

Optimal Threshold Selection for Segmentation of Dense Homogeneous Objects in Tomographic Reconstructions

Wim van Aarle*, Kees Joost Batenburg, and Jan Sijbers

Abstract—In this paper, we present a novel approach to segment dense, homogeneous objects in a tomographic reconstruction (or tomogram). A popular method to extract such objects from a tomogram is global thresholding, in which the threshold value is determined from the image histogram. However, accurate threshold selection is not straightforward, since, due to noise or artefacts in the reconstruction, the histogram does not always contain a clear, separate peak for the dense object. We propose a new threshold estimation approach, segmentation inconsistency minimization, that exploits the available projection data to determine the optimal global threshold. The proposed algorithm was tested on simulation data and on experimental μ CT data. The results show that this method results in more accurate segmentations, compared to alternative threshold selection methods.

Index Terms—Computed tomography (CT), dense objects, iterative reconstructions, segmentation, thresholding, tomography.

I. INTRODUCTION

THIS paper deals with the segmentation of dense objects in images obtained from computed tomography (CT), also known as tomograms. It is assumed that the objects to be segmented have a constant density that is higher than that of the surrounding materials. In medical imaging, dense object segmentation is required in many applications. For example, suppression of streak or beam hardening artefacts caused by metal implants requires segmentation of the medical tomograms [1]–[3]. Object identification and motion estimation (e.g., for coregistration) often requires detection of implanted markers on a device or guide wire [4], [5]. Accurate localization of individual cochlear implant electrodes within the inner ear is important to model the electrical field of the cochlea [4], [6]. Extraction of the trabecular bone features, such as the cortical thickness or the cortical area, from the surrounding marrow spaces involves segmentation of the dense bone with respect to the background [7], [8].

Manuscript received November 19, 2010; revised December 22, 2010; accepted December 23, 2010. Date of publication January 06, 2011; date of current version April 01, 2011. This work was supported in part by IBBT and in part by the SBO-project QUANTIVIAM (060819) of the Institute for the Promotion of Innovation through Science and Technology in Flanders. *Asterisk indicates corresponding author.*

*W. van Aarle is with the IBBT-Visionlab, University of Antwerp, B-2610 Antwerp, Belgium (e-mail: wim.vanaarle@ua.ac.be).

J. Sijbers is with the IBBT-Visionlab, University of Antwerp, B-2610 Antwerp, Belgium (e-mail: jan.sijbers@ua.ac.be).

K. J. Batenburg is with the Centrum Wiskunde & Informatica, 1098 XG Amsterdam, The Netherlands (e-mail: joost.batenburg@ua.ac.be).

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

Digital Object Identifier 10.1109/TMI.2010.2104328

Without noise or artefacts in the CT reconstruction, segmentation of dense homogeneous objects would be trivial. Unfortunately, in practice, accurate separation of such objects from the surroundings within a tomographic reconstruction is a non-trivial task for several reasons.

- *Limited number of projection images.* In many cases, the number of available projections is not sufficient to guarantee a unique reconstruction. Therefore, the computed reconstruction most likely does not correspond entirely to the underlying, unknown object.
- *Noise or artefacts in the measured projection data.* Real-world data is inevitably polluted by noise and artefacts for example caused by malfunctioning detector elements or scatter, leading to inaccuracies in the reconstruction.
- *Approximations in the reconstruction algorithm.* Common reconstruction algorithms typically do not compute an exact inverse of the Radon transform, resulting in discrepancies between the reconstruction and the original object.

As a consequence, dense objects do not always show up as clear peaks in the histogram of a tomographic reconstruction.

For dense object segmentation, it is a common choice to set a global threshold somewhere between the grey level of the pixels belonging to the object and those of the maximum value of the other pixels, which we call the background in the remainder of this paper [9]. Typically, this threshold is selected based on the histogram of the tomogram [10]. If only a few materials are present and each of these correspond to a distinct grey level peak in the histogram, it is possible to accurately determine appropriate thresholds, for example by analyzing the concavity points on the convex hull of the histogram [11] or by modeling the histogram as a mixture of a series of Gaussian distributions [12]. The most popular global threshold selection method is the clustering method of Otsu [13]. It minimizes the weighted sum of intra-class variances of the different segmentation partitions.

The problem with histogram-based methods in the context of segmenting a homogeneous object in a continuous grey level image, however, is that there are no guaranteed histogram peaks representing the continuous background. Histogram-based methods are particularly inadequate if the object of interest is only slightly more dense than the surrounding materials.

Different approaches to segmentation of dense objects also exist, e.g., region-based algorithms such as region growing [14] and watershed segmentation [15]. These methods, however, are also solely based on the reconstructed image and are therefore very susceptible to reconstruction artefacts.

Ideally, reconstruction algorithms for tomography should be “invertible,” so that computed projections of the reconstructed image would equal the measured projection data. For filtered backprojection, the most common reconstruction algorithm used in practice, this assumption does not hold, mainly due to various interpolation steps involved in the algorithm. Iterative algebraic methods (e.g., SIRT [16]), which are less commonly used due to their computational requirements, only satisfy the invertibility assumption for the case of noiseless projection data and an infinite number of iterations. As the reconstructed image does not correspond accurately with the measured projections, using the projection data for the segmentation can potentially result in a segmentation that is more faithful to the original measurements.

Recently, a new method was proposed for global [17] and local [18] threshold selection in tomograms, called projection distance minimization (PDM). This approach is based on the assumption that the scanned object contains a small number of different densities, each corresponding to a constant grey level in the reconstruction. By segmenting the reconstructed image, this property is restored in the tomogram. To measure the quality of this segmentation, projections of the segmentation are computed and compared to the measured projection data. An optimal segmentation will result in maximal correspondence between the simulated projections and the measured dataset. However, this approach requires that the scanned object contains only a few different densities and does not allow for segmentation of objects with a constant greylevel in a continuously varying surrounding.

In this paper, we will introduce a new global thresholding method, the segmentation inconsistency minimization (SICM) method, for dense object segmentation that employs similar concepts as the mentioned PDM methods, while allowing the remaining part of the image to vary freely. For each candidate segmentation, the projections of the segmented object are subtracted from the measured projection data, after which the remaining part of the image is reconstructed and checked for consistency with the residual projections. The threshold for which minimal inconsistency is obtained is selected for the segmentation. The only assumptions that are made, is that the density of the object is constant and that it is higher than all remaining densities in the scanned object.

In [19], preliminary work on this topic was published and an algorithm called segmentation consistency maximization (SCM) was introduced. The basic strategy behind the newly proposed SICM is similar to that of SCM, but there are significant differences as well, such as the addition of automatic grey level estimation and a more elaborate optimization technique.

We will focus on μ CT applications such as the segmentation of bones, cochlear implants, foams, etc. The applicability of the proposed method is not limited to these applications, however. Indeed, the method can also be applied in regular CT or in electron tomography.

The paper is structured as follows. In Section II, the tomography setting is introduced. Section III describes our threshold selection algorithm in detail. Experimental results are presented in Section IV. Section V concludes the paper.

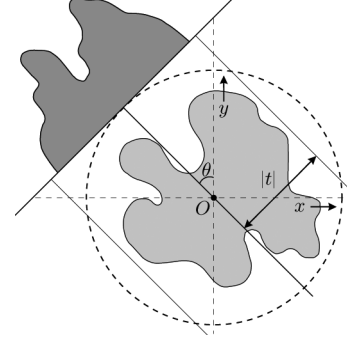


Fig. 1. Basic setting of transmission tomography.

II. NOTATION AND CONCEPTS

In this section, we introduce a notation for basic tomography operations such as forward projection and reconstruction. We also introduce a definition of sinogram inconsistency. These concepts are used in Section III, where they will be used to define a dense object segmentation algorithm.

A. Computed Tomography

Let $\mathbf{v} \in \mathbb{R}^n$ denote the discretized square image of an object, with n , the number of pixels. Assume that the object is completely contained in this square.

