Super-resolution for computed tomography based on discrete tomography

Wim van Aarle, K. Joost Batenburg, G. Van Gompel, E. Van de Casteele and Jan Sijbers

***Abstract*—In Computed Tomography (CT), partial volume effects impede accurate segmentation of structures that are small with respect to the pixel size. In this paper, it is shown that for objects consisting of a small number of homogeneous materials, the reconstruction resolution can be substantially increased without altering the acquisition process. A super-resolution reconstruction approach is introduced that is based on discrete tomography, in which prior knowledge about the materials in the object is assumed. Discrete tomography has already been used to create reconstructions from a low number of projection angles, but in this paper, it is demonstrated that it can also be applied to increase the reconstruction resolution. Experiments on simulated and real μCT data of bone and foam structures show that the proposed method indeed leads to significantly improved structure segmentation and quantification compared to what can be achieved from conventional reconstructions.**

***Index Terms*—computed tomography, segmentation, super-resolution, discrete tomography**

I. Introduction

In *X-ray Computed Tomography (CT)*, images are typically reconstructed on a voxel grid. Since each voxel is represented by a constant grey level, it is intrinsically assumed that the material within such a voxel is homogeneous. It is clear, however, that a voxel representation cannot properly represent structures that have a varying density within a voxel. Thus, each voxel in the images could contain more than one material or tissue type. This phenomenon is referred to as the *partial volume effect (PVE)*. PVEs will cause object boundaries to be smeared out across the boundary voxels. Also, if a feature of the scanned object is small relative to the nominal voxel size, PVEs reduce the contrast between the structure of interest and its background signal. Consequently, it is difficult to achieve the intrinsic resolution of the detector. Fig. 1a shows a *filtered backprojection (FBP)* reconstruction of a polyurethane foam for which the widths of the edges of the pores are comparable to the detector size. A globally thresholded segmentation of Fig. 1a, created with the commonly used clustering method of Otsu [1], is shown in Fig. 1b. Clearly, many thin structures remain undetected, whereas the thickness for some larger structures is overestimated.

To reduce PVEs, and hence to obtain sufficient contrast, a high resolution scan can be acquired. This, however, requires a much higher radiation dose and a longer scanning time [2]. In Fig. 2, an FBP reconstruction with a spatial resolution of 35μm is shown of a rat femur along with an FBP reconstruction with a spatial resolution of 9μm of the same femur. It is clear that the contrast in Fig. 2b is significantly better than that in Fig. 2a. Fig. 2b is therefore better suited for accurate segmentation and estimation of the of the morphometric bone parameters [3], which is crucial to understand the effects of drug trials in for example osteoporosis research [4]. Other applications where high resolution images are required include metrology (e.g. to inspect the surface of the blades of aerofoil turbines), structural biology and materials research [5].

The conventional approach to reduce PVEs without increasing the radiation dose is to upsample the reconstruction voxel grid, allowing for a more accurate representation and potentially improving the overall visualisation of small structures. This upsampling is also known as *super-resolution* [6]. It is important to note, however, that in CT, a unique reconstructed image can only be obtained if the projection domain is adequately sampled. The required amount of information is dependent on the number of voxels. Upsampling the reconstruction grid therefore typically leads to a limited data reconstruction problem: the number of ray-equations (measured projection data) remains the same while the number of unknowns (reconstruction voxels) increases significantly. To overcome this problem, additional information must be entered into the reconstruction problem. This can be done in numerous ways. In [7], [8], [9], information from multiple low resolution CT images is combined into a high resolution CT image [7], [8], [9], but these methods result in an increased scan time and radiation dose. In [10], additional detector samples are created by Fourier interpolation and a compressed sensing solution is used to solve a reconstruction with many projection images and high noise intensity.

In this paper, a super-resolution approach for CT is proposed that effectively solves the limited data problem by incorporating prior knowledge about the unknown object. In CT, such prior knowledge comes in many forms, e.g. sparsity of the reconstructed image [11] or its gradient [12], [13]. Here, the novel super-resolution scheme is based on the *Discrete Algebraic Reconstruction Technique (DART)*, an iterative reconstruction technique that can be applied if the scanned object is known to consist of a small set of materials, each corresponding to a different constant and priorly known grey level in the reconstruction [14]. Similar to [14], the focus in this work is on the reconstruction of images that consist of a small number of grey levels, typically up to 4 or 5. It should be noted that in practice, due to the polychromatic nature of the X-ray spectrum, there is no quantitative model

Wim van Aarle, K. Joost Batenburg, E. Van de Casteele and Jan Sijbers are affiliated with iMinds-Vision Lab, University of Antwerp, Antwerp, Belgium; K. Joost. Batenburg is also affiliated with Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands. G. Van Gompel is affiliated with the University Hospital of Brussels, Belgium.

Corresponding author: Wim van Aarle, E-mail: wim.vanaarle@ua.ac.be.

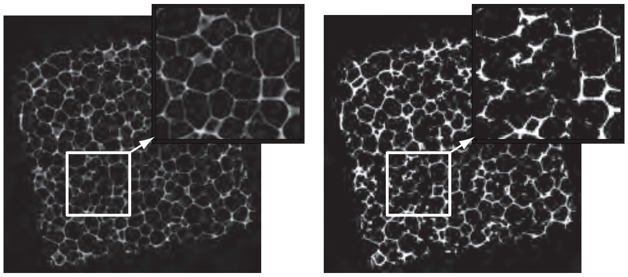

(a) reconstruction (b) segmentation

Fig. 1: (a) Reconstruction of a polyurethane foam, taken with a SkyScan 1172 μCT scanner at a pixel resolution of 17μm. (b) Otsu's segmentation of the reconstruction. Many cell walls remain undetected in the segmentation while other structures are overestimated.

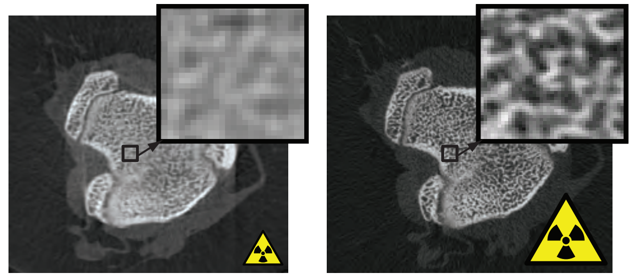

(a) 35μm reconstruction, low radiation dose (b) 9μm reconstruction, high radiation dose

Fig. 2: FBP reconstructions of the epiphyseal plate of a rat femur taken at two different resolutions in a SkyScan 1172 μCT scanner. The high dose reconstruction (b) is clearly much easier to segment. Note that as both slices were taken from different scans, the object was slightly displaced between the acquisition of both datasets. Even though image registration was performed, there is still a residual difference.

to exactly determine the grey levels based on the materials of the scanned object. However, it is typically possible to obtain a sufficiently accurate estimation by looking at a reconstructed image computed with a standard iterative technique such as SIRT [15]. Alternatively, in [16], [17], the authors provide methods to (semi-)automatically estimate grey level values for use in discrete tomography.

It has been shown already that by utilizing discrete tomography techniques, very accurate reconstructions can often be computed from only a few projection images [18], [19], [20]. In this paper, it will be demonstrated that discrete tomography can also be used for benefits in a different direction, namely to increase the resolution of the reconstructed images with the same (or only slightly less) number of projection angles. It will be shown that by upsampling the reconstruction grid and incorporating prior knowledge about the objects grey levels, the lack of high resolution projection data can be compensated. The proposed approach effectively increases the spatial resolution of the tomographic reconstructions [21], [22].

The paper is organized as follows. Section II introduces notation for algebraic and discrete tomography. Section III introduces the new super-resolution approach. In Section IV, experiments are described that were performed to evaluate the reconstruction accuracy for the proposed super-resolution approach. Results are presented for both simulated data and experimental μCT data. Finally, Section V concludes this work.

II. Concepts and Notation

In this section, general concepts and notations are introduced. In Section II-A, the algebraic tomography model is described. For simplicity, a monochromatic x-ray beam will be assumed. Note, however, that this does not preclude application of the method to polychromatic x-ray imaging since preprocessing methods can be applied to compute monochromatic from polychromatic projections [23], [24], [25], [26]. Section II-B discusses algebraic reconstruction techniques such as SIRT. Section II-C concerns discrete reconstruction techniques such as DART, an iterative reconstruction technique that exploits prior knowledge about the grey levels of each of the scanned materials [14]. DART effectively combines reconstruction and segmentation into a single tomographic algorithm. For clarity, all concepts will be presented on a 2D parallel beam projection geometry. However, the proposed methods can be generalized to any acquisition geometry.

