Computer Vision and Image Understanding 117 (2013) 306-318

Contents lists available at SciVerse ScienceDirect

ELSEVIER

Computer Vision and Image Understanding

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



Dynamic angle selection in binary tomography

K. Joost Batenburg^{a,b,*}, Willem Jan Palenstijn^b, Péter Balázs^c, Jan Sijbers^b

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

^b Vision Lab, University of Antwerp, Universiteitsplein 1, B-2610 Wilrijk, Belgium

^c Department of Image Processing and Computer Graphics, University of Szeged, Árpád tér 2, H-6720 Szeged, Hungary

ARTICLE INFO

Article history: Received 1 November 2011 Accepted 3 July 2012 Available online 8 December 2012

Keywords: Binary tomography Discrete tomography Angle selection

ABSTRACT

In this paper, we present an algorithm for the dynamic selection of projection angles in binary tomography. Based on the information present in projections that have already been measured, a new projection angle is computed, which aims to maximize the information gained by adding this projection to the set of measurements. The optimization model used for angle selection is based on a characterization of solutions of the binary reconstruction problem, and a related definition of information gain. From this formal model, an algorithm is obtained by several approximation steps. Results from a series of simulation experiments demonstrate that the proposed angle selection scheme is indeed capable of finding angles for which the reconstructed image is much more accurate than for the standard angle selection scheme. © 2012 Elsevier Inc. All rights reserved.

1. Introduction

Tomography deals with the reconstruction of an image from its projections, acquired along a range of angles. The Inverse Radon Transform provides a closed-form inversion formula for this reconstruction problem, provided that projections are available for *all* angles. Although this assumption is clearly not satisfied in practice, an accurate reconstruction can be computed if a large number of projections are available, for a full angular range, by using the well-known Filtered Backprojection algorithm [1,2].

In several applications of tomography, only few projections can be acquired. Such reconstruction problems are known as *limiteddata problems*. In electron tomography, for example, the electron beam damages the sample, limiting the number of projections that can be acquired [3]. In industrial tomography for quality assurance, cost considerations impose limitations on the duration of a scan, and thereby on the number of projections.

Applying classical reconstruction algorithms such as Filtered Backprojection to limited-data problems often results in inferior reconstruction quality. Several approaches have been proposed to overcome these problems, by incorporating various forms of *prior knowledge* about the object in the reconstruction algorithm. Recently, advances in the field of Compressed Sensing have demonstrated high potential in obtaining a reduction of the number of required projection images by exploiting sparsity of the image with respect to a certain set of basis functions [4–8]. Following a

E-mail address: joost.batenburg@cwi.nl (K.J. Batenburg).

different approach, the field of *discrete tomography* focuses on the reconstruction of images that consist of a small, discrete set of gray values [9,10]. By exploiting the knowledge of these gray values in the reconstruction algorithm, it is often possible to compute accurate reconstructions from far fewer projections than required by classical "continuous" tomography algorithms [11,12].

When reconstructing an image from a small set of projections (i.e., less than 20), the *particular set* of projection angles can have a large influence on the quality of the reconstruction. In [13,14], it was shown that the choice of the projection angles can have a crucial influence on the reconstruction quality in *binary tomography*, i.e., discrete tomography based on just two gray levels. The authors also presented algorithms to identify optimal projection angles based on a *blueprint* image, which is known to be similar to the scanned object. For more general grayscale tomography, a framework was recently proposed for optimizing the acquisition of projections, based on certain prior knowledge on the object to be scanned [15].

These findings naturally lead to the question if "optimal" angles can also be selected in cases where no blueprint image is available. As the optimal angles depend on the scanned object, they can certainly not be selected prior to the scanning procedure. Instead, we consider an *on-line* variant of the problem, where projections are measured one-by-one, and the new angle is selected based on the information present in the currently available projections.

In this paper, we present an algorithm for angle selection in binary tomography that is based on a concise model of the available projection information and prior knowledge of the binary character of the unknown object. The algorithm depends on two key ingredients: sampling of the set of images that adhere to the currently available projections, and determining the amount of

^{*} Corresponding author at: Centrum Wiskunde & Informatica, Science Park 123, NL-1098 XG Amsterdam, The Netherlands.

^{1077-3142/\$ -} see front matter © 2012 Elsevier Inc. All rights reserved. http://dx.doi.org/10.1016/j.cviu.2012.07.005

information that can be gained by adding a particular angle to the set of measurements.

This paper is structured as follows. In Section 2, we introduce a formal model for the tomography problem, as well as several related concepts, which are necessary to define the angle selection problem. Although our formulation of the angle selection problem characterizes the optimal angle, it does not directly lead to an algorithm for computing this angle, which requires several approximations. Section 3 describes how the current set of binary solutions to the reconstruction problem can be approximated by a sampling procedure, leading to so-called surrogate solutions. In Section 4, we describe how the information gain can be approximated for a particular candidate angle. These two approximation steps are then combined into an algorithm for the angle selection problem. which is presented in Section 5. An example is discussed in Section 6, where the concepts involved in the angle estimation algorithm are related to actual images involved in the computations. Section 7 describes a series of simulation experiments that have been performed to investigate how the reconstructions obtained by the proposed angle selection scheme compare to several alternative angle selection strategies. In the experimental results, present in Section 8, it is shown that the proposed method is indeed capable of finding angles for which the reconstructed image is much more accurate than for the standard angle selection scheme. The approach and experimental findings are discussed in Section 9, followed by conclusions in Section 10.