Projections of this image are measured along lines $l_{\theta,t} = \{(x,y) \in \mathbb{R} \times \mathbb{R} : x \cos \theta + y \sin \theta = t\}$, where θ represents the angle between the line and the y -axis and t represents the coordinate along the projection axis; see Fig. 1. In practice, a projection is measured at a finite set of projection angles and at a finite set of detector cells, each measuring the integral of the object density along a ray. Let m denote the total number of measured detector values for all angles and let $\mathbf{p} \in \mathbb{R}^m$ denote the measured projection data.

The forward projection of the object for a finite set of angles can be modeled as a linear operator \mathbf{W} , called the projection operator, that maps the image \mathbf{v} to the projection data \mathbf{q}

$$\mathbf{q} := \mathbf{W}\mathbf{v}. \quad (1)$$

In (1), $\mathbf{W} = (w_{ij})$ is an $m \times n$ matrix where w_{ij} represents the contribution of image pixel v_j to detector value p_i . The vector \mathbf{q} is called the forward projection or sinogram of \mathbf{v} .

The reconstruction problem in transmission CT concerns the recovery of \mathbf{v} from a given vector \mathbf{p} of projection data, such that

$$\mathbf{W}\mathbf{v} = \mathbf{p}. \quad (2)$$

Many reconstruction algorithms exist, such as filtered backprojection (FBP), Feldkamp (FDK), etc. In the remainder of this paper, however, we will focus on a single algorithm, namely the simultaneous iterative reconstruction technique (SIRT) [16].

B. SIRT Reconstruction

The SIRT algorithm is an iterative reconstruction technique that finds the least squares solution $\hat{\mathbf{v}}$ to the system of equations in (2). It has several favorable mathematical and computational properties, such as guaranteed convergence and linearity.

We describe the iteration update for the SIRT algorithm. Let $\tilde{\mathbf{v}}^{(0)} = \mathbf{0}$. For $q = 1, 2, \dots$, let $\mathbf{r}^{(q)} = \mathbf{p} - \mathbf{W}\tilde{\mathbf{v}}^{(q-1)}$ be the projection difference before the q^{th} iteration. In each new iteration q , the current reconstruction $\tilde{\mathbf{v}}^{(q-1)}$ is updated, yielding a new reconstruction $\tilde{\mathbf{v}}^{(q)}$, as follows:

$$\tilde{\mathbf{v}}_j^{(q)} = \tilde{\mathbf{v}}_j^{(q-1)} + \frac{1}{\sum_{i=1}^n w_{ij}} \sum_{i=1}^m \frac{w_{ij} r_i^{(q)}}{\sum_{j=1}^n w_{ij}}. \quad (3)$$

In the remainder of this paper, we will use the symbol \mathbf{S} to denote the linear operator that creates a SIRT-reconstruction of a sinogram \mathbf{p} with a certain fixed number of iterations, i.e., $\tilde{\mathbf{v}} = \mathbf{S}\mathbf{p}$.

It can be shown [16] that the SIRT-algorithm as described in (3) converges to the solution $\tilde{\mathbf{v}}$ where the weighted squared projection difference $\|\mathbf{W}\tilde{\mathbf{v}} - \mathbf{p}\|_{\mathbf{R}}^2 = (\mathbf{W}\tilde{\mathbf{v}} - \mathbf{p})^T \mathbf{R} (\mathbf{W}\tilde{\mathbf{v}} - \mathbf{p})$ is minimal with \mathbf{R} , a diagonal matrix that contains the inverse row sums of \mathbf{W} : $r_{ii} = 1/\sum_j w_{ij}$.

The SIRT algorithm can also be performed on a subset $A \subset \{1, \dots, n\}$ of the image pixels by removing the columns of \mathbf{W} that are not in A . In this way, a reconstruction can be computed for which the projection difference is minimal in the set of all reconstruction images that are zero outside of A . We denote the SIRT reconstruction operator restricted to the set A by \mathbf{S}_A .

C. Sinogram Inconsistency

Not all vectors \mathbf{p} are valid sinograms. The set of all valid continuous sinograms has been characterized by Ludwig and Helgason in [20], [21]. They describe a set of conditions that must be satisfied by all sinograms, known as consistency conditions. In a discretized setting, where projection data is available only for a limited set of angles, a measured sinogram is called consistent if (2) is a consistent system, i.e., $\mathbf{p} \in \text{span}\{\mathbf{w}_i : 0 < i \leq n\}$. In practice, a sinogram will rarely be consistent due to noise, discretization, partial volume effects, etc. We therefore also introduce the inconsistency of a sinogram \mathbf{p} given by

$$\min_{\mathbf{x} \in \mathbb{R}^n} \|\mathbf{W}\mathbf{x} - \mathbf{p}\|$$

(with $\|\cdot\|$ a vector norm, defined by the reconstruction algorithm), i.e., the distance between \mathbf{p} and the nearest consistent sinogram.

An important role in our proposed algorithm is played by sets of sinograms that correspond to images where certain pixels are known to be zero. A sinogram $\mathbf{p} \in \mathbb{R}^m$ is called A -consistent if $\mathbf{p} \in \text{span}\{\mathbf{w}_a : a \in A\}$. Thus, for each A -consistent sinogram \mathbf{p} there exists a reconstructed image $\tilde{\mathbf{v}}$ with $\mathbf{W}\tilde{\mathbf{v}} = \mathbf{p}$ and $\tilde{v}_i = 0$ for each $i \notin A$. The distance between \mathbf{p} and the nearest A -consistent vector is called the A -inconsistency of \mathbf{p} . Note that when $A = \{1, \dots, n\}$, the concepts of inconsistency and A -inconsistency are equivalent.

It is not possible to compute the A -inconsistency directly. However, to approximate the A -inconsistency of a vector \mathbf{p} with respect to the norm $\|\cdot\|_{\mathbf{R}}$, one can compute a SIRT reconstruction of \mathbf{p} restricted to A with a fixed number of iterations, compute the forward projection of this reconstruction and compare it to the vector \mathbf{p} of measured projections. Define the A -pseudo-inconsistency of \mathbf{p} by

$$IC_A(\mathbf{p}) = \|\mathbf{W}\mathbf{S}_A\mathbf{p} - \mathbf{p}\|_{\mathbf{R}}. \quad (4)$$

III. COMPUTATIONAL APPROACH

In this section, we introduce the segmentation inconsistency (SIC) measure for determining the quality of a thresholded segmentation of the dense objects in the reconstructed image. We will explain this metric 1) by arguing that if a threshold is chosen too low, the segmentation inconsistency will be higher and 2) by experimentally showing that if the threshold is chosen too high, the measured segmentation inconsistency will also increase.

Subsequently, we present the SICM algorithm. This algorithm is an optimization technique to minimize the segmentation inconsistency. In each function evaluation the SIC will be computed on the residual sinogram that is created by subtracting the forward projection of the dense object from the measured projection data. Note that this is possible only if the grey level value of the dense object is known. We will therefore also discuss a method for the automatic estimation of this value.

A. Segmentation Inconsistency

Although measured sinograms will always be polluted by noise and other errors, let us assume, for the sake of clarity, that a “perfect” sinogram \mathbf{p} has been measured of an unknown image \mathbf{v} , i.e., $\mathbf{W}\mathbf{v} = \mathbf{p}$. All that is known of \mathbf{v} is that it contains one or more dense objects with a constant, maximal grey level value ρ and thus satisfies the prerequisites of the suggested algorithm. Let $\tilde{\mathbf{v}}$ be an image, the tomogram of \mathbf{v} , such that $\mathbf{W}\tilde{\mathbf{v}} \approx \mathbf{p}$. The tomogram may have been computed by any reconstruction algorithm, e.g., FBP or SIRT.