A. Computed tomography

Let f represent the 2D attenuation of a certain object, which will be referred to as the *object function*. A parallel beam projection geometry defines the tomographic projection of f as the line integrals of f along the lines $l_{\theta,t} = \{(x,y) \in \mathbb{R}\times\mathbb{R} : x\cos\theta + y\sin\theta = t\}$, where $\theta \in [0,\pi)$ represents the angle between the line and the y-axis and where $t \in \mathbb{R}$ represents the coordinate along the projection axis. For a finite set of lines $l_{\theta,t}$, the X-ray beam intensity at the detectors, $I(\theta,t)$, are measured as

$$I(\theta,t) = I_0 e^{-\int_{l_{\theta,t}} f(x,y)ds}, \tag{1}$$

with I_0 the incident beam intensity. Define the *attenuation projection function* $p(\theta,t)$ as follows:

$$p(\theta,t) = -\ln\left(\frac{I(\theta,t)}{I_0}\right) = \int_{l_{\theta,t}} f(x,y)ds \quad , \tag{2}$$

also called the *forward projection* or *sinogram* of $f(x,y)$.

In practice, a projection is measured at a set of projection angles and at a set of detector elements with a width Δt. Let $\boldsymbol{I} \in \mathbb{R}^m$ denote the measured intensity data, with m the number of detector values multiplied by the number of projection angles. For $j \in 1,\ldots,m$, I_j can then be modelled as

$$I_j = \int_{-\frac{\Delta t}{2}}^{\frac{\Delta t}{2}} I_0 e^{-p(\theta,t+t')}dt', \tag{3}$$

with t and θ the detector coordinate and projection angle of the measured detector value I_j, respectively. The attenuation projection data $\boldsymbol{p} \in \mathbb{R}^m$ can then be defined as follows:

$$p_j = -\ln\left(\frac{I_j}{I_0}\right) \quad . \tag{4}$$

Note that, due to the logarithmic operation in Eq. (4), the contribution of a pixel to the measured projection values does not only depend on the average value of that pixel, but also on the *distribution* of the attenuation *within that pixel*.

Tomography deals with the reconstruction of $f(x, y)$ based on $\boldsymbol{p}$. This reconstructed function is represented by an image, a grid of square pixels with a finite width and height, Δs. Let $\boldsymbol{v} \in \mathbb{R}^n$ denote a discretized square image of the function $f(x, y)$, where n denotes the number of pixels. v_i can then be modelled as the total value of f, taken over the square pixel:

$$v_i = \int_{-\frac{\Delta s}{2}}^{\frac{\Delta s}{2}} \int_{-\frac{\Delta s}{2}}^{\frac{\Delta s}{2}} f(x_i + x', y_i + y')dx'dy', \tag{5}$$

with x_i and y_i the coordinates of the center point of pixel v_i.

The value of a certain pixel v_i will thus depend on an entire area of values of the real object function. If the object has an edge running through the area of pixel v_i, or if the object is not homogeneous inside the pixel boundaries, the value of v_i will not represent the attenuation coefficient of any of the materials of the object, but will represent an average of all attenuation coefficients. This is called the *partial volume effect (PVE)*. Note that for object functions that consist of piece-wise constant regions, the fraction of pixels for which PVEs occur is directly related to the size of Δs.

B. Algebraic tomography model

Using the discretized definitions of projection data (Eq. (4)) and reconstructed image (Eq. (5)), a computational model — approximating the mathematical projection model — can be constructed. The forward projection of the object for a finite set of angles is modelled as a linear operator $\boldsymbol{W}$, called the *projection operator*, which maps the image $\boldsymbol{v}$ to the projection data $\boldsymbol{q}$:

$$\boldsymbol{q} := \boldsymbol{W}\boldsymbol{v}. \tag{6}$$

In Eq. (6), $\boldsymbol{W} = (w_{ij})$ is an $m \times n$ matrix where w_{ij} represents the contribution of image pixel v_j to detector value q_i. The vector $\boldsymbol{q}$ is called the *forward projection* of $\boldsymbol{v}$. The reconstruction problem in CT can then be modelled as the recovery of $\boldsymbol{v}$ from a given vector $\boldsymbol{p}$ of projection data, such that:

$$\boldsymbol{W}\boldsymbol{v} = \boldsymbol{p}. \tag{7}$$

Many reconstruction algorithms have been proposed to solve Eq. (7) without any constraints on $\boldsymbol{v}$ [27]. One of these methods is the *Simultaneous Iterative Reconstruction Technique (SIRT)*, which will be used, as defined in [15], throughout this paper. SIRT is a linear algorithm that finds a solution $\tilde{\boldsymbol{v}}$ such that the weighted squared projection difference $||\boldsymbol{W}\tilde{\boldsymbol{v}} - \boldsymbol{p}||_R = (\boldsymbol{W}\tilde{\boldsymbol{v}} - \boldsymbol{p})^T \boldsymbol{R}(\boldsymbol{W}\tilde{\boldsymbol{v}} - \boldsymbol{p})$ is minimal. $\boldsymbol{R} \in \mathbb{R}^{m \times m}$ is a diagonal matrix that contains the inverse row sums of $\boldsymbol{W}$: $r_{ii} = 1/\sum_j w_{ij}$. In each iteration k, the current reconstruction $\boldsymbol{v}^{(k-1)}$ is updated, yielding a new reconstruction $\boldsymbol{v}^{(k)}$, as follows:

$$v_j^{(k)} = v_j^{(k-1)} + \lambda \frac{1}{\sum_{i=1}^n w_{ij}} \sum_{i=1}^m \frac{w_{ij}\left(p_i - w_{ij} v_i^{(k-1)}\right)}{\sum_{j=1}^m w_{ij}}. \tag{8}$$

In Eq. (8), λ is a relaxation parameter.

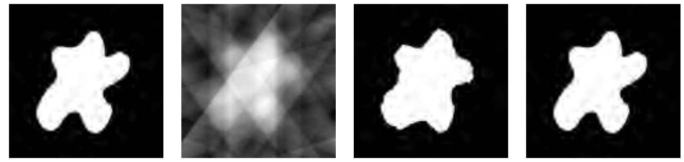

(a) Phantom (b) SIRT (c) S-SIRT (d) DART

Fig. 3: (a) 256×256 phantom image. (b) SIRT reconstruction using 5 equiangular projections. (c) Segmentation of (b) using Otsu's method [1] (S-SIRT). (d) DART reconstruction using 5 equiangular projections.

C. Discrete tomography

In some reconstruction problems m is much smaller than n (e.g. when the number of projection directions is very low or the data is truncated), which leads to an underdetermined system of linear equations: so-called *limited data problems*. Fig. 3b shows a SIRT reconstructed image of the phantom image in Fig. 3a from only 5 equiangular projections. The segmentation of Fig. 3b using Otsu's method [1], which will be referred to by S-SIRT in the remainder of this paper, is shown in Fig. 3c. Note that mainly the pixels near the border or the object are incorrectly segmented.

The observation is used by the *Discrete Algebraic Reconstruction Technique (DART)* to reduce the size of Eq. (7). DART uses prior knowledge about the discrete grey levels to iteratively solve Eq. (7) under the constraint that v_i can only take values that are elements of a set $\boldsymbol{\rho} = \mathbb{R}^l$. Each element of $\boldsymbol{\rho}$ contains the grey level value of one of the l different materials of the scanned object. It is chosen by the user based on the available prior knowledge.

The DART algorithm has shown great potential [5], [28], [29]. Here, a concise summary of the algorithmic steps is given. For any $A \subset \{1, ..., n\}$, and any $\bar{\boldsymbol{v}} \in \mathbb{R}^n$, let $\bar{\boldsymbol{v}}_A \in \mathbb{R}^{n_A}$ be a vector that contains a subset of the entries of $\bar{\boldsymbol{v}}$, where $\bar{v}_i$ is included iff $i \in A$. Furthermore, let $\boldsymbol{W}_A \in \mathbb{R}^{m \times n_A}$ be the matrix that contains the columns $i \in A$ of the matrix $\boldsymbol{W}$. The DART algorithm consists of the following steps:

1) Create an initial reconstruction $\boldsymbol{v}^{(0)}$ using SIRT. Put $k = 0$, the iteration number.
2) If $k > 0$, apply a smoothing filter to $\boldsymbol{v}^{(k)}$. This can be done by application of a convolution with the 2D stencil $\frac{b}{8}[1\ 1\ 1;\ 1\ (1-b)\ 1;\ 1\ 1\ 1]$, where b is the intensity of the smoothing, which is typically chosen at $b = 0.20$. This smoothing step is required because if pixels can very independently of each other, great variations near the border tend to occur. Blurring then regularizes the data.