2. Notation and concepts

Our description is restricted to the reconstruction of twodimensional images from one-dimensional projections, but can be generalized to higher-dimensional settings in a straightforward manner. The reconstructed image is represented on a rectangular grid of size $n = w \times h$.

For setting up the tomography model, we make the idealized assumption that the unknown *original object* is a binary image that can also be represented on this grid, even though the proposed algorithm can still be used if the grid assumption is not satisfied. The unknown original image can be represented by a vector $\bar{\boldsymbol{v}} = (\bar{\boldsymbol{v}}_i) \in \{0, 1\}^n$, where the entries $\bar{\boldsymbol{v}}_i$ correspond to the pixel values of the reconstruction.¹

Projections are measured as sets of detector values for various angles, rotating around the object. For each angle, the detectors register a parallel projection of the object. The finite set $\Theta = \{\theta_1, \ldots, \theta_d\}$ of angles for which projection data has been measured is gradually expanded: each time, an angle is selected based on the projections available so far, and the projection corresponding to that angle is added to the set of measurements.

We denote the number of detector values for each projection by k. For any angle $\theta \in [0, \pi)$, the projection process in tomography, assuming noiseless measurements, can be modeled as a linear operator $\mathbf{W}^{(\theta)}$, which maps the image $\bar{\boldsymbol{v}}$ to the vector $\boldsymbol{p}^{(\theta)}$ of measured data:

$$\boldsymbol{p}^{(\theta)} = \boldsymbol{W}^{(\theta)} \, \bar{\boldsymbol{v}}.\tag{1}$$

The $k \times n$ matrix $\mathbf{W}^{(\theta)} = (w_{ij}^{(\theta)})$ is called the *projection matrix* for angle θ . The entry $w_{ij}^{(\theta)}$ determines the weight of the contribution of pixel *j* to measurement *i*, which represents the area of the intersection between the image pixel and a strip that intersects with the image and projects onto the *i*th detector pixel. These concepts are illustrated in Fig. 1.



Fig. 1. Discretized strip projection model.

From this point on, we assume that each matrix $\mathbf{W}^{(\theta)}$ has the property that $\sum_{i=1}^{k} w_{ij}^{(\theta)} = 1$ for all j = 1, ..., n. Assuming that the image is completely covered within the field of view of the detector, this property is satisfied for the strip projection model that we use here, as the total pixel weight for each projection angle is equal to the area of a pixel, which is 1. For most other projection models commonly used in tomography, such as the line model, where the weight of a pixel is determined by the length of its intersection with a line, this property is approximately satisfied, but not always exactly.

For a *set* of projection angles $\Theta = \{\theta_1, \theta_2, ..., \theta_d\}$, the projection matrix $\mathbf{W}^{(\Theta)}$ consists of a stack of projection matrices for the individual angles, resulting in measurements $\mathbf{p}^{(\Theta)}$ of the form

$$\boldsymbol{p}^{(\Theta)} = \begin{pmatrix} \boldsymbol{p}^{(\theta_1)} \\ \vdots \\ \boldsymbol{p}^{(\theta_d)} \end{pmatrix} = \begin{pmatrix} \boldsymbol{W}^{(\theta_1)} \\ \vdots \\ \boldsymbol{W}^{(\theta_d)} \end{pmatrix} \bar{\boldsymbol{v}}.$$
 (2)

For a given projection matrix $\boldsymbol{W}^{(\Theta)}$ and projection data $\boldsymbol{p}^{(\Theta)} = \boldsymbol{W}^{(\Theta)} \bar{\boldsymbol{v}}$, let $S_{\boldsymbol{W}^{(\Theta)}}(\boldsymbol{p}^{(\Theta)}) = \{\boldsymbol{x} \in \mathbb{R}^n : \boldsymbol{W}^{(\Theta)} \boldsymbol{x} = \boldsymbol{p}^{(\Theta)}\}$ be the set of all real-valued solutions corresponding with the projection data, and let $\overline{S}_{\boldsymbol{W}^{(\Theta)}}(\boldsymbol{p}^{(\Theta)}) = S_{\boldsymbol{W}^{(\Theta)}}(\boldsymbol{p}^{(\Theta)}) \cap \{0,1\}^n$ be the set of *binary solutions* of the system. We focus on the case where the total number of measurements m = dk is small with respect to n, such that the real-valued reconstruction problem is severely underdetermined.

For any two binary images $\bar{\mathbf{x}}, \bar{\mathbf{y}} \in \{0, 1\}^n$, define the *image distance* by $\|\bar{\mathbf{x}} - \bar{\mathbf{y}}\|_2$. For any set $V \subset \{0, 1\}^n$, define the *diameter* of this set by diam $(V) = \max\{\|\bar{\mathbf{x}} - \bar{\mathbf{y}}\|_2 \mid \bar{\mathbf{x}}, \bar{\mathbf{y}} \in V\}$. If the diameter of V is small, all images in the set must be quite similar, whereas a large diameter indicates that strong variations occur within the set.

We now turn to the problem of angle selection. Let $\Theta = \{\theta_1, \ldots, \theta_d\}$ be the current set of *d* directions for which projection data $\mathbf{p}^{(\Theta)} = \mathbf{W}^{(\Theta)} \bar{\mathbf{v}}$ of the unknown original image $\bar{\mathbf{v}}$ have already been measured. The task is now to select the next angle for which a projection will be measured in such a way that as much "information" as possible is gained by the new measurement.

