for SIRT (Fig. 4e), as is the underestimation of the particle sizes for that algorithm.

Table 1

Optimal values for thresholds and gray levels, and the corresponding projection and phantom distances.

Phantom	Algorithm	τ	ρ	d_{pr}	d_{ph}
Discrete, 90 projections		N/A	1	0	0
	SIRT	0.772	0.928	641	32.0
	DART	0.710	1.00	130	8.54
Discrete, 36 projections		0.368	0.99	161	9.79
	SIRT	0.843	0.977	334	38.6
	DART	0.608	1.00	42.4	5.55
Partially discrete, 90 projections		0.368	1.00	103	15.8
	SIRT	N/A	1	0	0
	DART	0.766	0.926	654	31.9
Partially discrete, 90 projections		0.516	0.964	702	54.9
	PDART	0.494	1.00	266	23.0




Fig. 4. A discrete and a partially discrete phantom and several reconstructions. The first two rows show the discrete phantom and reconstructions, using 90 and 36 projections for the first and second rows, respectively. The third row shows the partially discrete phantom and reconstructions, using 90 projections. (a) Discrete phantom. (b) SIRT, (c) DART, and (d) PDART from 90 projections. (e) SIRT, (f) DART, and (g) PDART from 36 projections. (h) Partially discrete phantom. (i) SIRT, (j) DART, and (k) PDART from 90 projections.




Fig. 5. HAADF STEM projection image from the tilt series of the Catalyst sample.



Fig. 6. DART reconstruction of a slice of the Catalyst dataset in the xz -plane. The brightness of the support material was increased for clarity.

The numerical results in Table 1 (rows 5–7) confirm this, and show that the difference between DART and PDART has increased, both for d_{pr} and for d_{ph} . The results for both synthetic datasets based on the discrete phantom suggest that, for a fully discrete sample, DART should be the algorithm of choice.

From the partially discrete phantom (Fig. 4h), a third synthetic dataset using 90 projections at evenly spaced 2° intervals was created. The resulting reconstructions are shown in the last row of Fig. 4. For this third dataset, the results are markedly different. It is clear that the quality of the DART reconstruction (Fig. 4j) is quite bad. The reason that DART fails is that the model that is imposed by DART, namely that the sample is fully discrete, does not apply. The model that PDART (Fig. 4k) assumes is correct: the dense particles are discrete, and they are embedded in a continuously varying material. The visual results are confirmed by the numerical results in Table 1 (rows 8–11). Both d_{pr} and d_{ph} are now higher for DART than for SIRT (Fig. 4i), while the values for PDART are still lower, as for the preceding experiments. This implies that PDART should be the algorithm of choice for a partially discrete sample that contains homogeneous dense particles.

3.2. Sample 1: heterogeneous catalyst

The sample is a heterogeneous catalyst, consisting of metal nanoparticles on a mesoporous silica support (called Catalyst hereafter) [16].

An HAADF STEM tilt series was acquired using an FEI Tecnai F20 ST microscope operated at an acceleration voltage of 200 kV. The sample was mounted on a Fischione Model 2020 Advanced Tomography Holder. The series was recorded using 2° tilt angle increments over a range of -70° to $+72^\circ$. Fig. 5 shows a single projection image from the tilt series. The size of the projection images is 1280×1280 pixels. To increase the SNR, the projection images were downsampled by a factor of 2 in both dimensions.

In principle, the Catalyst dataset appears to be suitable for a fully discrete DART reconstruction, as it has just two compositions. However, as Fig. 6 shows, it was not possible to reconstruct the supporting




Fig. 7. Detail of the reconstruction of the catalyst dataset at different numbers of iterations. The brightness of the support material was increased for clarity. The first row shows (a) SIRT and (b) PDART reconstructions using 50 iterations. The second row shows (c) SIRT and (d) PDART reconstructions using 100 iterations. The third row shows (e) SIRT and (f) PDART reconstructions using 150 iterations.

particle as a uniform structure using DART. The fine porous structure of the supporting particle results in heavy partial volume effects, such that the particle cannot be properly represented by a constant gray

level. As a result of the mismatch within the support, the segmentation of the catalyst particles also degrades. If one would only assume discreteness of the catalyst particles, the support would be reconstructed using continuous gray levels, thereby mitigating this problem. We therefore expect that PDART will be more suitable.