Put $B = \{b : v_b = \rho\}$, the set of pixels belonging to the dense object, i.e., the set of pixels to be found by the algorithm. We seek to approximate B by applying a threshold operation to $\tilde{\mathbf{v}}$. Let $\tilde{\tau}$ be a candidate threshold. Define $\tilde{B} = \{b : \tilde{v}_b \geq \tilde{\tau}\}$, the estimated set of pixels of the homogenous, dense object. Define $\tilde{A} = \{a : \tilde{v}_a < \tilde{\tau}\}$, complementary to \tilde{B} , the estimated set of pixels of the background.

For any image $\tilde{\mathbf{v}} \in \mathbb{R}^n$ and any set $\tilde{A} \subset \{1, \dots, n\}$, let $\tilde{\mathbf{v}}_{\tilde{A}}$ be an image that is equal to $\tilde{\mathbf{v}}$ for all pixels $i \in \tilde{A}$ and 0 for all pixels $i \notin \tilde{A}$.

Let $\tilde{\mathbf{s}} = (\tilde{s}_i) \in \mathbb{R}^n$ denote the segmentation of the dense object, based on the estimated segmentation \tilde{B}

$$\tilde{s}_i = \begin{cases} 0, & \text{if } i \in \tilde{A} \\ \rho, & \text{if } i \in \tilde{B}. \end{cases} \quad (5)$$

Based on this choice, we can now divide \mathbf{p} into two parts: a part that belongs to the dense object: $\mathbf{W}\tilde{\mathbf{s}}$ and a part that belongs to the background of image \mathbf{v} . The second part is called the residual sinogram of the region \tilde{A} and is given by

$$\mathbf{p}_{\tilde{A}} := \mathbf{p} - \mathbf{W}\tilde{\mathbf{s}}. \quad (6)$$

Definition: The segmentation inconsistency of any segmentation image $\tilde{\mathbf{s}}$ with only two grey level values (0 for background pixels and ρ for foreground pixels), is defined as

$$\text{SIC}(\tilde{\mathbf{s}}) = IC_{\tilde{A}}(\mathbf{p}_{\tilde{A}}) := \|\mathbf{W}\mathbf{S}_{\tilde{A}}(\mathbf{p} - \mathbf{W}\tilde{\mathbf{s}}) - (\mathbf{p} - \mathbf{W}\tilde{\mathbf{s}})\|_{\mathbf{R}}. \quad (7)$$

We will now argue and demonstrate why $\text{SIC}(\tilde{\mathbf{s}})$ is a useful measure for determining the quality of a segmentation $\tilde{\mathbf{s}}$.

- 1) Suppose that the reconstructed set \tilde{B} is an overestimation of B , i.e., $B \subset \tilde{B}$. This typically happens if $\tilde{\tau}$ is chosen too low.

Put $\tilde{\mathbf{e}} := \tilde{\mathbf{s}} - \mathbf{v}_{\tilde{B}}$, a vector that contains the exact overestimation of the dense object. This vector is nonzero only for pixels in $\tilde{B} \setminus B$ and is strictly positive in this region because $\forall i \in \tilde{B} \setminus B : v_i < \rho$. The residual sinogram can then also be computed as follows:

$$\mathbf{p}_{\tilde{A}} := \mathbf{W}\mathbf{v}_{\tilde{A}} - \mathbf{W}\tilde{\mathbf{e}}. \quad (8)$$

As the set of \tilde{A} -consistent sinograms is a linear subspace of the set of all sinograms, $\mathbf{p}_{\tilde{A}}$ is \tilde{A} -consistent if and only if $\mathbf{W}\tilde{\mathbf{e}}$ is \tilde{A} -consistent, i.e., there exists an image $\tilde{\mathbf{x}} \in \mathbb{R}^n$ such that $\mathbf{W}\tilde{\mathbf{x}} = \mathbf{W}\tilde{\mathbf{e}}$ and $x_i = 0$ for all $i \notin \tilde{A}$.

If both the number of projections and the region $\tilde{B} \setminus B$ are small, it may occur that such an image $\tilde{\mathbf{x}}$ exists, resulting in two different segmentations for which the residual sinogram is consistent. However, as the number of projections increases and the difference between \tilde{B} and B becomes larger, it becomes highly unlikely that the residual sinogram is still consistent. In fact, the \tilde{A} -inconsistency will typically increase as \tilde{B} is made larger, which will be demonstrated in Section IV.

- 2) Now suppose that the reconstructed set \tilde{B} is an underestimation of B , i.e., $\tilde{B} \subset B$. This typically happens if $\tilde{\tau}$ is chosen too high. In that case, the residual sinogram $\mathbf{p}_{\tilde{A}}$ will be \tilde{A} -consistent. Here we can make effective use of the experimental convergence properties of the SIRT algorithm. As the segmented dense object becomes smaller (i.e., the threshold $\tilde{\tau}$ is increased), the background \tilde{A} becomes larger, resulting in slower convergence for the iterative SIRT algorithm that is applied to those pixels. If we terminate after a fixed number of iterations, the computed $\text{SIC}(\tilde{\mathbf{s}})$ will therefore generally increase along with the threshold $\tilde{\tau}$. Fig. 2 shows an experimental confirmation of this algorithm property. For a Shepp–Logan phantom image of 256×256 pixels using 100 projections, an increasingly large random subset of pixels was kept fixed at their true values, while computing the $\text{SIC}(\tilde{\mathbf{s}})$ for the remaining pixels. Fig. 2 shows a strictly decreasing relation between the size of \tilde{A} and the computed segmentation inconsistency.
- 3) Of course, as a result of reconstruction errors and artifacts, there may be a threshold interval where neither $B \subset \tilde{B}$ nor $\tilde{B} \subset B$. In that case, there will be an increase of segmentation inconsistency due to the false positive pixels, the effect proven in 1). However, the effect explained in 2) will be somewhat cancelled out due to a mixture of false negatives and false positives, where the size of the segmented region is about the same as the original dense object. Our experimental results in the next section suggest that the segmentation inconsistency measure can still be used as an effective measure for the segmentation quality within this range as well.

In summary, to determine the optimal dense segmentation threshold, we exploit two properties of the pseudo-inconsistency measure. If the size of the dense object is overestimated,

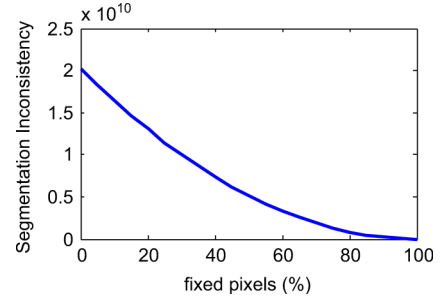


Fig. 2. The relation between the number of fixed pixels and the SIC after 10 iterations, for a Shepp–Logan phantom of size 256×256 and 100 projections.

the residual sinogram will typically be \tilde{A} -inconsistent, which is detected by the pseudo-inconsistency measure. If, on the other hand, the dense object is underestimated, the residual sinogram will always be \tilde{A} -consistent, yet convergence properties of the SIRT algorithm will favor a larger segmented region. We therefore expect that the SIC will reach a minimum when \tilde{B} and B are almost equal.

B. Grey Level Estimation

The SIC concept of the previous subsection is based on the assumption that the grey level ρ of the dense object is known. In practice, this assumption is generally not valid. In this section, we will therefore extend the SIC concept by including the estimation of ρ .

Define $\tilde{\mathbf{t}} = (\tilde{t}_i) \in \mathbb{R}^n$ as the binary image corresponding to the segmentation \tilde{B}

$$\tilde{t}_i = \begin{cases} 0, & \text{if } i \in \tilde{A} \\ 1, & \text{if } i \in \tilde{B}. \end{cases} \quad (9)$$

Then, $\mathbf{W}\tilde{\mathbf{s}} = \rho\mathbf{W}\tilde{\mathbf{t}}$.