A principal obstacle in constructing a formal model of this problem is the fact that the original image can be any element of the set $\overline{S}_{W^{(\Theta)}}(\boldsymbol{p}^{(\Theta)})$, unless additional prior knowledge is available about the original image. A projection angle that yields much information for a particular element of this set, might yield very little information for other images in the same set.

To quantify the amount of information that is contained in a certain set of projections $\boldsymbol{p}^{(\Theta)}$, we propose to use $\operatorname{diam}\left(\overline{S}_{\boldsymbol{W}^{(\Theta)}}(\boldsymbol{p}^{(\Theta)})\right)$: the diameter of the set of binary solutions that adhere to all current projections. Note that even when large switching components exist for the given set of projection

¹ Throughout this paper, we indicate a vector or a scalar that has a binary domain by putting a bar above its symbol.

directions (i.e., pairs of binary images that have the same projections in all directions), diam $(\overline{S}_{W^{(\Theta)}}(p^{(\Theta)}))$ may still be small, or even 0, as p may uniquely determine a binary solution. For any binary image \bar{x} and set of angles $\Theta = \{\theta_1, \theta_2, \ldots\}$, define the *uncertainty* of (\bar{x}, Θ) by

$$U(\bar{\boldsymbol{x}}, \boldsymbol{\Theta}) = \operatorname{diam}\left(\bar{\boldsymbol{S}}_{\boldsymbol{W}^{(\boldsymbol{\Theta})}}(\boldsymbol{W}^{(\boldsymbol{\Theta})}\bar{\boldsymbol{x}})\right). \tag{3}$$

The uncertainty corresponds with the diameter of the set of binary images that have the same projections as \bar{x} for all angles in Θ . Similarly, we define the *information gain* of $(\bar{x}, \Theta, \theta)$ by

$$G(\bar{\mathbf{x}}, \Theta, \theta) = U(\bar{\mathbf{x}}, \Theta) - U(\bar{\mathbf{x}}, \Theta \cup \{\theta\}), \tag{4}$$

which can be used as a measure for the information gained by measuring the projection for angle θ , if the projections for all angles in Θ are already available and \bar{x} is the original object. Clearly, the information gain is always nonnegative, and is zero for any θ that is already in Θ . We extend this concept to the *mean information gain* of a *set* of binary images $V \subset \{0, 1\}^n$, defined by

$$G(V, \Theta, \theta) = \sqrt{\frac{1}{|V|} \sum_{\bar{\mathbf{x}} \in V} G(\bar{\mathbf{x}}, \Theta, \theta)^2}.$$
(5)

We have now defined all concepts and notation required for formulating the angle estimation problem:

Problem 1. Let $\Theta = \{\theta_1, \dots, \theta_d\}$ be the current set of *d* directions for which projection data $\mathbf{p}^{(\Theta)} = \mathbf{W}^{(\Theta)} \bar{\mathbf{v}}$ of the unknown original image $\bar{\mathbf{v}}$ have already been measured. Find

$$\theta_{d+1} = \arg \max_{\theta \in [0,\pi]} G\left(\overline{S}_{\boldsymbol{W}^{(\Theta)}}(\boldsymbol{p}^{(\Theta)}), \boldsymbol{\Theta}, \theta\right).$$

This problem can be interpreted as follows: we seek the new projection angle θ_{d+1} , such that the total information gained by adding this angle, over all binary images that adhere to the current set of known projections, is maximized.

Although Problem 1 effectively captures all concepts involved in the angle estimation problem, its formulation is not suitable for direct translation into an algorithm. It is well known that certain variants of the problem of reconstructing binary images from more than two projections are NP-hard [16], and for the strip model we are not aware of any feasible approach to enumerate the set of binary solutions for a large-scale reconstruction problem. Such an enumeration is needed at two levels: (1) to compute the total gain over all binary solutions that adhere to the available projections, and (2) to compute the diameter of the solution set within the computation of the gain for an individual image.

To construct an algorithm based on the concepts in Problem 1, we propose to use three approximations. Firstly, the summation over all binary solutions within the computation of the total gain is replaced by a summation over a set of *surrogate solutions*, which are real-valued solutions to the reconstruction problem for which all pixel values are in the interval [0,1]. Secondly, within the computation of the gain for an individual image, the diameter of the binary solution set is approximated by an *upper bound* of the diameter, which was recently proposed in [17]. Based on the experimental results presented in that paper, we expect that this upper bound can be used effectively as a substitute for the true diameter. Finally, the continuous domain of the candidate angle θ_{d+1} is replaced by a finite discretization, with an angular step of 1 degree.

In the next two sections, we provide a description of these approximation steps, which are then combined into a complete algorithm for angle selection in Section 5.

3. Surrogate solutions

To approximate the evaluation of $G(\bar{S}_{\boldsymbol{W}^{(\Theta)}}(\boldsymbol{p}^{(\Theta)}), \Theta, \theta)$, we resort to computing the mean information gain for a set of *surrogate solutions*, which are not necessarily binary. The surrogate solutions are real-valued solutions of the reconstruction problem for which all entries have values in the interval [0, 1]. These surrogate solutions are then used as samples to represent the true set $\bar{S}_{\boldsymbol{W}^{(\Theta)}}(\boldsymbol{p}^{(\Theta)})$ of binary solutions.