The experimental dataset was reconstructed in 3D using both SIRT and PDART. Optimal parameters τ and ρ were determined from the projection data, using the procedure outlined in Section 2.3. Both algorithms were run for 100 iterations. For experimental samples, the effect of noise has to be taken into account when determining the optimal number of iterations. This is due to the well-known effect that for iterative algorithms, such as SIRT, the influence of noise in the projection data starts dominating the reconstruction after a certain number of iterations (known as semi-convergence). This effect carries over to PDART, since it uses SIRT. Fig. 7 shows that the number of iterations is less important for PDART than it is for SIRT. For PDART at 50 iterations (Fig. 7b), there is still a hint of the typical blur that surrounds the particles in the PDART reconstruction during the early stages of reconstruction (also see Fig. 3d for an extreme example of this effect). At 100 iterations (Fig. 7d), this effect has greatly diminished. The difference with the reconstruction at 150 iterations is small, although the effect of the noise starts to become apparent.

Fig. 8 shows isosurface renderings of the catalyst particles in the resulting reconstructions. From Fig. 8, it can be observed that the thresholded SIRT reconstruction is lacking many of the small particles. Fig. 9 shows a set of 2D slices through the reconstructed volume, both for a large catalyst particle (top row), and for a small one (bottom row). The same thresholds were used as for the isosurface renderings from Fig. 8. From the middle column of Fig. 9, it can be clearly observed that the threshold for SIRT, which appears suitable for the large particle, is not at all suitable for segmenting the small particle. As a consequence, no single threshold can be found for which the entire volume is segmented with reasonable accuracy. On the other hand, the PDART reconstructions for both particles, which were also computed using a single pair (τ, ρ) , do not suffer from this problem.




Fig. 8. Isosurface rendering of the (a) SIRT and (b) PDART reconstructions of the Catalyst sample. The PDART reconstruction is shown behind the SIRT reconstruction for comparison.




Fig. 9. Three orthogonal slices through a large catalyst particle (top row) and a small particle (bottom row) in the reconstruction of the Catalyst sample. The same threshold was used for the large and small particles. (a) SIRT, large, (b) segmented SIRT, large, (c) PDART, large, (d) SIRT, small (e) segmented SIRT, small and (f) PDART, small.




Fig. 10. Isosurface rendering of the Catalyst (a) phantom and reconstructions for (b) SIRT and (c) PDART. The phantom is shown behind the SIRT and PDART reconstructions for comparison.

Table 2

Optimal values for thresholds and gray levels, and the corresponding projection and phantom distances.

Sample	Algorithm	τ	ρ	d_{pr}	d_{ph}
Catalyst	SIRT	1.16	1.81	10 259	N/A
	PDART	0.403	2.03	10 018	N/A
Phantom		N/A	1	0	0
	SIRT	0.531	0.828	3311	210
	PDART	0.164	0.977	2880	149




Fig. 11. HAADF STEM projection image from the tilt series of the Pb-Si sample, showing the needle shaped structure of the sample.

To validate the results for the experimental dataset, a detailed simulation of a mesoporous support particle with catalyst particles on its surface was performed in cooperation with the Fraunhofer




Fig. 12. DART reconstruction of a slice of the Pb-Si dataset in the xz -plane. The brightness of the Si matrix material was increased for clarity.

ITWM, Germany. The dimensions of the phantom are $548 \times 548 \times 325$ pixels. From the phantom, a synthetic dataset was created, consisting of 72 projections at evenly spaced 2° intervals, producing the same 38° missing wedge as in the original tilt series of the Catalyst dataset. Poisson noise was applied to the simulated projection images. This setup resembles the conditions under which the experimental dataset was recorded. The SIRT and PDART algorithms were both run for 100 iterations. The segmented volumes for the phantom image and for both algorithms are shown in Fig. 10. In accordance with the reconstructions from the experimental data, the segmentation computed by PDART contains particles of all sizes, whereas from the segmented SIRT reconstruction a number of small particles are missing.

The numerical results for the reconstruction quality are summarized in Table 2. The projection distance d_{pr} for PDART is

lower than for the other reconstructions, which means that it corresponds more accurately to the projection data. This suggests that the value for ρ that was found by PDART is also closer to the true value than the value that was found by SIRT. For the

phantom, the numbers for d_{ph} confirm the results for d_{pr} . Moreover, the value for ρ is indeed closer to the correct value of 1.

3.3. Sample II: Pb in Si inclusions

The second sample consists of Pb nanoinclusions in a crystalline Si matrix (called Pb–Si hereafter). A micro-pillar sample was prepared by FIB milling, using an FEI Nova Nanolab 200 DualBeam system. The procedure to prepare these dedicated micro-pillar samples is explained in more detail in [26].