Let $\tilde{\rho}$ be a candidate grey value, that will be used to compute the residual sinogram. Assume that $\tilde{B} \subset B$. The residual sinogram is then given by

$$\mathbf{p}_{\tilde{A}} = \mathbf{p} - \tilde{\rho}\mathbf{W}\tilde{\mathbf{t}} = \mathbf{p} - \rho\mathbf{W}\tilde{\mathbf{t}} - (\tilde{\rho} - \rho)\mathbf{W}\tilde{\mathbf{t}} = \mathbf{W}\mathbf{v}_{\tilde{A}} - (\tilde{\rho} - \rho)\mathbf{W}\tilde{\mathbf{t}}. \quad (10)$$

In (10), $\mathbf{W}\mathbf{v}_{\tilde{A}}$ is \tilde{A} -consistent by definition. Therefore, for the residual sinogram to be \tilde{A} -consistent, $\mathbf{W}\tilde{\mathbf{t}}$ must be \tilde{A} -consistent. That is, there must exist an image $\tilde{\mathbf{x}} \in \mathbb{R}^n$ such that $\mathbf{W}\tilde{\mathbf{x}} = \mathbf{W}\tilde{\mathbf{t}}$ and $x_i = 0$ for all $i \in \tilde{B}$. Again, if the number of projections and the region \tilde{B} are both small, it may occur that such an image $\tilde{\mathbf{x}}$ exists, resulting in two different grey levels for which the residual sinogram is consistent. However, as the number of projections is increased and the region \tilde{B} becomes larger, it becomes highly unlikely that the residual sinogram is consistent. This means that if $\tilde{B} \subset B$, the SIC concept defined in the previous subsection can also be used to estimate the grey level ρ , by computing the SIC over all grey levels $\tilde{\rho}$ and choosing the grey level for which it is minimal. If $\tilde{B} \not\subset B$, characterizing the SIC as above becomes more complex. Still, also in this case, the SIC measure can be employed for threshold selection with unknown ρ , as demonstrated in the experimental results of Section IV.

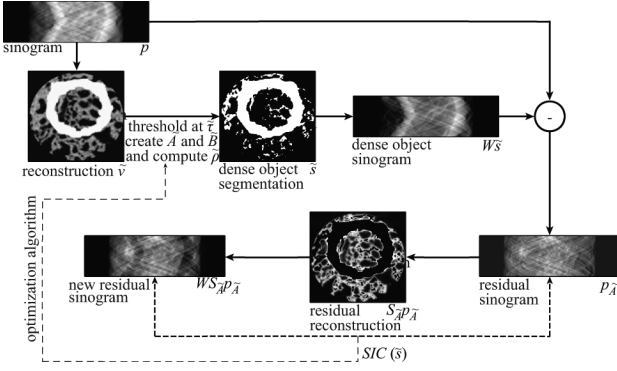


Fig. 3. Schematic overview of the SICM algorithm. In the shown example, the threshold $\tilde{\tau}$ is chosen too low. It can be observed that the residual sinogram $\mathbf{p}_{\tilde{A}}$ is quite different from $\mathbf{WS}_{\tilde{A}}\mathbf{p}_{\tilde{A}}$, indicating a large segmentation inconsistency.

The grey level ρ for which the SIC is minimal can be computed efficiently by exploiting the linearity of the SIRT algorithm

$$\begin{aligned} \text{SIC}(\tilde{\rho}\tilde{\mathbf{t}}) &= \|\mathbf{WS}_{\tilde{A}}(\mathbf{p} - \tilde{\rho}\mathbf{W}\tilde{\mathbf{t}}) - (\mathbf{p} - \tilde{\rho}\mathbf{W}\tilde{\mathbf{t}})\|_{\mathbf{R}} \\ &= \|(\mathbf{WS}_{\tilde{A}}\mathbf{p} - \mathbf{p}) - \tilde{\rho}(\mathbf{WS}_{\tilde{A}}\mathbf{W}\tilde{\mathbf{t}} - \mathbf{W}\tilde{\mathbf{t}})\|_{\mathbf{R}} \\ &= (\mathbf{W}\tilde{\mathbf{t}} - \mathbf{WS}_{\tilde{A}}\mathbf{W}\tilde{\mathbf{t}})\mathbf{R}(\mathbf{W}\tilde{\mathbf{t}} - \mathbf{WS}_{\tilde{A}}\mathbf{W}\tilde{\mathbf{t}})^T \tilde{\rho}^2 \\ &\quad + (\mathbf{WS}_{\tilde{A}}\mathbf{p} - \mathbf{p})\mathbf{R}(\mathbf{WS}_{\tilde{A}}\mathbf{p} - \mathbf{p})^T \\ &\quad + 2(\mathbf{W}\tilde{\mathbf{t}} - \mathbf{WS}_{\tilde{A}}\mathbf{W}\tilde{\mathbf{t}})\mathbf{R}(\mathbf{WS}_{\tilde{A}}\mathbf{p} - \mathbf{p})^T \tilde{\rho}. \end{aligned} \quad (11)$$

The optimal value $\tilde{\rho}_{\text{opt}}$ is found where the first derivative of (11) vanishes. As $\text{SIC}(\tilde{\rho}\tilde{\mathbf{t}})$ is always nonnegative, the minimum of this quadratic polynomial in $\tilde{\rho}$ can be found where the derivative is zero

$$\tilde{\rho}_{\text{opt}} = \frac{(\mathbf{W}\tilde{\mathbf{t}} - \mathbf{WS}_{\tilde{A}}\mathbf{W}\tilde{\mathbf{t}})(\mathbf{p} - \mathbf{WS}_{\tilde{A}}\mathbf{p})^T}{(\mathbf{W}\tilde{\mathbf{t}} - \mathbf{WS}_{\tilde{A}}\mathbf{W}\tilde{\mathbf{t}})(\mathbf{W}\tilde{\mathbf{t}} - \mathbf{WS}_{\tilde{A}}\mathbf{W}\tilde{\mathbf{t}})^T}. \quad (12)$$

C. The SICM Algorithm

The SICM algorithm combines grey level estimation with segmentation inconsistency computation. It uses the segmentation inconsistency found after a fixed number of SIRT iterations as a quantitative measure for the quality of the selected threshold. To find the threshold τ_{opt} where $\text{SIC}(\tilde{\mathbf{s}})$ is minimal, we use the simplex search method of Lagarias [22], which is an unconstrained derivative-free optimization method.