The starting point for generating a surrogate solution is a gray level *template image* $\mathbf{c} \in [0,1]^n$, which is randomly generated from a given parameterized family of gray level images. The template image is used as a starting point for an iterative algorithm that computes a real-valued surrogate solution $\boldsymbol{v} \in S_{\boldsymbol{W}^{(\Theta)}}(\boldsymbol{p}^{(\Theta)}) \cap [0,1]^n$. The generation of surrogate images can be considered as sampling based on a *prior distribution*, where the particular family of template images determines the prior distribution. Features present in the template images are partially preserved within the corresponding surrogate solutions, and therefore control the approximation of the current set of binary solutions.

In preliminary experiments, we also considered the case that the gray levels for surrogate solutions are not constrained. In that case, surrogate solutions are sometimes formed with values that are considerably greater than 1, or smaller than 0, and therefore far away from any binary image (which they are supposed to represent). This behavior is effectively avoided by including the gray level constraint.

In our implementation, the basic algorithm for computing a real-valued surrogate solution is the iterative SIRT algorithm [18], defined as follows. Let $\mathbf{v}^{(0)} = \mathbf{c}$. For q = 1, 2, ..., let $\mathbf{r}^{(q)} = \mathbf{p}^{(\Theta)} - \mathbf{W}^{(\Theta)} \mathbf{v}^{(q-1)}$ be the projection difference before the *q*th iteration. In each iteration *q*, the current reconstruction $\mathbf{v}^{(q-1)}$ is updated, yielding a new reconstruction $\mathbf{v}^{(q)}$, as follows:

$$\nu_{j}^{(q)} = \nu_{j}^{(q-1)} + \frac{1}{\sum_{i=1}^{m} w_{ij}^{(\Theta)}} \sum_{i=1}^{m} \frac{w_{ij}^{(\Theta)} r_{i}^{(q)}}{\sum_{j=1}^{n} w_{ij}^{(\Theta)}}.$$
 (6)

with *m* the total number of detector measurements. It can be shown that for a consistent system of equations, the SIRT-algorithm as described in Eq. (6) converges to the solution $\tilde{\boldsymbol{\nu}}$ that is closest to the initial image $\boldsymbol{\nu}^{(0)}$ with respect to a certain norm (a weighted sum of squares, see [18]). Here, we use an adaptation of this algorithm which confines the solution to the interval [0, 1]. After each iteration, the pixels that have a negative value are set to 0, and those that are greater than 1 are set to 1. As this truncation operation is the projection onto a convex set, the resulting algorithm still converges to a solution of the tomography problem, yet it is not guaranteed to be the solution that is closest to the initial image.

We remark that the real-valued equation system $W^{(\Theta)}x = p$ is severely underdetermined, such that the algorithm can typically reach a solution that still resembles the template image. As a consequence, allowing sufficient variation within the set of template images will result in variations in the set of surrogate solutions formed.

The surrogate solutions are typically *not* binary images, and it is not at all clear that by sampling these grayscale images, the properties of the set of binary solutions can be "represented" with sufficient detail to allow for effective angle selection. Yet, the experimental results presented in Section 8 suggest that for the phantoms considered, the mean information gain for a set of surrogate solutions can often approximate the information gain with respect to the actual unknown object quite accurately. The example that is presented in Section 6 shows that particular properties of the binary solution set, such as the presence of binary *switching* *components* can indeed be observed in the set of surrogate solutions.

4. Computing an upper bound for the diameter

Once a surrogate solution **v** has been computed, the information gain needs to be computed for each of the candidate angles. We recall that the information gain for a candidate angle θ is defined as the difference between the diameter of the current set of binary solutions, and the diameter of the set of binary images that have the same projections as **v** for all angles in $\Theta \cup \{\theta\}$. In this section, we derive an upper bound on the diameter of $\overline{S}_{\mathbf{W}^{(\Theta)}}(\mathbf{p}^{(\Theta)})$, which can be used effectively as a substitute for the true diameter.

An important concept in the derivation of this upper bound is the *central reconstruction*: the shortest real-valued solution in $S_{W^{(\Theta)}}(p^{(\Theta)})$ with respect to the ℓ_2 norm. Throughout this section, we denote this solution by x^* . The central reconstruction can be computed by standard linear algebra methods. In our implementation, we use an iterative Krylov subspace method for solving the system $W^{(\Theta)}x = p^{(\Theta)}$, called *CGLS* (Conjugate Gradient Least Squares, [19]). The CGLS algorithm can effectively exploit the sparse structure of the projection matrix to reduce the required computation time, and does not require storage of large, dense matrices. Apart from numerical errors, applying CGLS to the system $W^{(\Theta)}x = p^{(\Theta)}$ results, after convergence, in the computation of x^* .

Although the derivation of the upper bound on the diameter is already included in [17], we include a slightly modified version here, for easy reference and clarity. We start by showing that the Euclidean norm of any binary solution of $W^{(\Theta)}x = p^{(\Theta)}$ is completely determined by $p^{(\Theta)}$. This property follows from two basic observations: (1) summing the measured values over all detector elements for a projection angle yields the sum of the pixel values in the image, and (2) the squared Euclidean norm of a binary image is equal to the sum of its pixel values.

Lemma 2. Let $\bar{\boldsymbol{x}} \in \overline{S}_{\boldsymbol{W}^{(\Theta)}}(\boldsymbol{p}^{(\Theta)})$. Then, $\|\bar{\boldsymbol{x}}\|_2^2 = \frac{\|\boldsymbol{p}^{(\Theta)}\|_1}{d}$.