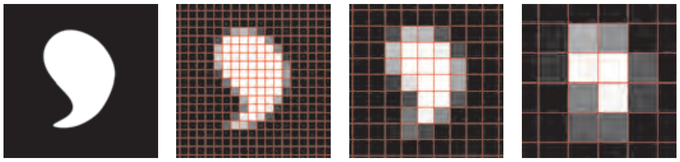

Fig. 4: (a) A binary object. (b-d) The same binary object, represented on a pixel grid with increasing pixel sizes.

3) Segment $\boldsymbol{v}^{(k)}$. A simple scheme with fixed global thresholds, $\boldsymbol{\tau} \in \mathbb{R}^{l-1}$, is used to replace the grey level of each pixel $v_i^{(k)}$ by that of the corresponding value of $\boldsymbol{\rho}$. The values $\boldsymbol{\tau}$ are typically chosen exactly in the middle of two grey levels.
4) Determine $A \subset \{1, \dots, n\}$, the set of border pixels. A pixel is defined as a border pixel if its value is different from any of its neighbours (defined by an 8-connectivity window). To allow non-border pixels to be updated as well and to further reduce the impact of noisy projection data, a small number of additional random pixels are added to A (typically about 10%).
5) Compute $\boldsymbol{p}^{(k+1)}$, the residual projection data, by subtracting the forward projection of all pixels $v_i^{(k)}$ with $i \notin A$ from the measured data $\boldsymbol{p}$.
6) Create the reconstruction $\boldsymbol{v}^{(k+1)}$ using SIRT by solving $\boldsymbol{W}_A \boldsymbol{v}_A^{(k+1)} = \boldsymbol{p}^{(k+1)}$. This system of equations has a much smaller number of unknowns than the original system and is therefore better determined, even when few projection angles are available.
7) Increase k by 1 and return to step 2 until some termination criterion has been reached.

For a more in depth description of DART, we refer to the original publication [14]. Fig. 3d shows a DART reconstruction of the phantom image in Fig. 3a from only 5 projections.

Note that, from its design, DART is especially suited for structures that are large with respect to Δs. If the object to be reconstructed consists of many small structures, such as foams or trabecular bone, two effects limit the possible improvements of DART over standard techniques. For one, the PVE breaks the assumption that the number of grey levels is small. Even binary objects can then no longer be accurately depicted on a grid with only black or white pixel values (Fig. 4). Also, for small objects, the number of elements in A will still be large, thereby insufficiently reducing the reconstruction problem size.

III. Super-resolution

To achieve the intrinsic detector resolution and to counter the PVE, the reconstruction grid must be upsampled (Fig. 5b). Let a be the upsampling factor in each dimension. Each pixel of width Δs is then subdivided into a^2 pixels of width $\frac{\Delta s}{a}$. Denote the upsampled reconstruction image by $\boldsymbol{v}' \in \mathbb{R}^{a^2 n}$.

Note that, typically, $\frac{\Delta s}{a}$ is different from Δt, the width of the detector cell. If the projection weights w_{ij} are computed by intersection of a single ray with the upsampled image, some pixels will not have a ray going through them for each projection angle and the projection data will not be computed correctly. Two methods are investigated to overcome this problem: sinogram upsampling (Fig. 5c, Section III-A) and detector supersampling (Fig. 5d, Section III-B).

Define the *relative reconstruction resolution* as the ratio of the detector width, Δt, to the pixel size, Δs:

$$R_v = \frac{\Delta t}{\Delta s}. \tag{9}$$

A. Sinogram upsampling

With *sinogram upsampling (SU)*, the number of detector cells is artificially increased by subdividing each detector of size Δt into a detectors of size $\frac{\Delta t}{a}$. Fig. 5c shows a schematic overview of this geometry. The value of each detector point is determined by linear interpolation of $\boldsymbol{p}$. Let $\boldsymbol{p'} \in \mathbb{R}^{am}$ be the upsampled sinogram and let $\boldsymbol{W}_{SU} \in \mathbb{R}^{am \times a^2 n}$ be the corresponding projection operator. The reconstruction equation then becomes:

$$\boldsymbol{W}_{SU} \boldsymbol{v'} = \boldsymbol{p'}. \tag{10}$$

Note that, the relative reconstruction resolution, $R_{v'}$, has remained the same. Furthermore, when interpolating the projection data, a certain smoothness in the projection data is assumed.

B. Detector supersampling

With *detector supersampling (DS)*, the sinogram $\boldsymbol{p}$ remains unaltered. However, the number of virtual rays targeting each detector cell is increased by a factor a, each $\frac{\Delta t}{a}$ apart. The relative reconstruction resolution increases by the same factor. Fig. 5d shows a schematic overview of this geometry. The reconstruction equation is:

$$\boldsymbol{W}_{DS} \boldsymbol{v'} = \boldsymbol{p}, \tag{11}$$

where each row in the projection operator $\boldsymbol{W}_{DS} \in \mathbb{R}^{m \times a^2 n}$ is the summation of the a corresponding rows of $\boldsymbol{W}_{SU} \boldsymbol{v'}$.

In Eq. (10) and Eq. (11) the number of unknowns has been increased by a factor a^2 while the number of equations has been increased by a factor a (Eq. (10)) and remained unaltered (Eq. (11)), respectively. Solving the reconstruction equation is now a limited data problem. As was noted in Section II, prior knowledge about the scanned objects can be used to solve Eq. (10) and Eq. (11) with the DART algorithm.

Note that there is a non-linear relationship between the measured projection data $\boldsymbol{p}$ and the actual attenuation projection data $p(\theta, t)$ (Eq. (3) and Eq. (4)). As DART uses a linear projection model, the proposed super-resolution approaches do not accurately model the PVE. In the next section, however, it will be experimentally demonstrated that even with this limited model, super-resolution on piecewise homogeneous objects with known attenuation coefficients can indeed be achieved, leading to significant improvements in reconstruction accuracy.

IV. Experiments

In this section, the proposed super-resolution method is demonstrated and its effectiveness is evaluated on various simulated images (Section IV-A) and on real datasets (Section IV-B).

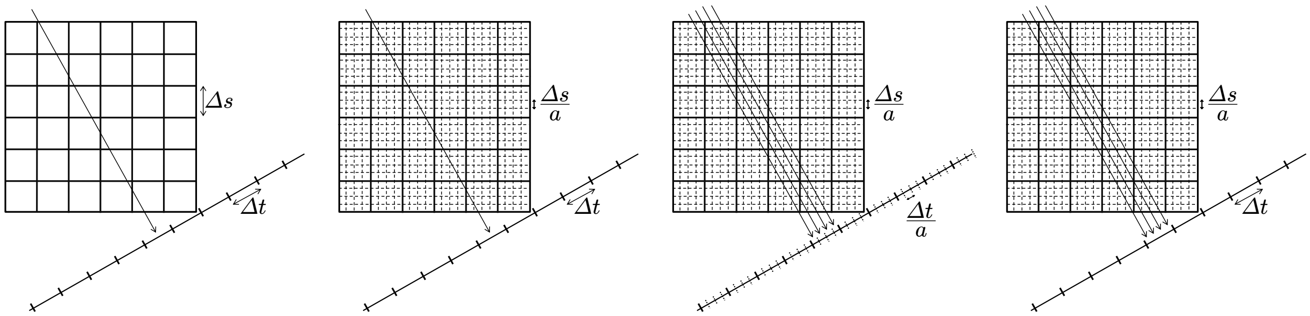

(a) Basic geometry (b) Upsampled grid (c) Sinogram upsampling (SU) (d) Detector supersampling (DS)

Fig. 5: (a) Basic projection geometry. Each detector cell corresponds to a single ray. (b) Upsampled reconstruction grid. Certain pixels are not hit by a ray. (c) Sinogram upsampling. Each detector is subdivided into multiple detectors with interpolated values (d) Detector supersampling. Each detector corresponds to multiple rays. The contribution of each ray is summed.

A. Simulation experiments

Experiments were performed on five simulation phantoms (Fig. 6a and Fig. 8). Some phantoms were generated analytically (Fig. 6a and Fig. 8c), while others were generated based on high resolution rasterized images.