Fig. 3 shows a flowchart of the calculation of the segmentation inconsistency for a certain threshold $\tilde{\tau}$. In Fig. 4, a pseudo code description of a single SIC evaluation is given. It should be noted that for each threshold evaluation three SIRT-reconstructions are required: two for the calculation of $\tilde{\rho}$: $\mathbf{S}_{\tilde{A}}\mathbf{W}\tilde{\mathbf{t}}$ and $\mathbf{S}_{\tilde{A}}\mathbf{p}$ and one for the calculation of $\text{SIC}(\tilde{\mathbf{s}})$: $\mathbf{S}_{\tilde{A}}\mathbf{p}_{\tilde{A}}$. However, because $\mathbf{p}_{\tilde{A}} = \mathbf{p} - \tilde{\rho}\mathbf{W}\tilde{\mathbf{t}}$, we can easily compute $\mathbf{S}_{\tilde{A}}\mathbf{p}_{\tilde{A}}$ by subtracting $\tilde{\rho}\mathbf{S}_{\tilde{A}}\mathbf{W}\tilde{\mathbf{t}}$ from $\mathbf{S}_{\tilde{A}}\mathbf{p}$, thereby reducing the number of SIRT reconstructions per threshold evaluation to two.

```

Input: reconstructed image  $\tilde{\mathbf{v}} = \mathbf{Sp}$  such that  $\mathbf{W}\tilde{\mathbf{v}} \approx \mathbf{p}$ ;
Put  $\tilde{A} := \{a : \tilde{v}_a < \tilde{\tau}\}$ ;
for  $j = 1, \dots, n$ 
    if  $j \in \tilde{A}$  then  $\tilde{t}_j := 0$ ; else  $\tilde{t}_j := 1$ ;
end
 $\mathbf{p}_{\tilde{t}} := \mathbf{W}\tilde{\mathbf{t}}$ ;
 $\mathbf{p}_{\tilde{A}} := \mathbf{p} - \mathbf{p}_{\tilde{t}}$ ;
Compute  $\mathbf{q} = \mathbf{WS}_{\tilde{A}}\mathbf{p}$  and  $\mathbf{r} = \mathbf{WS}_{\tilde{A}}\mathbf{p}_{\tilde{t}}$ ;
 $\tilde{\rho} := \frac{(\mathbf{p}_{\tilde{t}} - \mathbf{r})(\mathbf{p}_{\tilde{t}} - \mathbf{r})^T}{(\mathbf{p}_{\tilde{t}} - \mathbf{r})(\mathbf{p}_{\tilde{t}} - \mathbf{r})^T}$ ;
 $\mathbf{p}'_{\tilde{A}} := \mathbf{p}_{\tilde{A}} - \tilde{\rho}\mathbf{r}$ ;
 $\tilde{\mathbf{s}} := \tilde{\rho}\tilde{\mathbf{t}}$ ;
 $\text{SIC}(\tilde{\mathbf{s}}) := \|\mathbf{p}_{\tilde{A}} - \mathbf{p}'_{\tilde{A}}\|_{\mathbf{R}}$ ;

```

Fig. 4. Pseudo code for the computation of a single SIC measure.

IV. EXPERIMENTS

A. Simulation Studies

Simulation experiments were performed based on various phantom images of size 512×512 , (Fig. 5). All phantom images contain at least one area with a constant, maximal grey level and a background with a continuous set of lower grey level values. Phantom (a) represents a slice of a femur in a gradient-filled surrounding object. The femur contains large as well as small structures (trabeculae). Phantom (b) represents a slice of a human head filled with three differently shaped large objects (implants). Phantom (c) again represents a slice of a human head, but has only a few very small dense objects (dental implants). Finally, Phantom (d) represents a slice of a foam-like structure with a large amount of small metal particles. Each phantom image thus represents a different application and indeed a different segmentation problem (small objects versus large objects, many objects versus only a few, different type of background, etc.)

First, we evaluated the SIC and the relative number of misclassified pixels (rNMP), as a function of the threshold value. The rNMP is defined as the total number of misclassified pixels divided by the total number of pixels belonging to the dense objects. To be a useful measure for threshold selection, the minimum of the SIC should correspond well with the minimum of the rNMP. For all phantoms of Fig. 5, parallel-beam sinograms were simulated using 180 equally spaced projection angles between 0° and 180° . Grey level reconstructions were computed using 300 iterations of the SIRT algorithm described in [16]. Then, for a range of global threshold values, a segmentation was created and the rNMP's and SIC's were computed. For the SIRT reconstructions of the SIC measurements, 300 iterations were used. We also computed the rNMP with the popular Otsu's [13], K-means [23], and expectation maximization (EM) [24] segmentation methods. All these methods require that the user specifies the number of segmentation partitions. This additional prior knowledge is not always available and depends on the scanned object. In these experiments, the number of partitions was different for each phantom image and was chosen such that the segmentation methods from the literature, in general, generated the lowest rNMP. We used six partitions for phantom (a), four for phantom (b), eight for phantom (c), and seven for phantom (d).

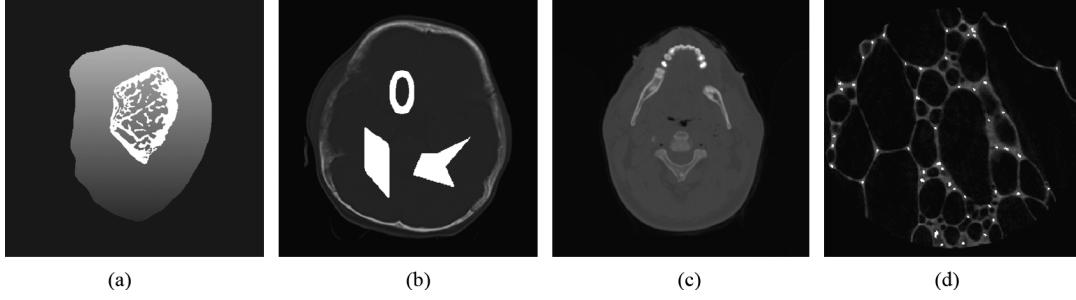


Fig. 5. Simulated phantom images of size 512×512 . (a) Femur of a rat in a surrounding of various densities (b) Slice of human head with differently shaped objects (c) Slice of human head with a few small dental fillings (d) Foam object with metal marker particles that are often used for image registration.

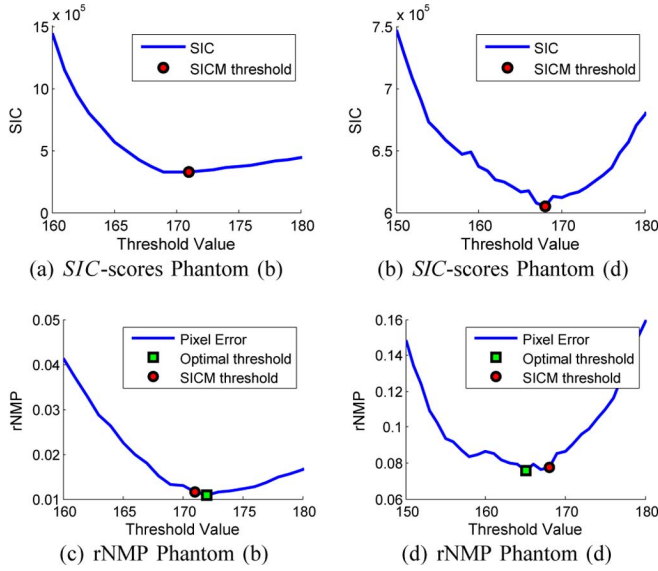


Fig. 6. (a), (b) SIC-scores as a function of the threshold values (c), (d) relative number of misclassified pixels (rNMP) as a function of the threshold values. For reliable threshold estimation, the minimum of the SIC-score should lie close to the minimum of the rNMP.

Fig. 6 shows the SIC-score and rNMP graphs for phantom images (b) and (d). It can be observed that the distance between the minimum of the SIC-curve and the minimum of the rNMP-curve (the squares and circles in Fig. 6, respectively) is very small. In Table I, the results are shown numerically. The thresholds suggested by the SICM method are good approximations of the optimal thresholds, i.e., the thresholds for which the rNMP is minimal (when using any global threshold technique, a lower number is not possible). Furthermore, the NMP of the SICM method are close to the optimal scenario. For phantoms (a), (b), and (d), the other segmentation methods result in a substantially higher number of misclassified pixels, whereas for phantom (c) the results of SICM are comparable to the best performing alternatives.

In the following three experiments we investigated the effect of various tomographic conditions on the segmentation performance of the SICM algorithm. The resulting rNMP was compared with that of the optimal threshold and those of the other previously used segmentation methods.

- Firstly, the contrast between the maximal density of the background and the density of the continuous dense object was varied. This contrast is defined as $\rho / \max_{i \in A} v_i$.

TABLE I
NUMERICAL RESULTS FOR THE FIRST SIMULATION EXPERIMENT.
RNMP: RELATIVE NUMBER OF MISCLASSIFIED PIXELS

Phantom	Threshold		rNMP * 100				
	Optimal	SICM	Optimal	SICM	Otsu	Kmeans	EM
a	190.38	195.22	4.11	4.57	6.87	6.92	7.18
b	172.31	171.80	1.08	1.15	1.55	1.76	17.44
c	208.72	205.30	10.20	12.25	89.80	11.22	10.71
d	164.80	167.40	7.55	7.55	98.83	9.08	22.39

As this contrast decreases, the grey level distribution of the dense object in the tomographic reconstruction overlaps more and more with that of the background. We therefore expect that the accuracy of the segmentations will also decrease.