Proof. By the definition of the
$$\ell_1$$
-norm, $\|\boldsymbol{p}^{(\Theta)}\|_1 = \sum_{i=1}^m |p_i^{(\Theta)}| = \sum_{i=1}^m p_i^{(\Theta)}$, since $p_i^{(\Theta)} \ge 0$ $(i = 1, ..., m)$. Also,

$$\sum_{i=1}^{m} p_i^{(\Theta)} = \sum_{i=1}^{m} \left(\sum_{j=1}^{n} w_{ij}^{(\Theta)} \bar{x}_j \right) = \sum_{j=1}^{n} \left(\sum_{i=1}^{m} w_{ij}^{(\Theta)} \right) \bar{x}_j = \sum_{j=1}^{n} d\bar{x}_j, \tag{7}$$

and therefore $\|\boldsymbol{p}^{(\Theta)}\|_1 = d\sum_{j=1}^n \bar{x}_j$. As $\bar{\boldsymbol{x}} \in \{0,1\}^n$, we have $\|\bar{\boldsymbol{x}}\|_2^2 = \|\bar{\boldsymbol{x}}\|_1 = \sum_{j=1}^n \bar{x}_j = \frac{\|\boldsymbol{p}^{(\Theta)}\|_1}{d}$. \Box

Define the *central radius* by $R = \sqrt{\left(\frac{\|\boldsymbol{p}^{(\theta)}\|_1}{d} - \|\boldsymbol{x}^*\|_2^2\right)}$. The following Lemma states that all binary solutions of the reconstruction problem lie on the hypersphere centered in \boldsymbol{x}^* with radius R:

Lemma 3. Let
$$\bar{\mathbf{x}} \in \overline{S}_{\mathbf{W}}(\mathbf{p}^{(\Theta)})$$
. Then $\|\bar{\mathbf{x}} - \mathbf{x}^*\|_2 = \sqrt{\frac{\|\mathbf{p}^{(\Theta)}\|_1}{d} - \|\mathbf{x}^*\|_2^2}$

Proof. Let $\mathcal{N}(W)$ denote the nullspace of W. Then $(\bar{x} - x^*) \in \mathcal{N}(W)$. As the shortest solution of a linear system is orthogonal to the nullspace of that system, we have $x^* \perp (\bar{x} - x^*)$. Applying Pythagoras' Theorem and Lemma 2 yields

$$\|\bar{\boldsymbol{x}} - \boldsymbol{x}^*\|_2^2 = \frac{\|\boldsymbol{p}^{(\Theta)}\|_1}{d} - \|\boldsymbol{x}^*\|_2^2. \qquad \Box$$
(8)

Supposing the existence of at least two different binary solutions, Lemma 3 allows us to derive an upper bound on the number of pixel differences between those solutions.

Theorem 4. Let $\bar{\boldsymbol{x}}, \bar{\boldsymbol{y}} \in \overline{S}_{\boldsymbol{W}^{(\Theta)}}(\boldsymbol{p}^{(\Theta)})$. Then $\|\bar{\boldsymbol{x}} - \bar{\boldsymbol{y}}\|_2 \leq 2R$.

Proof. According to Lemma 3, we have $\|\bar{\boldsymbol{x}} - \boldsymbol{x}^*\|_2 = \|\bar{\boldsymbol{y}} - \boldsymbol{x}^*\|_2 = R$. Therefore,

 $\|\bar{\boldsymbol{x}}-\bar{\boldsymbol{y}}\|_2 \leqslant \|\bar{\boldsymbol{x}}-\boldsymbol{x}^*\|_2 + \|\bar{\boldsymbol{y}}-\boldsymbol{x}^*\|_2 = 2R.$

Fig. 2. Algorithmic steps for computing the mean information gain for a candidate angle θ , based on K surrogate solutions.







Fig. 3. Illustration of various steps in the angle selection procedure, where two projections (horizontal and vertical) are already available, and the third projection angle must be selected.

The upper bound from Theorem 4 can be computed simply by evaluating the radius R of the sphere centered in x^* that contains the binary solutions. We remark that there is no guarantee that different binary solutions are indeed so far apart, or even that one or more binary solutions exist. Still, the experimental results in Section 7 demonstrate that the radius R correlates strongly with the reconstruction error that is made when reconstructing a binary image from a small number of projections. It can therefore serve as a substitute of the true diameter, to calculate the amount of "information" present in the projection data.

5. Angle selection algorithm

Combining the ingredients from Sections 3 and 4, we can now define an algorithm for approximating the mean information gain for a candidate angle θ , over the current set of binary solutions. Fig. 2 shows each of the algorithmic steps. Note that this description is formulated for maximum clarity and includes unnecessary recomputation steps, which can be optimized in the actual implementation.

Based on this algorithm for computing the mean information gain, the angle selection algorithm is formed by iterating over all



Fig. 4. Original phantom images used for the experiments.

possible candidate angles, and selecting the angle that yields the highest mean information gain.

6. Example of an angle selection problem

To illustrate the concepts of surrogate solutions and information gain, we now consider an example. Fig. 3(a) shows a binary phantom image \bar{x} that has two principal orientations. Let us suppose that the projection of this image has already been measured for the horizontal and vertical directions. The central reconstruction x^* corresponding to these two projections is shown in Fig. 3b. The difference image $(\bar{x} - x^*)$ is shown in Fig. 3c. The Euclidean norm of this image corresponds with the central radius, and therefore provides, after multiplication by two, an upper bound on the diameter of the set of binary solutions that adhere to the two given projections.