Analytical rings phantom 1: In the first experiment, the efficacy of a discrete super-resolution technique was examined as a function of the size of a structure with respect to Δt. To this end, a simulated analytical binary phantom containing 11 rings with a varying width, Δq, was created. A rasterized rendering of this phantom is depicted in Fig. 6(a). As a measure of the magnitude of the PVE, the notion of *relative projection resolution*, R_p, is introduced. It is defined as the ratio of the object width, i.e the thickness of the ring, to the detector width:

$$R_p = \frac{\Delta q}{\Delta t} \tag{12}$$

For the phantom in Fig. 6a, the R_p of the outer three rings is 10, 5 and 3. The R_p of the fourth ring is 1 and can thus be used to measure if the intrinsic detector resolution is achieved. The seven most inner rings have an R_p of $\frac{1}{2}$ to $\frac{1}{8}$.

Projection data was analytically generated (using Eq. (4)) for a parallel beam geometry with 60 equiangular projection angles and 256 detector pixels. Reconstructions were computed for both S-SIRT and DART and with both the sinogram upsampling approach and the detector supersampling approach, with increasing levels of super-resolution: a = 1, 2, 4 and 10.

From Fig. 6b-d, it is clear that by increasing a (combined with DS and DART), the spatial resolution improves, as indicated by the appearance of the rings in the center. This effect is less pronounced if SU is used (Fig. 6f). Furthermore, the thin rings can not be seen at all if no prior knowledge is included in the reconstruction (Fig. 6e and Fig. 6g). These results can also be observed in Fig. 6i-l, where for each a the *relative Number of Misclassified Pixels (rNMP)* of each ring is plotted. The rNMP measures the total number of pixels that are classified in a wrong partition (false negatives as well as false positives) with respect to the total number of pixels of that object. For analytical phantom images, this rNMP value was approximated by comparing the reconstruction to a very high resolution rasterization of the phantom image. For this experiment, the rNMP was computed for each ring separately. The false negatives of each ring can be easily counted, but counting false positive pixels is more difficult as it is not clear to which ring such pixel belongs. In the results shown in Fig. 6, each false positive pixel is accounted to the ring that it is closest to. Additionally, all experiments in this paper have also been evaluated with the Root Mean Square Error and the Structural Similarity validation metrics. The results from these metrics provide the same insights as the rNMP. For the sake of brevity, only rNMP is therefore covered here.

Analytical rings phantom 2: In a second experiment, it is explored how the number of required projection angles is related to the level of volume upsampling, used in combination with the DART method. Projection data was analytically simulated of a single ring, with varying thickness $\Delta q \in [\frac{1}{16}, 4]$. The number of projection directions was varied from 2 to 45 and data was generated for 64 detectors of width $\Delta t = 1$ (hence, $R_p = q$ in this experiment). In Fig. 7d, a ring of size $q = \frac{1}{2}$ is visualized, along with its projection data.

DART reconstructions were performed with various levels of DS. These reconstructions were validated with the rNMP metric, measured against a high resolution rasterization of the phantom. In Fig. 7a-c, the resulting rNMP values are plotted as a function of the number of projection angles for different rings. From these plots, it is clear that for low angle counts, the reconstructions greatly benefit from additional angles. However, after a certain point, ”enough” information is available and additional angles offer no improved accuracy any more. Obviously, this ”point of sufficient information” depends on the complexity of the structures to be reconstructed (in this experiment this can be clearly seen that thicker rings require much less projections than thinner rings). Besides that, however, also the level of upsampling plays a role. As larger levels of upsampling are applied, the reconstruction grid becomes finer and finer and smaller rings can be reconstructed with improved accuracy, which requires additional information, i.e. projections.

In Fig. 7e, the minimal R_p that can be reconstructed with

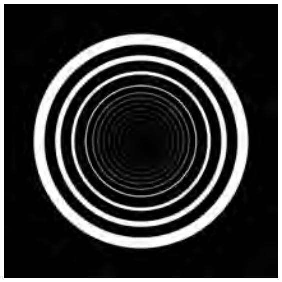

(a) Rasterized ground truth image

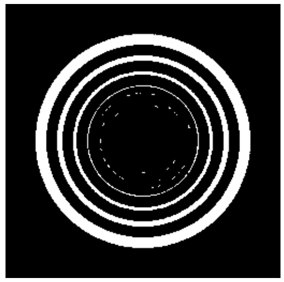

(b) DS, $a = 1$ DART

(c) DS, $a = 2$ DART

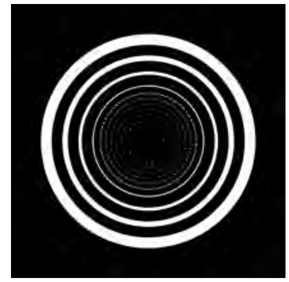

(d) DS, $a = 4$ DART

(e) SU, $a = 10$ S-SIRT

(f) SU, $a = 10$ DART

(g) DS, $a = 10$ S-SIRT

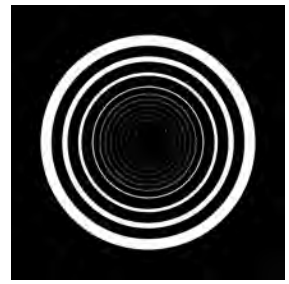

(h) DS, $a = 10$ DART

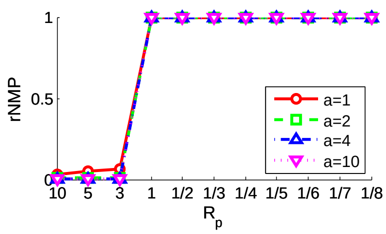

(i) SU, S-SIRT

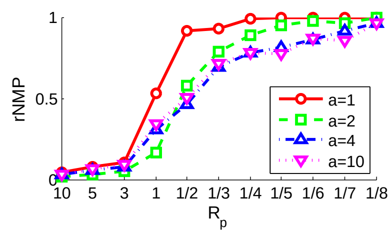

(j) SU, DART

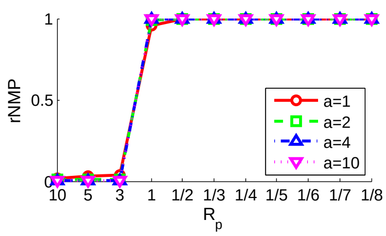

(k) DS, S-SIRT

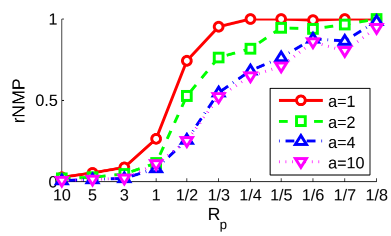

(l) DS, DART

Fig. 6: Experimental results from a simulated analytical phantom containing 11 rings of varying width. (a) High resolution rendering of the phantom image, also used as ground truth image. (b-d) The inner rings become more visible as a increases. (e-h) Reconstructions of each proposed super-resolution approach, with (DART) and without (S-SIRT) prior knowledge. (i-l) For increasing values of a, plotting the relative Number of Missclassified Pixels (rNMP) in function of the widths of each ring.

rNMP < 0.30 is plotted as a function of the number of projections, for increasing levels of upsampling. To be able to reconstruct the smallest rings, a large upsampling factor is clearly required and many more projection angles must be used.

Other simulated phantoms: Experiments were also performed on the simulated datasets presented in Fig. 8. The three bone phantoms (Fig. 8a,b,d) are 1024×1024 pixel phantoms based on actual reconstructions of rat femurs and where $R_v = \frac{1}{4}$, i.e. $\Delta s = \frac{\Delta t}{4}$. Fig. 8c represents a set of 20 randomly generated, analytically defined polyurethane foam phantom images. The width of each cell wall was chosen randomly in the interval $[\frac{\Delta t}{2}, \frac{5\Delta t}{2}]$. Fig. 8a,c are binary images whereas Fig. 8b,d contain three distinct grey level values. It should be noted that the number of grey levels to be used in the reconstruction should not be too high. Otherwise, the prior knowledge is no longer sufficiently strong to optimally restrict the solution space. We have found that, depending on the shape of the objects, typically up to 4 or 5 unique grey levels may be present.