Fig. 7 shows the rNMP as a function of the phantom contrast for the optimal global threshold, the SICM thresholding method and the other clustering algorithms. In general, the SICM curve approximates the curve of the optimal threshold, whereas the other methods are much less stable and often do not provide a good approximation.

- Secondly, the number of projection angles was reduced. In practice, such a reduction decreases both scan time and radiation dose. However, it is expected that for a low angle count the results will be unreliable as the system will become more and more underdetermined.

The results, shown in Fig. 8, indeed confirm that the rNMP increases drastically as the number of projection angles is decreased beyond a certain minimum number that is specific for each phantom. However, the SICM method still leads to significantly more accurate segmentations than other methods.

- Finally, Poisson noise with a varying source intensity was applied to the sinograms. The intensity of the noise is related to the measured detector count when there is no object between the source and the detector. This determines the signal-to-noise ratio. In this experiment, therefore, we effectively simulated low dose, low photon count scans. The results, shown in Fig. 9, indicate that the noise level has little effect on the ability of the SICM algorithm to estimate accurate global threshold levels.

It should be noted that in all reported experiments, both on simulated and on experimental data, 300 iterations were computed for each SIRT-reconstruction that is part of the SIC computations. This number is based on empirical findings. If the number of iterations is too low, the pseudo-inconsistency (computed by a fixed number of SIRT iterations) does not match well

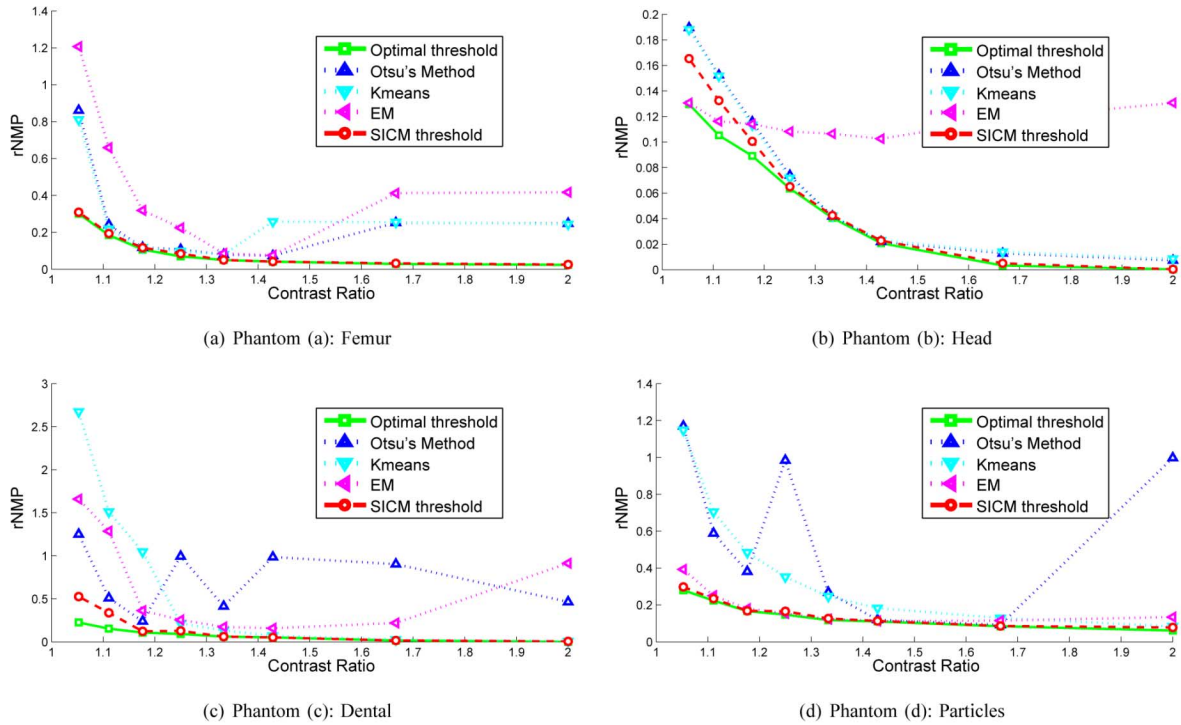


Fig. 7. The relative number of misclassified pixels (rNMP) as a function of the contrast between the background and the dense object in the original image phantom. On the y -axis $\rho / \max_{i \in A} v_i$ is shown.

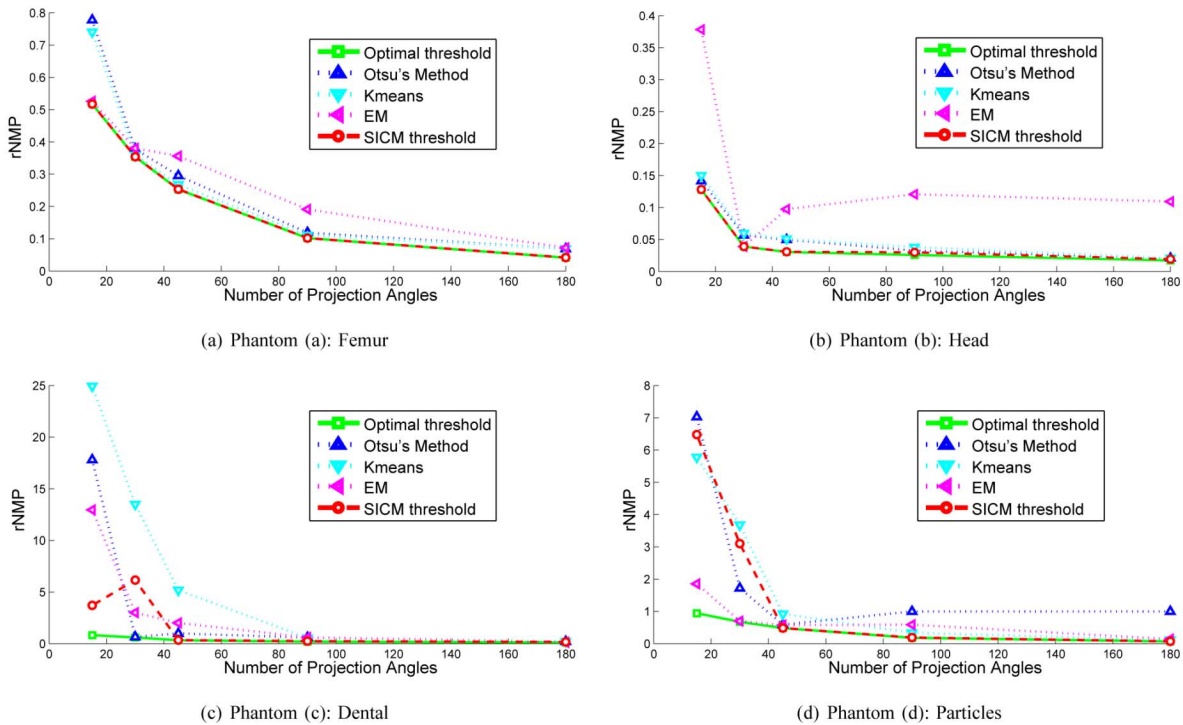


Fig. 8. The relative number of misclassified pixels (rNMP) as a function of the number of projection angles.

with the true inconsistency. On the other hand, if the number of iterations is too high, we can no longer make use of the SIRT-reconstruction technique to determine if $\tilde{B} \subset B$. In future work we will therefore look into other iterative reconstruction techniques that can be used to estimate the inconsistency of a sinogram.