Fig. 3d–f shows three template images that were formed as a random superposition of Gaussian elliptic blobs. The template images are set to 0 outside the circular reconstruction region. The corresponding surrogate solutions are shown in Fig. 3g–i, respectively. The surrogate solutions adhere to the two available projections and contain only pixel values in the interval [0, 1]. Although the surrogate solutions share the same horizontal and vertical projections, there are substantial differences between the three images. In particular, the white arrows in Fig. 3g and i indicate our visual impression of where the major axis of the large ellipse should be. Both surrogate solutions suggest an opposite



Fig. 5. Information gain and rNMP as a function of the projection angle for Phantom 1.



Fig. 6. Information gain and rNMP as a function of the projection angle for Phantom 2.

orientation, which corresponds with the fact that the phantom contains a large binary switching component: two identical blocks of white pixels can be selected in the bottom left and top right part of the large ellipse. When these blocks are flipped vertically (moving them to the top left and bottom right, respectively, and replacing them by background pixels), a new image is formed that has identical projections to the phantom, yet has a general orientation that is rotated 90° compared to the phantom. The phantom image and this second image will likely have quite different information gain characteristics, yet the angle selection algorithm has no way to determine which one is the actual unknown original image. In Fig. 3j, we see that the information gain computed based on the different surrogate solutions has an additional peak around 55°, reflecting the uncertainty about the principal feature direction of the original object. For high quality angle selection, it is important that the different possibilities for binary solutions are properly represented within the set of surrogate solutions, thereby averaging among the information gains for different solutions.

Although it may be possible to gain information about favorable directions using the central reconstruction, it cannot be used as a surrogate solution in our approach. The central reconstruction is the shortest solution that adheres to the available projections, using the Euclidean metric. If one would compute the projections of the central reconstruction for all candidate angles and determine the information gain for each of these angles, this gain will also be 0, as the central reconstruction does not change: the current central reconstruction will also be the shortest solution of the new equation system that is formed by adding the extra projection.

After computing the surrogate solutions, the information gain can now be computed for each angle, and for each surrogate solution, by first computing the projection of the surrogate solution for the new angle, and then computing an upper bound on the diameter of the set of binary solutions for the reconstruction problem that is formed by adding this projection to the current measured data. Fig. 3(j) shows plots of the information gain for each candidate angle, based on knowledge of the phantom image (in red²), and the mean information gain based on the three surrogate solutions (in blue). We observe that, even though the phantom image is hardly recognizable in the three surrogate solutions, the peak of the information gain for the phantom can also be seen in the information gain for the set of surrogate images. A secondary peak can be observed in the plot based on the surrogate solutions, caused by the uncertainty about the original binary object, as the original image contains a large switching component.

7. Experiments

Simulation experiments have been performed to assess the ability of the proposed algorithm to select favorable projection angles. Starting with a general description of the methodology of the experiments in Section 7.1, two sets of experiments are described in Sections 7.2 and 7.3. The results are presented in Section 8.

 $^{^{2}\,}$ For interpretation of color in Fig. 3, the reader is referred to the web version of this article.



Fig. 7. Information gain and rNMP as a function of the projection angle for Phantom 3.

7.1. General methodology

The experiments are all based on simulated projection data obtained by computing the projections of the test images (so-called *phantoms*) in Fig. 4: (a) a single ellipse; (b) a double ellipse; (c) a hand phantom; (d) a foam phantom. All phantoms have a size of 128×128 pixels. The reconstruction area, however, was confined to the pixels within a disk of radius 64.

Ideally, a quantitative evaluation of the proposed angle selection algorithm should include an experiment that shows how well the approximate information gain, computed based on a surrogate solutions and an upper bound on the diameter of the current binary solution set, can approximate the actual information gain as defined by Eq. (5). However, we resorted to this approximation in the first place, *because* of the complexity of this evaluation.

As an alternative, we evaluate the quality of the selected angles with respect to the actual unknown object, based on the assumption that a good angle selection scheme should lead to an accurate reconstruction of the object from fewer angles than the number of angles that would be needed for the standard equi-angular scheme. We remark that this problem should be approached from a statistical point-of-view: in particular cases, an angle for which limited information is gained with respect to the actual unknown object, might yield significantly more information for other candidate solutions that adhere to the currently known projections.

As discussed in Section 3, the class of images from which the template images are sampled can have a substantial impact on the sampling of surrogate solutions. We believe that some form of prior knowledge about the object, represented in the template images, might help to achieve better angle choices. We consider a full exploration of the various possibilities for generating template images to be outside the scope of the present paper, yet recommend such a study for future research. In preliminary experiments, we found that for the type of phantom images we consider here, each template image should contain a substantial number of randomly generated "features", representing a variety of orientations. In particular, for all experimental results presented here, the template images are generated as a superposition of 2D gaussian blobs. Each template image is formed by adding the intensity for 50 such blobs, where the orientation of the blob is chosen randomly and the standard deviation along both the major and minor axis are chosen randomly between 3 and 10 pixels. Examples of such template images are shown in Fig. 3d-f. For all experiments, K = 10 surrogate solutions where generated for the selection of each angle. In preliminary experiments we observed that for the considered phantoms, increasing the number of surrogate images beyond this number did not result in a clear improvement of the angle selection quality. As the computation time increases linearly with the number of surrogate solutions, we decided to use this number. We remark that the information from the available projection data is incorporated in the surrogate solutions that are derived from the template images. So, starting with a template image (that does not contain information related to the current projections), a solution of the current tomography problem is computed that is "close to" the template image. In that way, the sampling diversity of the template images is retained within the set of surrogate solutions, while the available projection data are incorporated as well.