The sample was mounted on a Fischione Model 2050 On-Axis Rotation Tomography Holder, which allowed to acquire a series with a tilt range of $\pm 90^\circ$. An HAADF STEM tilt series was recorded using 2° tilt angle increments, using a JEOL JEM-3000F microscope operated at an acceleration voltage of 300 kV. Fig. 11 shows a single projection image from the tilt series. The size of the projection images is 512×512 pixels.

As for the Catalyst dataset, the Pb–Si dataset appears to be suitable for a fully discrete DART reconstruction, as it has just two compositions. However, Fig. 12 demonstrates that also in this case discrete tomography failed to reconstruct the matrix material as a uniform structure. Artifacts related to the inherent difficulty of aligning the projections of micro-pillar samples have as a result that the Si matrix cannot be properly represented by a constant gray level. If we assume that only the Pb particles are discrete, we can avoid degrading the accuracy of the reconstruction of those particles by the artifacts in the Si matrix. Hence, PDART is again expected to be more suitable.

Using 70 iterations for both SIRT and PDART, the experimental dataset was reconstructed in 3D. Stopping the reconstruction at




Fig. 13. Isosurface rendering of the (a) SIRT and (b) PDART reconstructions of the Pb–Si sample. The PDART reconstruction is shown behind the SIRT reconstruction for comparison.




Fig. 14. Three orthogonal slices through the reconstructions of the Pb–Si sample for (a) SIRT, (b) segmented SIRT and (c) PDART; xz-plane (squares), xy-plane (horizontal rectangles), yz-plane (vertical rectangles).




Fig. 15. Isosurface rendering of the Pb-Si (a) phantom and reconstructions for (b) SIRT and (c) PDART. The phantom is shown behind the SIRT and PDART reconstructions for comparison.

Table 3
Optimal values for thresholds and gray levels, and the corresponding projection and phantom distances.

Sample	Algorithm	τ	ρ	d_{pr}	d_{ph}
Pb-Si	SIRT	0.631	0.829	1832	N/A
	PDART	0.347	0.855	1704	N/A
Phantom	SIRT	N/A	1	0	0
	PDART	0.221	0.989	751	30.5

70 iterations avoided emphasizing the noise, while producing clearly delineated particles. As before, optimal parameters τ and ρ were determined from the projection data, using the procedure outlined in Section 2.3. Fig. 13 shows isosurface renderings of the Pb particles in the resulting reconstructions. It is apparent that the segmented SIRT reconstruction is lacking many of the small particles. This is confirmed by Fig. 14, which shows three orthogonal slices through the reconstruction.

The results for this experimental sample were validated using a 3D phantom modeled after the original Pb-Si dataset. The dimensions of the phantom are $332 \times 332 \times 73$ pixels. From the phantom, a synthetic dataset was created, consisting of 90 projections at evenly spaced 2° intervals. Poisson noise was applied to the simulated projection images. This setup resembles the conditions under which the experimental dataset was recorded. The SIRT and PDART algorithm were both run for 70 iterations. The segmented volumes for the phantom image and for both algorithms are shown in Fig. 15. As for the reconstructions from the experimental data, the smaller particles seem to be missing from the segmented SIRT reconstruction.

The numerical results for the reconstruction quality are summarized in Table 3. As for the Catalyst dataset, the projection distance for PDART is lower than for the other reconstructions, which means that it has closer correspondence to the projection data. This again suggests that the value for ρ that was found by PDART is closer to the true value than the value that was found by SIRT.

4. Discussion

The case studies in this paper show two substantially different experimental datasets that could not be reconstructed successfully using discrete tomography. Since they both contain dense particles as their main composition of interest, they are suitable for application of the PDART algorithm. The algorithm is straightforward to implement, which makes it applicable in practice. The results from PDART seem to improve on the segmentation of SIRT reconstructions that were created using a threshold that was optimized to match the projection data.

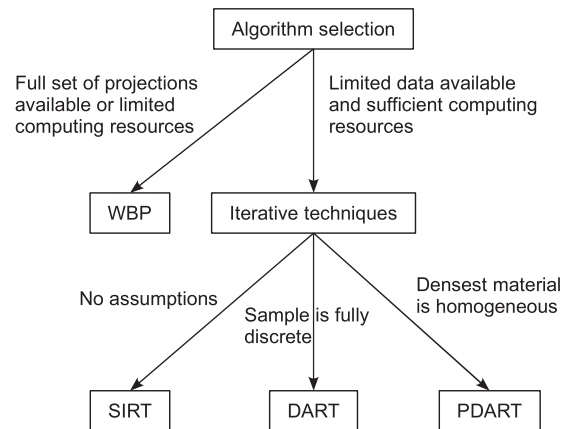


Fig. 16. Algorithm selection tree.