For each dataset, projection data was generated on a parallel beam projection geometry with 180 equiangular projection

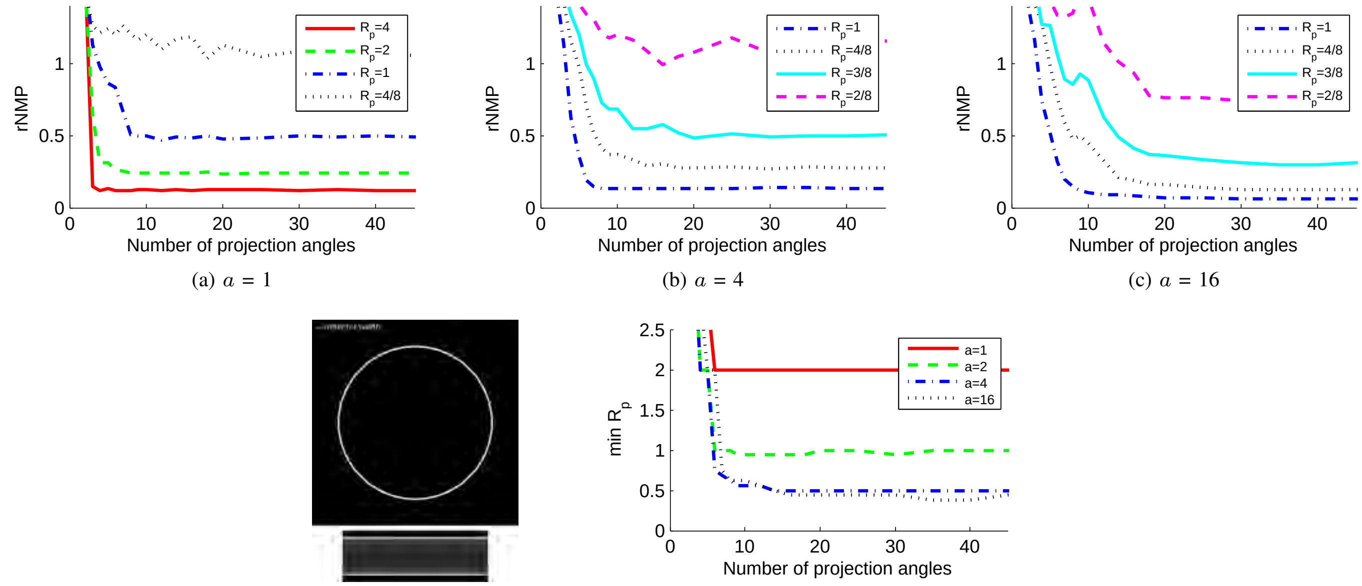

(a) a = 1 (b) a = 4 (c) a = 16

(d) phantom, $R_p = \frac{1}{2}$ (e) a = 16

Fig. 7: Experimental results from simulated analytical rings with a varying thickness Δq. (a-c) For different levels up upsampling, the rNMP of rings with a varying ring thickness as a function of the number of projection angles. After a certain point, additional projections offer no further accuracy improvements. This point is dependent on the thickness of the rings and on the level of upsampling. (d) High resolution rendering of a phantom ring with $R_p = \frac{1}{2}$. (e) The minimal R_p of the ring that can be reconstructed with rNMP < 0.30 as a function of the projection count. Higher levels of upsampling clearly result in better resolution but to reach the point of "sufficient information", more projections are then required.

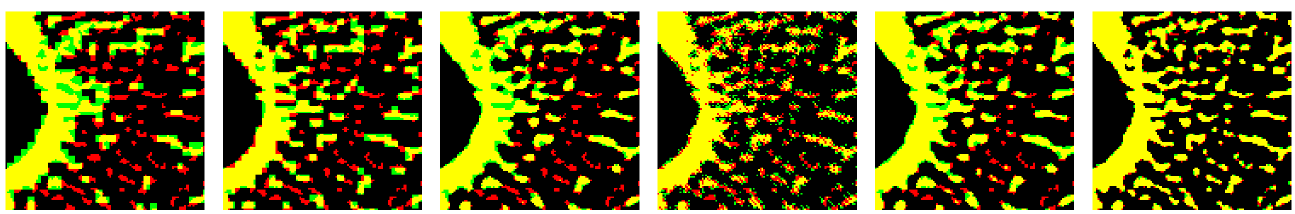

(a) $a = 1$ S-SIRT rNMP=0.388 (b) $a = 1$ DART rNMP=0.460 (c) SU, $a = 4$ S-SIRT rNMP=0.249 (d) SU, $a = 4$ DART rNMP=0.302 (e) DS, $a = 4$ S-SIRT rNMP=0.208 (f) DS, $a = 4$ DART **rNMP=0.086**

Fig. 9: Region of the various reconstructions of Fig. 8a on a parallel beam geometry with 180 projection direction and $I_0 = 20000$. The ground truth image is displayed in red and the reconstructions in green. Where both images overlap, i.e. where the segmentation is correct, the corresponding pixel is yellow.

Phantom	$a = 1$		$a > 1$				
			a	SU		DS	
	S-SIRT	DART		S-SIRT	DART	S-SIRT	DART
Fig. 8a	0.388	0.460	4	0.249	0.302	0.208	**0.086**
Fig. 8b	0.519	0.036	4	0.519	0.026	0.519	**0.003**
Fig. 8c	0.468	0.347	8	0.463	0.381	0.389	**0.107**
Fig. 8d	0.168	0.172	4	0.112	0.090	0.091	**0.068**

Fig. 10: Numerical results (rNMP) for all phantom experiments of Fig. 8. For the set of phantom Fig. 8(c), the average rNMPs are given. A parallel beam geometry with 180 projection direction and $I_0 = 20000$ was used.

angles and 256 detector cells with $\Delta t = 1$. For the pixel based phantoms (Fig. 8a,b,d), the PVE was induced by simulating high resolution projection data (with 1024 detector cells with $\Delta t = \frac{1}{4}$) in the intensity domain (i.e. $\boldsymbol{I}$, Eq. (3)). The detector bins were then summed 4 by 4 after which the resulting data was converted to the attenuation domain (i.e. to $\boldsymbol{p}$, Eq. (4)). For the analytical phantoms (Fig. 8c), projection data was computed analytically, inherently modelling the PVE. For every dataset, Poisson noise was applied; the intensity of which is defined by the incident beam intensity, I_0. In these experiments, $I_0 = 20000$.

To quantify the segmentation accuracy, the rNMP measure

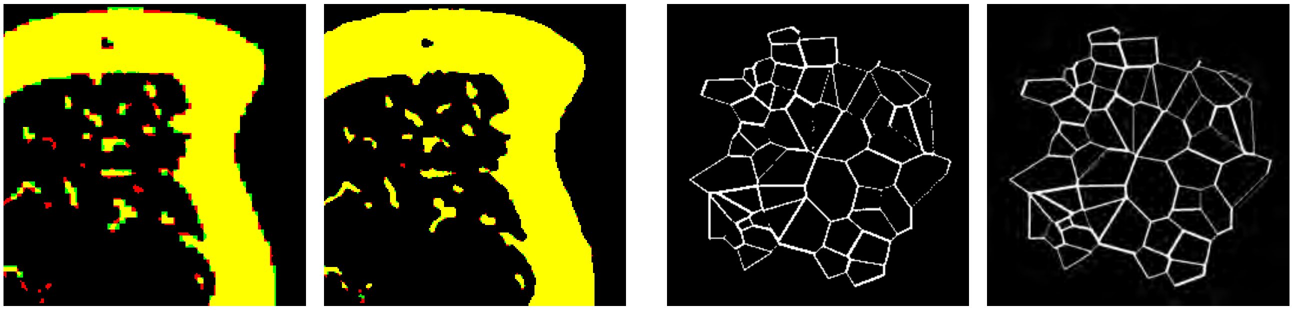

(a) $a = 1$, DART
rNMP = 0.036

(b) DS, $a = 4$, DART
rNMP = 0.003

(c) $a = 1$, DART
avg(rNMP) = 0.347

(d) DS, $a = 8$, DART
avg(rNMP) = 0.107

Fig. 11: (a,b) Detector supersampling for phantom Fig. 8b. (c,d) Detector supersampling for phantom Fig. 8c.