B. Experimental Data Studies

We also applied the proposed algorithm on μ CT data. Fig. 10 shows a reconstructed image of a cochlear implant in surrounding tissue, acquired with a SkyScan 1076 μ CT scanner using 360 projection angles at a detector resolution of

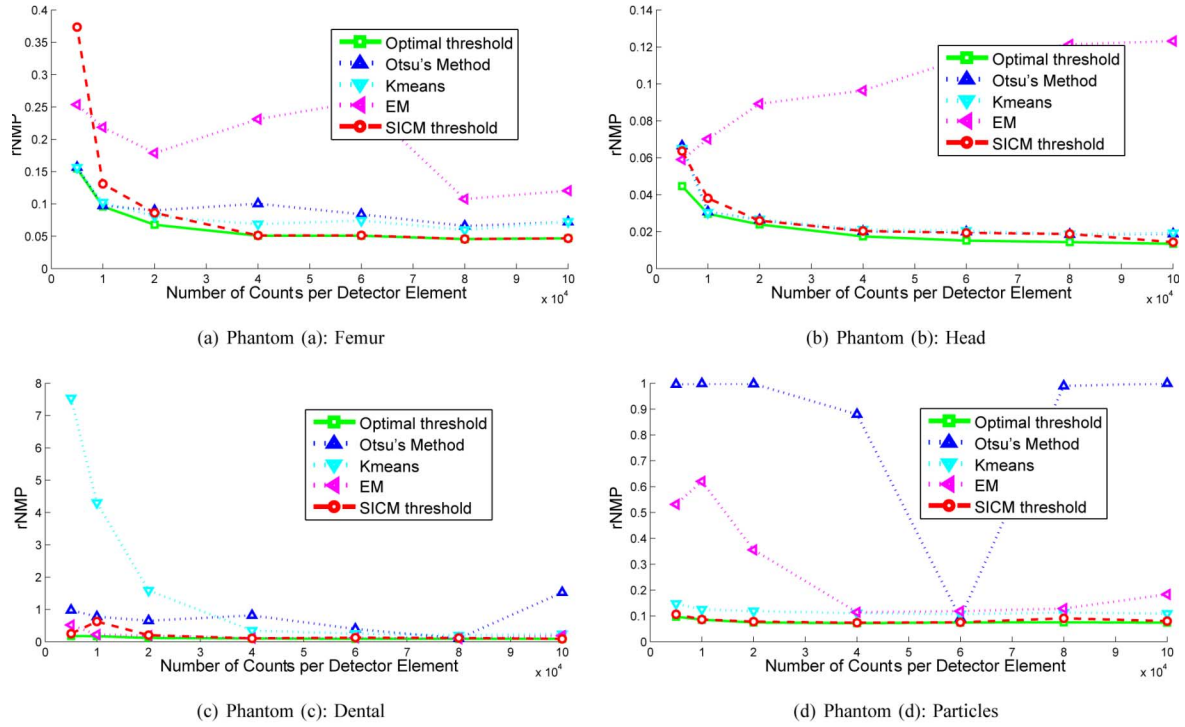


Fig. 9. The relative number of misclassified pixels (rNMP) as a function of the number of counts per detector pixel.

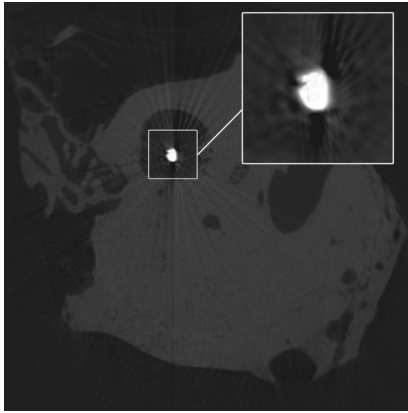






Fig. 10. SIRT reconstruction of a cochlear implant in surrounding tissue.

TABLE II
RESULTS FOR THE COCHLEAR IMPLANT EXPERIMENT

Otsu		SICM	
			
360 angles ground truth	72 angles rNMP: 0.4452	360 angles rNMP: 0.0750	72 angles rNMP: 0.2074

12 μm . The standard SkyScan NRecon software package has been used to correct for ring- and beam-hardening artefacts.

The goal is to accurately locate the small cochlear implant. However, validating the quality of these segmentations is difficult. Notice that in Table II the segmentations provided by

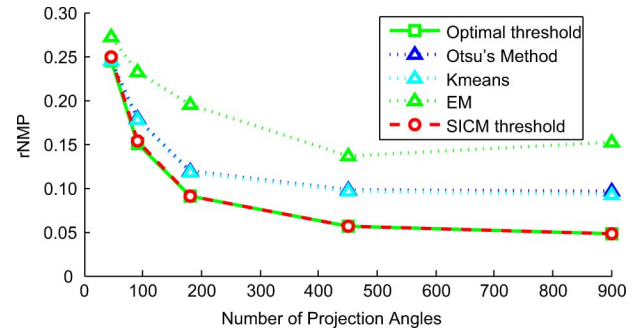


Fig. 11. Results for the mandible experiment.

Otsu's method and SICM with 360 projection angles are very similar, indicating that for this high number of angles the segmentation result does not depend strongly on the particular segmentation method. We therefore consider Otsu's segmentation with 360 projection angles as the ground truth image to which we compare the SICM and Otsu's segmentation using only 72 projection angles. In Table II, it is clear that, for this example, if the number of projection angles is lowered, an SICM segmentation is much more accurate than a classical Otsu segmentation.

Fig. 11 shows an FBP reconstruction of a slice through a human mandible. This image was recorded using a SkyScan 1173 μCT scanner using 900 projection angles at a detector resolution of 50 μm . Important in this application is the accurate segmentation of the cortical and trabecular bone. We again apply the SICM algorithm and the other standard methods from the literature. We compared our results by computing the rNMP with respect to a manually segmented image, shown in Fig. 12(b). The experiments were repeated using a decreasing number of

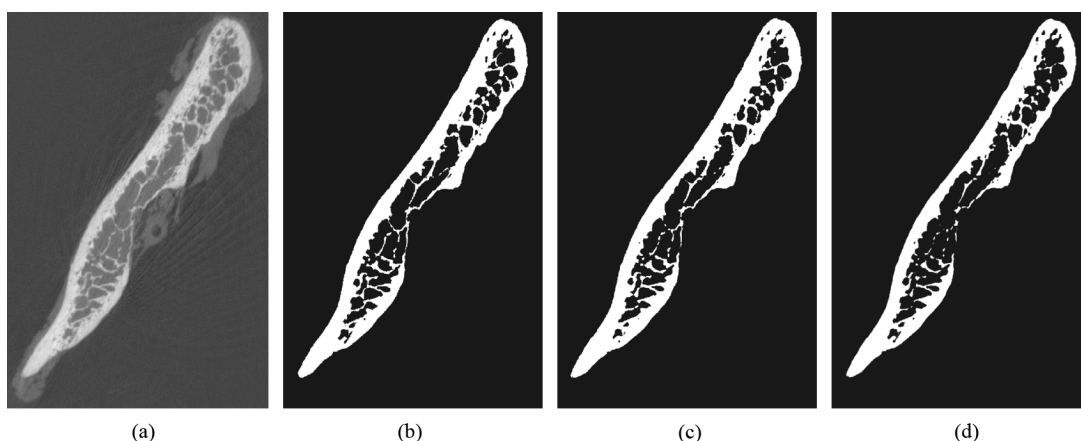


Fig. 12. Results for the mandible experiment. (a) FBP reconstruction, 900 angles. (b) Manual segmentation, 900 angles. (c) SICM segmentation, 450 angles. (d) Kmeans segmentation, 450 angles.

projection angles, simulating low dose scans. Fig. 12 clearly shows that the SICM method outperforms the other methods. Fig. 12(c) and (d), respectively, shows the SICM and Kmeans segmentation using 450 projection angles. The SICM segmentation is an accurate approximation of the manual segmentation and indeed of the reconstructed FBP image as it locates almost every detail of the trabeculae present.