In both studies, the particles are accurately segmented, regardless of their size. The reconstructions of the different samples and phantoms also show that the parameters of the method can be optimized in an objective manner. This makes the algorithm an alternative to manual segmentation, since, even if a comparable segmentation is created manually, there will always be subjective judgment involved.

The results for the experimental datasets, together with the phantom study, suggest that PDART is a useful alternative to DART, since the algorithm allows samples for which a fully discrete reconstruction is not possible to still benefit from the techniques of discrete tomography. The four algorithms that were applied in this paper are shown in an algorithm selection tree in Fig. 16. If a full set of projections is available, or if computing resources are limited, WBP can still be the algorithm of choice. However, if sufficient computing resources are available, iterative techniques have compelling advantages, like a lower sensitivity to noise. Another important advantage is that they allow to exploit prior knowledge, which leads to (partially) discrete algorithms. If a sample is fully discrete, DART is the best option. If only the densest material is homogeneous, PDART retains a number of the advantages of discrete tomography, while providing a SIRT reconstruction of the background. If no assumptions can be made, SIRT is the most generally applicable algorithm, which is also widely available.

5. Conclusions

We have presented a novel reconstruction algorithm for partially discrete tomography. The new method has two main benefits.

First, a partially discrete technique is useful to expand the set of samples for which concepts of discrete tomography can be applied. The method presented in this paper is a practical one in this regard, since it is applicable to the common problem of dense particle segmentation. The method retains two advantages of discrete tomography: the densest composition is automatically segmented, and it can produce reconstructions that are more accurate than SIRT.

Second, the method automatically determines the gray level that should be used for the densest composition in an objective manner, by locating the value that makes the resulting reconstruction closest to the original projection images. This property of the algorithm should make it easier to draw quantitative conclusions from the reconstructions.

Acknowledgments

The authors would like to thank M. Van Bael and H. Wang from the K.U. Leuven (Belgium) for providing the Pb–Si sample. They would also like to thank R. Cieslinski and S. Rozeveld from The Dow Chemical Company (USA) for providing the Catalyst sample. The authors also gratefully acknowledge the assistance of K. Schladitz and M. Godehardt from the Fraunhofer ITWM (Germany) with the analysis of the Catalyst sample. This work was financially supported by the Research Foundation (FWO), Flanders, Belgium (Contract no. G.0247.08). KJB acknowledges the Netherlands Organisation for Scientific Research (NWO) for support (research programme 639.072.005).