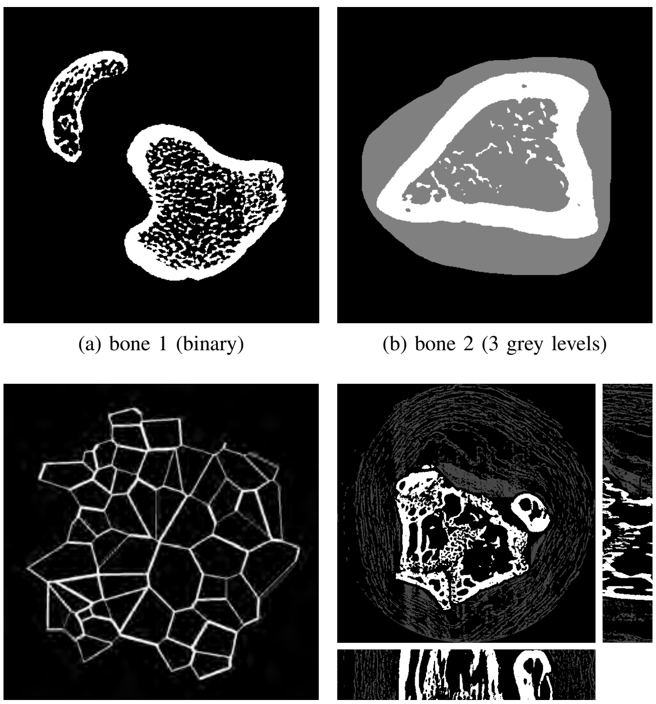

(a) bone 1 (binary)

(b) bone 2 (3 grey levels)

(c) one of 20 foams (binary)

(d) bone 3 (3 grey levels, 3D)

Fig. 8: Simulated phantom images. Projection data of phantom (a), (b) and (d) was generated from high resolution pixel images, based on actual reconstructions of rat femurs. The set of phantoms (c) were analytically defined and their projection data were also calculated analytically.

was computed. As the experiments were performed at varying pixel or voxel sizes, the reconstructions were first rescaled to the size of the original, high resolution ground truth images. For the analytical phantoms, high resolution rasterizations were used as the ground truth. For phantoms Fig. 8b and Fig. 8d, which contain an additional distinct grey level value representing soft-tissue, the rNMP was computed with respect to the most dense partition, i.e. the bone structures.

The following reconstruction methods were evaluated: S-SIRT (visualised for phantom Fig. 8a in Fig. 9a,c,e) versus DART (Fig. 9b,d,f); no super-resolution approach ($a = 1$, Fig. 9a,b) versus a super-resolution approach ($a > 1$, Fig. 9c,d,e,f); and sinogram upsampling (Fig. 9c,d) versus detector supersampling (Fig. 9e,f). In Fig. 10 the rNMP values for all phantoms are shown.

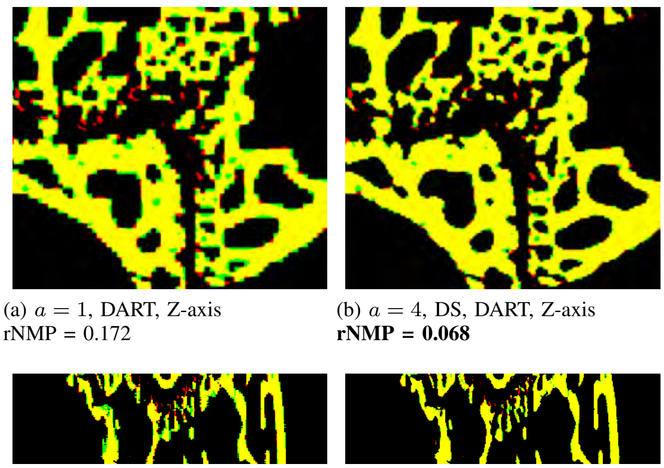

(a) $a = 1$, DART, Z-axis
rNMP = 0.172

(b) $a = 4$, DS, DART, Z-axis
rNMP = 0.068

(c) $a = 1$, DART, X-axis

(d) $a = 4$, DS, DART, X-axis

Fig. 12: The improvement of detector supersampling on 3D-DART reconstructions of phantom Fig. 8d is clearly visible from slices through the Z-axis and the X-axis.

For phantom Fig. 8b and Fig. 8c, the advantage of using DART and detector supersampling can be seen in Fig. 11a-d. Small trabecular structures are properly segmented only on an upsampled reconstruction grid. Similar results can be seen for the foam segmentation, where it is clear that especially the thinnest cell edges benefit the most from the proposed super-resolution approach.

Fig. 12 shows the improvement of detector supersampling on 3D DART reconstructions of phantom Fig. 8d for two orthogonal viewing directions. It can be seen that by applying super-resolution, the small three-dimensional trabecular structures are segmented much more accurately, also in the XZ-slices.

Limited view problem: A method to reduce the radiation dose is to reduce the number of projection angles. This, however, leads to limited data reconstruction problems. To demonstrate the effect of this limited view problem on the

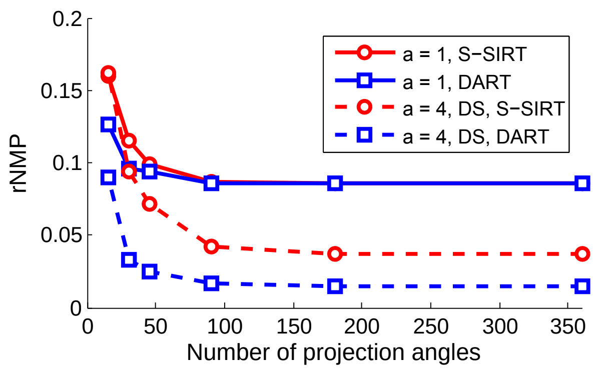

(a) phantom Fig. 8b

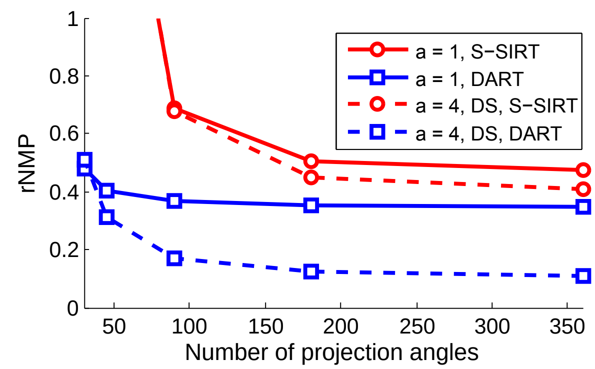

(b) phantom Fig. 8c

Fig. 13: The rNMP in function of the number of projection angles for phantoms Fig. 8b and Fig. 8c and for $a = 1$ and $a = 4$.

proposed super-resolution, projection data of Fig. 8b and Fig. 8c was generated several times, with the same I_0 and the downsampling strategy as before, with a decreasing number of projection angles, effectively simulating scans with a reduced radiation dose. For each set of projection data, DART and S-SIRT reconstructions were created with and without the detector supersampling approach. The rNMP values are plotted in Fig. 13a and Fig. 13b. One can conclude that, even with a drastically lowered number of projection angles, the combination of detector supersampling with the exploitation of prior knowledge results in reconstructed images that are more accurate than conventional S-SIRT reconstructions without a super-resolution approach and with a high number of projection angles.

Robustness of assumed principles: In the proposed super-resolution approach, it is assumed that the object has a homogeneous density and that this density is known in advance. A study was performed to investigate what happens if one of these assumptions are only approximately satisfied.

To demonstrate the robustness of our algorithm with respect to deviations from the first assumption, each pixel of phantom Fig. 8b was multiplied with a normally distributed random number with mean = 1. This was done multiple times with an increasing standard deviation. For each such image, projection data was generated with 30 projection angles, downsampled by a factor 4 - as explained before - and with $I_0 = 20000$. Reconstructions were made with S-SIRT and DART and with detector supersampling ($a = 1$ and $a = 4$). In Fig. 14a, the rNMP values are plotted in function of the standard deviation of the applied noise. While the rNMP of DART with the super-resolution approach indeed increases as the objects grows more and more inhomogeneous, improvements over the conventional methods are still achieved.

For phantom Fig. 8a, projection data was created with 30 projection angles, again downsampled by a factor 4 and with $I_0 = 20000$. Multiple DART reconstructions were created where the assumed grey level was varied between 0.8 and 1.2 times the correct grey level. Fig. 14b plots the rNMP values for these DART reconstruction with detector supersampling ($a = 1$ and $a = 4$) and for S-SIRT. It can be seen that the rNMP of the DART reconstructions indeed increases as the assumed grey level is incorrect. However, drastic improvements over the conventional S-SIRT method without super-resolution can be achieved even if the chosen grey level is just an approximation of the correct grey level.

In [16] and [17], methods were proposed to automatically estimate the optimal grey levels of piece homogeneous objects.