7.2. Experiments I: selecting one new angle

In the first set of experiments, we follow the proposed angle selection algorithm as it chooses consecutive angles for a single phantom. Each time a new angle has to be selected, the following error measures are evaluated for each candidate angle θ :

Phantom gain. The approximate information gain $G(\bar{\mathbf{x}}, \theta)$ for the actual phantom image $\bar{\mathbf{x}}$, computed using the central radius.

Sample gain. The approximate mean information gain $G(V, \theta)$ for a set of surrogate solutions V, computed using the central radius.

Relative number of misclassified pixels (rNMP). A thresholded CGLS (TCGLS) reconstruction is first computed, with threshold 0.5, obtained after adding the projection for angle θ . This value is computed by determining the number of pixel differences between the TCGLS reconstruction and the phantom image, and dividing this by the total number of pixels in the circular reconstruction region.

Although the TCLGS reconstruction is typically inferior to reconstructions computed by more specialized Discrete Tomography algorithms [11,12], its computation is straightforward as a byproduct of the proposed angle selection method, and the general shape of the rNMP curve seems to be similar compared to more advanced methods.

When interpreting the results, we focus on two comparisons. Firstly, if the (local) maxima of the information gain for the surrogate solutions coincide with the maxima for the information gain based on the phantom, the angle selection algorithm will choose an angle that yields substantial information about the phantom, as it results in a relatively strong reduction of the central radius. Secondly, the positions of the extremal points of the information gain curve for the phantom should correspond with the extremal points of the rNMP curve, such that local maxima of the gain curve correspond with local minima of the rNMP curve, and vice versa. If this property is approximately satisfied, a large information gain as defined by Eq. (4) will correspond with a relatively large reduction of the reconstruction error for the TCGLS algorithm.

7.3. Experiments II: selecting a series of angles

In the second set of experiments, we consider the selection of a complete sequence of angles, starting with a single angle and expanding the sequence up to 20 angles. The experiments have been performed using three different angle selection methods:

Standard. Angles are selected between 0° and 180°, with equiangular spacing. Note that for this scheme, changing the *number* of angles actually changes the entire set of selected angles. *Gap-angle*. Based on the current set of angles, a new angle is selected as the mid-point between the two consecutive angles that have the largest angular gap between them. If several pairs of angles have equal gaps, a uniformly random choice is made. *Dynamic*. Angles are selected using the dynamic angle selection algorithm of Section 5. To select an angle, K = 10 surrogate



Fig. 8. Information gain and rNMP as a function of the projection angle for Phantom 4.

solutions are generated, and their mean information gain is evaluated for all candidate angles, with steps of 1°.

To reduce the dependency of the results on the particular starting angle, as well as the dependency on the random seed (for the Dynamic and Gap-angle strategies), a series of runs has been performed for each angle selection method, using nine randomly selected starting angles, and five different random seeds for each starting angle.

8. Results

In this section, we report on the results of the two sets of experiments, described in Sections 7.2 and 7.3.

8.1. Results I: selecting one new angle

Figs. 5–8 show the Phantom gain and the Sample gain, for each candidate angle between 0° and 180° (in steps of 1°), and for each phantom image, respectively. The Phantom gain is shown along with its variation (shaded area), indicating the range from the minimum to the maximum value observed for a surrogate solution. As only the positions of the local optima of the gain curves are important for angle selection, and not the absolute values of the approximate information gain, the curve for the Sample gain is scaled such that the total area under the graph is the same as for the Phantom gain. In addition to the gain curves, the rNMP of the

TCGLS reconstruction for each candidate angle, obtained by reconstructing the image after adding the projection for the candidate angle, is shown in the same plot.

In Figs. 5a and 6a, it can be clearly observed that if the number of projections is extremely small (two, in this case), the dynamic angle selection algorithm does not select the angle that corresponds to the maximum Phantom gain. This can be explained by the fact that for d = 2, the reconstruction problem is highly underdetermined, and there are large differences between the surrogate solutions. An angle choice that is optimal with respect to the phantom does not appear to be optimal with respect to the combined set of surrogate solutions.

However, for more than two projections, the correlation between the Sample gain curve and the Phantom gain curve is much higher, for both Phantoms 1 and 2. Indeed, for three projections, the dynamic angle selection selects an angle that closely matches that of the maximum Phantom gain (see Figs. 5b and 6b). Similar observations can be made for four and five projections. For the more complex Phantoms 3 and 4, the number of projections that is needed to reach a reasonable agreement between the Phantom gain and the Sample gain is substantially higher. For Phantom 3, there is a high correlation between both curves for $d \ge 4$, whereas 14 projections are required to properly localize the "best" projection angle for Phantom 4.

Another observation that can be made, is that for all four phantoms, there is generally a good agreement between the rNMP that is computed with respect to the phantom, and the Phantom gain. Local minima of the rNMP curve tend to correspond with local



Fig. 9. rNMP as a function of the number of projections.

maxima of the Phantom gain, and vice versa. This correspondence is crucial for the ability of the dynamic angle selection algorithm to rapidly achieve a high reconstruction quality with respect to the phantom object, even though the phantom is not known to the algorithm.