References

- [1] P.A. Midgley, R.E. Dunin-Borkowski, Electron tomography and holography in materials science, *Nature Materials* 8 (2009) 271–280.
- [2] R.J. Spontak, M.C. Williams, D.A. Agard, Three-dimensional study of cylindrical morphology in a styrene–butadiene–styrene block copolymer, *Polymer* 29 (3) (1988) 387–395.
- [3] A.J. Koster, U. Ziese, A.J. Verkleij, A.H. Janssen, K.P. de Jong, Three-dimensional transmission electron microscopy: a novel imaging and characterization technique with nanometer scale resolution for materials science, *Journal of Physical Chemistry B* 104 (40) (2000) 9368–9370.
- [4] S. Bals, G. Van Tendeloo, C. Kisielowski, A new approach for electron tomography: annular dark-field transmission electron microscopy, *Advanced Materials* 18 (7) (2006) 892–895.
- [5] P.A. Midgley, M. Weyland, J.M. Thomas, B.F.G. Johnson, Z-contrast tomography: a technique in three-dimensional nanostructural analysis based on Rutherford scattering, *Chemical Communications* 10 (2001) 907–908.
- [6] E.P.W. Ward, T.J.V. Yates, J.-J. Fernández, D.E.W. Vaughan, P.A. Midgley, Three-dimensional nanoparticle distribution and local curvature of heterogeneous catalysts revealed by electron tomography, *Journal of Physical Chemistry C* 111 (31) (2007) 11501–11505.
- [7] M. Weyland, T.J.V. Yates, R.E. Dunin-Borkowski, L. Laffont, P.A. Midgley, Nanoscale analysis of three-dimensional structures by electron tomography, *Scripta Materialia* 55 (1) (2006) 29–33.
- [8] T.J.V. Yates, J.M. Thomas, J.-J. Fernández, O. Terasaki, R. Ryoo, P.A. Midgley, Three-dimensional real-space crystallography of MCM-48 mesoporous silica revealed by scanning transmission electron tomography, *Chemical Physics Letters* 418 (4–6) (2006) 540–543.
- [9] K. Kaneko, K. Inoke, K. Sato, K. Kitawaki, H. Higashida, I. Arslan, P.A. Midgley, TEM characterization of Ge precipitates in an Al–1.6at% Ge alloy, *Ultramicroscopy* 108 (3) (2008) 210–220.
- [10] G. Möbus, B.J. Inkson, Three-dimensional reconstruction of buried nanoparticles by element-sensitive tomography based on inelastically scattered electrons, *Applied Physics Letters* 79 (9) (2001) 1369–1371.
- [11] M. Weyland, P.A. Midgley, Extending energy-filtered transmission electron microscopy (EFTEM) into three dimensions using electron tomography, *Microscopy and Microanalysis* 9 (6) (2003) 542–555.
- [12] M.H. Gass, K.K.K. Koziol, A.H. Windle, P.A. Midgley, Four-dimensional spectral tomography of carbonaceous nanocomposites, *Nano Letters* 6 (3) (2006) 376–379.
- [13] B. Goris, S. Bals, W. Van den Broek, J. Verbeeck, G. Van Tendeloo, Exploring different inelastic projection mechanisms for electron tomography, *Ultramicroscopy* 111 (8) (2011) 1262–1267.
- [14] L.C. Gontard, R.E. Dunin-Borkowski, D. Ozkaya, Three-dimensional shapes and spatial distributions of Pt and PtCr catalyst nanoparticles on carbon black, *Journal of Microscopy* 232 (2) (2008) 248–259.
- [15] P. Gilbert, Iterative methods for the three-dimensional reconstruction of an object from projections, *Journal of Theoretical Biology* 36 (1) (1972) 105–117.
- [16] C. Kübel, D. Niemeyer, R. Cieslinski, S. Rozeveld, Electron tomography of nanostructured materials—towards a quantitative 3D analysis with nanometer resolution, *Materials Science Forum* 638–642 (2010) 2517–2522.
- [17] S. Van Aert, K.J. Batenburg, M.D. Rossell, R. Erni, G. Van Tendeloo, Three-dimensional atomic imaging of crystalline nanoparticles, *Nature* 470 (2011) 374–377.
- [18] S. Bals, M. Casavola, M.A. van Huis, S. Van Aert, K.J. Batenburg, G. Van Tendeloo, D. Vanmaekelbergh, Three-dimensional atomic imaging of colloidal core–shell nanocrystals, *Nano Letters* 11 (8) (2011) 3420–3424.
- [19] K.J. Batenburg, J. Sijbers, DART: a practical reconstruction algorithm for discrete tomography, *IEEE Transactions on Image Processing* 20 (9) (2011) 2542–2553.
- [20] K.J. Batenburg, S. Bals, J. Sijbers, C. Kübel, P.A. Midgley, J.C. Hernandez, U. Kaiser, E.R. Encina, E.A. Coronado, G. Van Tendeloo, 3D imaging of nanomaterials by discrete tomography, *Ultramicroscopy* 109 (6) (2009) 730–740.
- [21] S. Bals, K.J. Batenburg, J. Verbeeck, J. Sijbers, G. Van Tendeloo, Quantitative three-dimensional reconstruction of catalyst particles for bamboo-like carbon nanotubes, *Nano Letters* 7 (12) (2007) 3669–3674.
- [22] S. Bals, K.J. Batenburg, D. Liang, O. Lebedev, G. Van Tendeloo, A. Aerts, J.A. Martens, C.E.A. Kirschhock, Quantitative three-dimensional modeling of zeolite through discrete electron tomography, *Journal of the American Chemical Society* 131 (13) (2009) 4769–4773.
- [23] K.J. Batenburg, J. Sijbers, H.F. Poulsen, E. Knudsen, DART: a robust algorithm for fast reconstruction of three-dimensional grain maps, *Journal of Applied Crystallography* 43 (6) (2010) 1464–1473.
- [24] J. Gregor, T. Benson, Computational analysis and improvement of SIRT, *IEEE Transactions on Medical Imaging* 27 (7) (2008) 918–924.
- [25] J.C. Lagarias, J.A. Reeds, M.H. Wright, P.E. Wright, Convergence properties of the Nelder–Mead simplex method in low dimensions, *SIAM Journal on Optimization* 9 (1) (1998) 112–147.
- [26] X. Ke, S. Bals, A.R. Negreira, T. Hantschel, H. Bender, G. Van Tendeloo, TEM sample preparation by FIB for carbon nanotube interconnects, *Ultramicroscopy* 109 (11) (2009) 1353–1359.