B. Real-world experiments

The proposed method was applied to real μCT data. Fig. 15a shows an FBP reconstruction of a slice through a human mandible, which was recorded using a SkyScan 1173 μCT scanner with 900 equiangular projection angles in the interval $[0, \pi)$. The detector resolution was 50μm. The data was corrected for ring artefacts and beam hardening with the standard SkyScan NRecon software package. Only 100 projection angles were used in the experiments and the projection data was downsampled by summing the detector bins 4 by 4 in the intensity domain (Eq. (3)), such that many of the smaller structures were relatively small compared to the new detector sizes. Three distinct grey level values were used. One for air, one for soft tissue and one for bone. Validation was performed using the pixels of the latter category only.

Given the polychromatic nature of the X-ray spectrum, in practice, one can not know the exact grey levels of the materials in advance. In this experiment, the discrete grey levels $\boldsymbol{\rho}$ were manually approximated, guided by the grey levels present in the initial SIRT reconstruction, $v^{(0)}$. If the initial reconstruction is too erroneous to accurate estimate the grey level values, e.g. when there are only a few projections, one could use one of the grey level estimation algorithms that have been presented in the literature, such as the semi-automatic DGLS method [16] or an automatic method that combines grey level estimation into DART [17].

From the FBP reconstruction (Fig. 15a), created with the non-downsampled and the full set of projection data, a segmentation was manually created (Fig. 15b). This segmentation was used to validate the results by measuring the rNMP. This FBP reconstruction was also segmented using Otsu's clustering method (Fig. 15c). Fig. 15d-g indicate that also

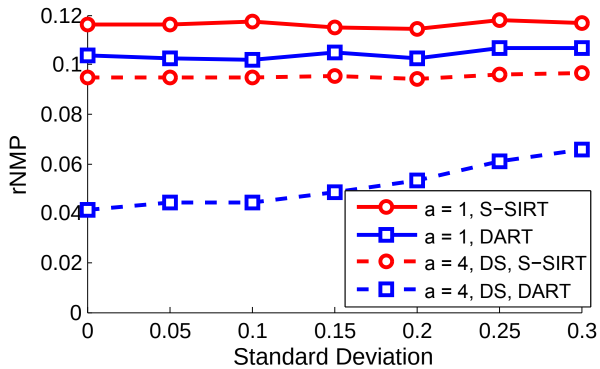

(a) phantom Fig. 8b

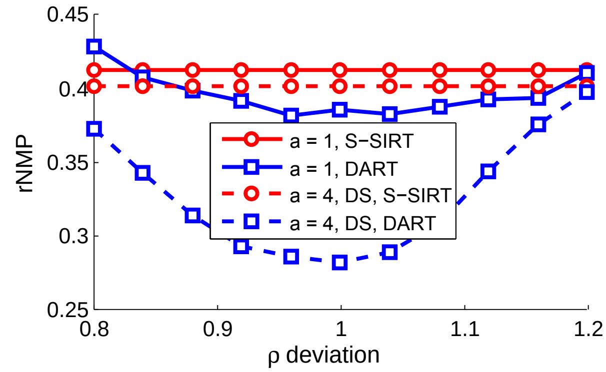

(b) phantom Fig. 8a

Fig. 14: (a) The rNMP in function of the standard deviation of the normally distributed noise that was multiplied with the phantom image Fig. 8b prior to simulating the projection data. (b) The rNMP in function of the deviation on the correct grey level ρ during DART reconstructions of Fig. 8c.

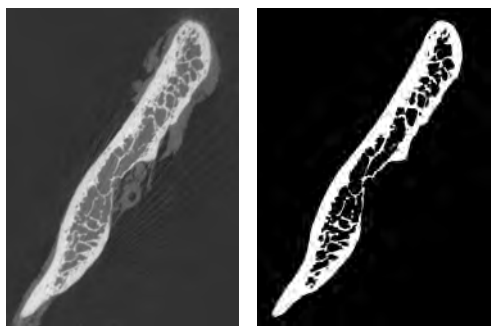

(a) $R_p = 1, a = 1$
FBP reconstruction

(b) $R_p = 1, a = 1$
manual segmentation

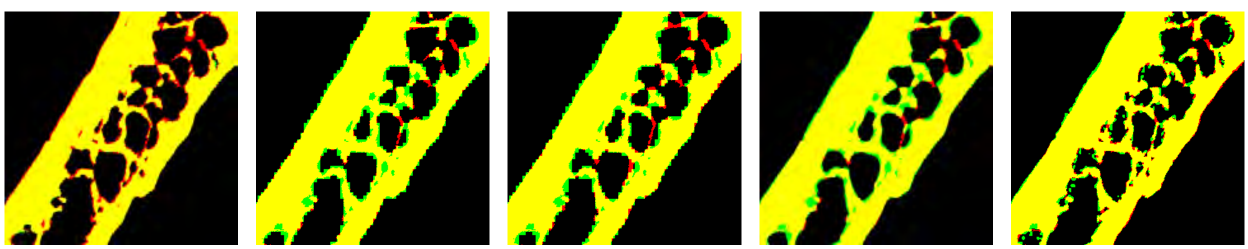

(c) $R_p = 1$
$a = 1$, FBP
rNMP=0.106

(d) $R_p = 4$
$a = 1$, S-SIRT
rNMP=0.175

(e) $R_p = 4$
$a = 1$, DART
rNMP=0.184

(f) $R_p = 4$
$a = 1$, DS, S-SIRT
rNMP=0.171

(g) $R_p = 4$
$a = 4$, DS, DART
rNMP=0.115

Fig. 15: Results for a real μCT dataset of a slice through a human mandible.

for real datasets, the addition of super-resolution and DART significantly improves the accuracy of the segmentation. When comparing Fig. 15c with Fig. 15g, it can be seen that with the proposed super-resolution approach, segmentations of low resolution projection data can be obtained that are of comparable quality to that of high resolution, high dose scans.

In Fig. 16, various reconstructions are shown of the polyurethane foam also shown in Fig. 1. In total, 500 projection images were taken in a SkyScan 1172 μCT scanner at a pixel resolution of 17μm. The projection data was downsampled by summing the detector bins 2 by 2 in the intensity domain and the SkyScan NRecon software package was used for ring artefact and beam hardening correction. Reconstructions were created with both S-SIRT and DART, and with $a = 1$ and $a = 4$ (which applies the DS super-resolution technique). In cut-outs, parts of the reconstructed images (green) are overlaid with a high resolution FBP reconstruction of the foam (red). As in the previous experiments, the application of DART without the use of a super-resolution technique, does not result in improved image quality. However, when discrete tomography and super-resolution are combined, even the thin cell walls become clearly visible.

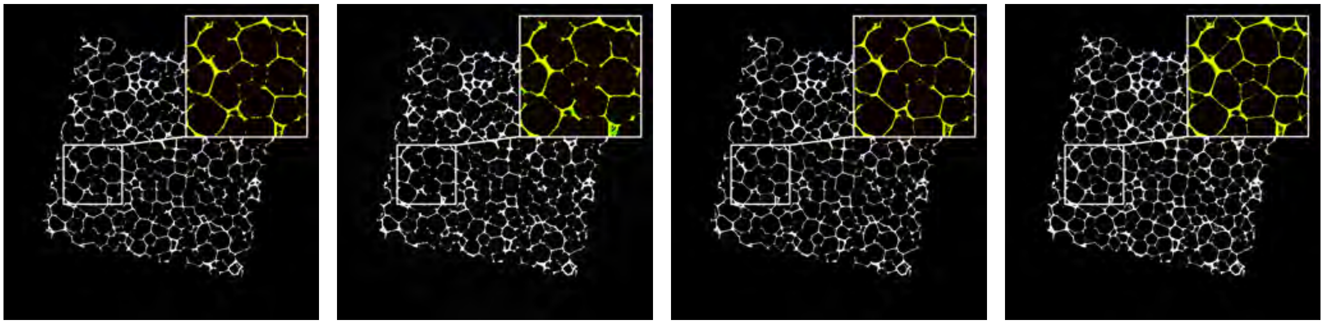

(a) $a = 1$, S-SIRT (b) $a = 1$, DART (c) $a = 1$, DS, S-SIRT (d) $a = 4$, DS, DART

Fig. 16: Results for a real μCT dataset containing a polyurethane foam.