V. DISCUSSION AND CONCLUSION

Accurate image segmentation is a difficult problem, in particular if one needs to segment objects that are very detailed or very small with respect to the image resolution. This is especially true in tomography if the image on which the segmentation is based is polluted by reconstruction errors or artefacts. Segmentation errors generally occur at the edges of the object where a hard edge is usually blurred in the reconstruction image. Popular methods, such as global clustering methods (Otsu, Kmeans, EM) or local neighborhood methods (region growing, watershed segmentation) all have different strategies to define where exactly the object edges are located, making it difficult to measure the accuracy of these segmentations in practice. Fortunately, tomography inherently provides a way to counter this problem as projection data is available that can be used to improve or optimize standard segmentation techniques.

In this paper, we have presented a novel method for finding a global threshold to accurately locate dense objects in a continuous surrounding in a tomographic reconstruction. Contrary to existing methods, the SICM method, is not only based on the reconstructed image, but also on the available projection data. For an optimal segmentation, the residual sinogram (i.e., the sinogram of the part of the image that does not belong to the dense object) is consistent. The inconsistency can be measured by applying a linear iterative reconstruction technique (such as SIRT) to the pixels not belonging to the dense object and by comparing the forward projection of this reconstruction to the residual sinogram.

Results have been generated for both simulated and experimental data and the SICM algorithm generally finds a good approximation of the optimal global threshold. In a large majority of the experiments performed, the SICM method outperforms

other thresholding methods. We have focused on μ CT for experimental data since it is a convenient and open research platform, where raw projection data is readily available. The same technique can be applied to other tomography fields (e.g., medical) as well.

One downside of the proposed method is its computational cost. Whereas methods that work only on the reconstruction image are typically very fast, the SICM method has an optimization routine where each function evaluation requires two SIRT reconstructions. For a full optimization of a 512×512 image with 180 projection angles, on a system with a modern 3 GHz Intel CPU, the typical computation time for a single SIC evaluation was 10 min and 221 min for a full optimization. To drastically lower the computation time, we resorted to GPU programming using the CUDA programming language. In our system with an NVIDIA GeForce GTX 480 GPU, on average we reduced the computation time for a single SIC evaluation to 13 s and for a full optimization to 4.5 min.

Future work will focus on combining the SIC approach with other, more advanced segmentation algorithms such as local thresholding.

REFERENCES

- [1] T. M. Link, W. Berning, S. Scherf, U. Joosten, A. Joist, K. Engelke, and H. E. Daldrup-Link, "CT of metal implants : Reduction of artifacts using an extended CT scale technique," *J. Comput. Assist. Tomogr.*, vol. 24, no. 1, pp. 165–172, 2000.
- [2] C. Lemmens, D. Faul, J. Hamill, S. Stroobants, and J. Nuyts, "Suppression of metal streak artifacts in CT using a MAP reconstruction procedure," in *IEEE Nucl. Sci. Symp. Conf. Rec.*, Oct. 2006, vol. 6, pp. 3431–3437.
- [3] W. Veldkamp, R. Joemai, and J. Geleijns, "Automated segmentation and interpolation in sinograms for metal artifact suppression in CT," *Med. Phys.*, vol. 36, no. 6, pp. 2443–2443, Jun. 2009.
- [4] B. R. Whiting, K. T. Bae, and M. W. Skinner, "Cochlear implants: Three-Dimensional localization by means of coregistration of CT and conventional radiographs," *Radiology*, vol. 221, pp. 543–549, 2001.
- [5] G. Schoonenberg, R. Florent, P. Lelong, O. Wink, D. Ruijters, J. Carroll, and B. t. H. Romeny, "Projection-based motion compensation and reconstruction of coronary segments and cardiac implantable devices using rotational X-ray angiography," *Med. Image Anal.*, vol. 13, pp. 785–792, 2009.
- [6] S. Yang, G. Wang, M. W. Skinner, J. T. Rubinstein, and M. W. Vannier, "Localization of cochlear implant electrodes in radiographs," *Med. Phys.*, vol. 27, no. 4, pp. 775–777, Apr. 2000.

- [7] *Mechanical Testing of Bone and the Bone—Implant Interface*, Y. An and R. A. Draughn, Eds. Boca Raton, FL: CRC, 2000.
- [8] W. L. Roque, A. C. A. de Souza, and D. X. Barbieri, "The euler-poincaré characteristic applied to identify low bone density from vertebral tomographic images," *Revista Brasileira De Reumatologia*, vol. 49, no. 2, 1983.
- [9] M. J. Eichler, C. H. Kim, R. Müller, and X. E. Guo, "Impact of thresholding techniques on micro-ct image based computational models of trabecular bone," *ASME Adv. Bioeng.*, vol. 48, pp. 215–216, Sep. 2000.
- [10] C. A. Glasbey, "An analysis of histogram-based thresholding algorithms," *Graph Models Image Process.*, vol. 55, no. 6, pp. 532–537, Nov. 1993.
- [11] A. Rosenfeld and P. Torre, "Histogram concavity analysis as an aid in threshold selection," *IEEE Syst., Man, Cybern. C*, vol. 13, pp. 231–235, Mar./Apr. 1983.
- [12] T. W. Ridler and S. Calvard, "Picture thresholding using an iterative selection method," *IEEE Syst., Man, Cybern. C*, vol. 8, no. 8, pp. 630–632, Aug. 1978.
- [13] N. Otsu, "A threshold selection method from gray level histograms," *IEEE Syst., Man, Cybern. C*, vol. 9, no. 1, pp. 62–66, Mar. 1979.
- [14] X. Yu and J. Ylaaski, "A new algorithm for image segmentation based in region growing and edge detection," in *Proc. Int. Symp. Circuits Syst.*, 1991, vol. 1, pp. 516–519.
- [15] S. Beucher and F. Meyer, "The morphological approach to segmentation: The watershed transformation," in *Dougherty E. R. Ed., Ser. Mathematical Morphology in Image Processing*. New York: Marcel Dekker, 1993.
- [16] J. Gregor and T. Benson, "Computational analysis and improvement of SIRT," *IEEE Trans. Med. Imag.*, vol. 27, no. 7, pp. 918–924, Jul. 2008.
- [17] K. J. Batenburg and J. Sijbers, "Optimal threshold selection for tomogram segmentation by projection distance minimization," *IEEE Trans. Med. Imag.*, vol. 28, no. 5, pp. 676–686, May 2009.
- [18] K. J. Batenburg and J. Sijbers, "Adaptive thresholding of tomograms by projection distance minimization," *Pattern Recognit.*, vol. 42, no. 10, pp. 2297–2305, 2009.
- [19] W. van Aarle, J. Batenburg, and J. Sijbers, "Threshold selection for segmentation of dense objects in tomograms," in *Int. Symp. Visual Computing, Ser. Lecture Notes Comput. Sci.*, G. Bebis, Ed., Berlin, Germany, Dec. 2008, vol. 5358, pp. 700–709.
- [20] D. Ludwig, "The radon transform on euclidean space," *Comm. Pure Appl. Math.*, vol. 19, pp. 49–81, 1966.
- [21] S. Helgason, *The Radon Transform*. Boston, MA: Birkhauser, 1980.
- [22] J. Lagarias, J. A. Reeds, M. H. Wright, and P. E. Wright, "Convergence properties of the nelder-mead simplex method in low dimensions," *SIAM J. Optim.*, vol. 9, no. 1, pp. 112–147, 1998.
- [23] T. Kanungo, D. Mount, N. Netanyahu, C. Piatko, R. Silverman, and A. Wu, "An efficient k-means clustering algorithm: Analysis and implementation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 7, pp. 881–892, Jul. 2002.
- [24] C. Carson, S. Belongie, H. Greenspan, and J. Malik, "Blobworld: Image segmentation using expectation-maximization and its application to image querying," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 8, pp. 1026–1038, Aug. 2002.