V. Discussion and conclusion

Accurate segmentation of structures that are small with respect to the reconstruction pixel size, poses a very complex and difficult problem as reconstructed images often lack contrast due to a partial volume effect. This often means that even the intrinsic detector resolution can not be achieved. High resolution reconstructions can provide a solution, but are often not feasible due to X-ray dose limitations, limited scanning time or hardware constraints.

To improve the detection of small structures in low resolution CT acquisitions, the use of a super-resolution approach has been proposed in which reconstructed images are computed on an upsampled reconstruction grid. Two geometrical methods for achieving super-resolution have been investigated: sinogram upsampling, where the projection data is upsampled by linear interpolation and detector supersampling, where multiple rays per detector element are cast through the reconstruction grid. Both methods result in a limited data problem. It was shown that with discrete tomography (DART) prior knowledge about the object materials, can be exploited to overcome this problem and thus increase the resolution of the reconstruction. Previously, DART has only been used to reduce the number of projection angles.

Experiments were performed on simulated as well as real data of objects containing small structures. Without using a super-resolution technique on objects containing small structures, the addition of prior knowledge (DART) sometimes resulted in less accurate segmentations when compared to the conventional S-SIRT algorithm. This effect was predicted in Fig. 4, where it was noted that DART is only suited for objects that are large with respect to the pixel size. However, if a super-resolution technique was applied, the use of prior knowledge with the DART method clearly resulted in more accurate reconstructions than the conventional S-SIRT approach. This effect was observed to be generally more profound if detector supersampling was chosen over sinogram upsampling.

In conclusion, the use of the detector supersampling super-resolution technique in which prior knowledge about the object density is exploited, can effectively increase the spatial resolution of a reconstructed image. In that way, small structures can be segmented more accurately with a shorter scan time and a lower radiation dose.

Acknowledgment

This work was financially supported by the iMinds-AIR project (Interdisciplinary Institute for Technology, a research institute founded by the Flemish Government). K.J.B. was supported by the Netherlands Organisation for Scientific Research (NWO), programme 639.072.005.

References

[1] N. Otsu, "A threshold selection method from gray level histograms," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 9, pp. 62–66, March 1979.

[2] P. L. Salmon and A. Y. Sasov, *Advanced Bioimaging Technologies in Assessment of the Quality of Bone and Scaffold Materials*. Springer Berlin Heidelberg, 2007.

[3] A. Parfitt, M. Drezner, F. Glorieux, J. Kanis, H. Malluche, P. Meunier, S. Ott, and R. Recker, "Bone Histomorphometry: standardization of nomenclature, symbols and units," *Journal of Bone and Mineral Research*, vol. 2, no. 6, pp. 595–610, 1987.

[4] E. Perilli, V. Le, B. Ma, P. Salmon, K. Reynolds, and F. N.L., "Detecting early bone changes using in vivo micro-ct in ovariectomized, zoledronic acid-treated, and sham-operated rats," *Osteoporosis International*, vol. 21, pp. 1371–1382, 2009.

[5] S. Bals, K. J. Batenburg, J. Verbeeck, J. Sijbers, and G. Van Tendeloo, "Quantitative Three-Dimensional Reconstruction of Catalyst Particles for Bamboo-like Carbon Nanotubes," *Nano Letters*, vol. 7, no. 12, pp. 3669–3674, 2007.

[6] S. C. Park, M. K. Park, and M. G. Kang, "Super-resolution image reconstruction: a technical overview," *IEEE Signal Processing Magazine*, vol. 20, no. 3, pp. 21–36, May 2003.

[7] J. A. Kennedy, O. Israel, A. Frenkel, R. Bar-Shalom, and H. Azhari, "Super-resolution in PET imaging," *IEEE Transactions on Medical Imaging*, vol. 25, pp. 137–147, 2006.

[8] G. Marai, D. Laidlaw, and J. Crisco, "Super-resolution registration using tissue-classified distance fields," *IEEE Transactions on Medical Imaging*, vol. 25, pp. 177–187, 2006.

[9] A. Patil, R. Singhai, and J. Singhai, "Curvelet transform based super-resolution using sub-pixel image registration," in *Computer Science and Electronic Engineering Conference (CEEC)*, 2010, pp. 1–6.

[10] E. Sidky, Y. Duchin, X. Pan, and C. Ullberg, "A constrained, total-variation minimization algorithm for low-intensity x-ray ct," *Medical Physics*, vol. 38, 2011.

[11] M. Li, H. Yang, and H. Kudo, "An accurate iterative reconstruction algorithm for sparse objects: application to 3D blood vessel reconstruction from a limited number of projections," *Physics in Medicine and Biology*, vol. 47, pp. 2599–2609, 2010.

[12] E. Sidky and X. Pan, "Image reconstruction by constrained, total-variation minimization," *Physics in Medicine and Biology*, vol. 53, pp. 4777–4807, 2008.

[13] 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," *Physics in Medicine and Biology*, vol. 55, pp. 6575–6599, 2010.

[14] K. J. Batenburg and J. Sijbers, "DART: A practical reconstruction algorithm for discrete tomography," *IEEE Transactions on Image Processing*, vol. 20, no. 9, pp. 2542–2553, 2011.

[15] J. Gregor and T. Benson, "Computational analysis and improvement of SIRT," *IEEE Transactions on Medical Imaging*, vol. 27, no. 7, pp. 918–924, 2008.

[16] K. J. Batenburg, W. van Aarle, and J. Sijbers, "A semi-automatic algorithm for grey level estimation in tomography," *Pattern Recognition Letters*, vol. 32, pp. 1395–1405, 2011.

[17] W. van Aarle, K. Batenburg, and J. Sijbers, "Automatic parameter estimation for the discrete algebraic reconstruction technique (dart)," *IEEE Transactions on Image Processing*, vol. 21, pp. 4608–4621, 2012.

[18] S. Van Aert, K. Batenburg, M. Rossell, R. Erni, and V. Tendeloo, "Three-dimensional atomic imaging of crystalline nanoparticles," *Nature*, vol. 470, pp. 374–377, 2011.

[19] E. Van de Casteele, K. Batenburg, P. Salmon, and S. Sijbers, "Bone segmentation using discrete tomography," in *Micro-CT User Meeting*, 2011.

[20] T. T. Roelandts, K. Batenburg, E. Biermans, K. C., S. Bals, and J. Sijbers, "Accurate segmentation of dense nanoparticles by partially discrete electron tomography," *Ultramicroscopy*, vol. 114, pp. 96–105, 2012.

[21] W. van Aarle, G. Van Gompel, K. J. Batenburg, E. Van de Casteele, and J. Sijbers, "A 3-dimensional discrete tomography approach for superresolution micro-ct images: application to foams," in *The first international conference on image formation in X-ray computed tomography*, Salt Lake City, 2010, pp. 45–48.

[22] G. Van Gompel, K. J. Batenburg, E. Van de Casteele, W. van Aarle, and J. Sijbers, "A discrete tomography approach for superresolution micro-CT images: application to bone," in *Biomedical Imaging: From Nano to Macro, IEEE International Symposium on*, 2010, pp. 816–819.

[23] P. M. Joseph and C. Ruth, "A method for simultaneous correction of spectrum hardening artifacts in ct images containing both bone and iodine," *Medical Physics*, vol. 24, pp. 1629–1634, 1997.

[24] E. Van de Casteele, D. Van Dyck, J. Sijbers, and E. Raman, "A model-based correction method for beam hardening artefacts in x-ray microtomography," *Journal of X-ray science and technology*, vol. 12, no. 1, pp. 53–57, 2004.

[25] M. Krumm, S. Kasperl, and M. Franz, "Reducing non-linear artifacts of multi-material objects in industrial 3D computed tomography," *Nondestructive Testing and Evaluation*, vol. 41, pp. 242–251, 2008.

[26] G. Van Gompel, K. Van Slambrouck, M. Defrise, K. J. Batenburg, J. de Mey, J. Sijbers, and J. Nuyts, "Iterative correction of beam hardening artifacts in CT," *Medical Physics*, vol. 38, pp. S36,1–S36,14, 2011.

[27] A. C. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging*, R. Cotellessa, Ed. New York, NY: IEEE Press, 1988, vol. Algorithms for reconstruction with non-diffracting sources, 49-112.

[28] 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, pp. 730–740, 2009.

[29] K. J. Batenburg, J. Sijbers, H. F. Poulsen, and E. Knudsen, "DART: a robust algorithm for fast reconstruction of three-dimensional grain maps," *Journal of Applied Crystalography*, vol. 43, no. 6, pp. 1464–1473, 2010.