8.2. Results II: selecting a series of angles

To compare the results for the three angle selection methods, two plots have been generated for each phantom. The first series of plots, presented in Fig. 9, shows the rNMP of TCGLS as a function of the number of projection angles. For each number d of angles, the TCGLS reconstruction was computed based on the projections for the first d angles in the sequence. The second series of plots, presented in Fig. 10 shows the *central radius* as a function of the number of projection angles. For each number d of projection angles, the central radius, which scales proportionally to the upper bound on the diameter of the set of binary solutions, is determined based on the first d angles in the sequence. The shaded regions denote the standard deviation originating from repeating the experiments with varying starting angles and random seeds.

The results show that for Phantoms 1, 2, and 4, the rNMP the Dynamic strategy is never clearly outperformed by the Standard and Gap-angle methods, while at substantial intervals for the number of projection angles, the rNMP are substantially better for the Dynamic strategy. For Phantom 1, the proposed angle selection method works particularly well if the number of angles is very small (between 3 and 10), whereas for the more complex Phantom 4, the key advantage of the Dynamic strategy shows up when using a larger number of projections (10 or more). The advantage of dynamic angle selection is smallest for Phantom 3 (the *hand* phantom), which has a relatively large number of features, each having a different orientation.

What is also clear from the plots in Figs. 9 and 10, is that the curves for the rNMP with respect to the unknown phantom, and for the central radius, have similar shapes and roughly show the same ordering between the performance of the different methods. Just as for the results in the first series of experiments, confirmation of this correspondence is important for establishing the ability of the dynamic angle selection algorithm to rapidly achieve a high reconstruction quality with respect to the phantom object, even though the phantom is not known to the algorithm.

In Fig. 11, a selection of TCGLS reconstructions is shown that was obtained after a run of the three angle selection strategies. For each phantom, the number of angles was selected such that a reasonable reconstruction could already be obtained, quite similar to the phantom, yet the number of projections was still too small to obtain an almost perfect reconstruction.



Fig. 10. Central radius as a function of the number of projections.



Fig. 11. Illustration of TCGLS reconstructions obtained after using the different angle selection algorithms, for the four phantoms. Next to each binary reconstruction, an image is shown that represents the difference between the reconstruction and the phantom. From left to right: standard, gap angle, dynamic; from top to bottom: Phantom 1 (6 angles), Phantom 2 (8 angles), Phantom 3 (12 angles), Phantom 4 (16 angles).

9. Discussion

The experimental results for the four phantom images suggest that the proposed dynamic angle selection strategy is indeed capable of finding angles for which the reconstructed image is substantially more accurate than for the standard angle selection scheme, if the number of angles is not too small. It seems that although the central radius provides an upper bound on the *diameter* of the set of binary solutions, a strategy that aims to reduce this upper bound also reduces the actual reconstruction error with respect to the unknown phantom.

It should also be noted that the proposed angle selection algorithm lacks computational efficiency. Within the context of this article, we focused on the capabilities of the algorithm with respect to reconstruction quality, and on avoiding the exponential time complexity of enumerating the set of binary solutions. However, in our implementation which utilizes a modern GPU for performing the SIRT and CGLS subroutines, selecting a single angle for an image of size 128×128 takes around 30 s if the current number of angles is 4.

By far the most computation time is spent in the CGLS subroutine, which computes the central reconstruction within the evaluation of the information gain for a surrogate solution. As this computation is performed for each surrogate solution, and for each angle, it is called around 180 K times just to select a single angle. We expect that the algorithm can be made far more efficient if the complete recomputation of the central radius can be avoided by using the results from earlier computations.

Despite these computational obstacles, we feel that the type of algorithm proposed here should be seen as a proof-of-concept for more efficient, or more advanced algorithms that utilize the concepts of information gain. Formal modeling of the amount of information that is present in a set of projection data is not straightforward, in particular for underdetermined reconstruction problems, where many solutions may exist. The results also demonstrate that dynamic angle selection is not only useful in the domain of an extremely small number of projection angles (i.e. less than 10). In fact, for the more complex Phantom 4, which clearly has a certain "preferential orientation", the dynamic angle selection scheme mainly outperforms the standard method if there are *more than* 10 projections. This means that the approach can potentially be useful for angle selection in the reconstruction of quite complex objects, such as trabecular bones, where it is crucial to use as few projection images as possible, due to radiation damage.

For the derivation of our angle selection algorithm, we made the idealized assumption that the projection data contains no noise. The actual algorithm can also be used when there is noise, as both SIRT (used for generating the surrogate solutions) and CGLS (used for computing the information gain) are convergent even if the righthand side of the linear system is perturbed. We consider a complete exploration of the capabilities and limitations to handle noisy data to be outside of the scope of the current article.

10. Conclusions

In this article, we proposed a formal model for the angle selection problem in binary tomography, and an actual *algorithm* that was obtained by introducing several approximation steps with respect to the idealized model.

It was shown that for a set of phantom images, the proposed dynamic angle selection strategy compares favorably to other strategies that do not take specific properties of the binary reconstruction problem into account.

Although the computational requirements for the dynamic angle selection method are currently restricting its use for large experimental dataset, we consider this work as a starting point for developing more efficient algorithm variants, based on the same underlying ideas. An important question that was not addressed in this work, is how the class of template images influences the results. Ideally, these images should represent properties that can occur in the true scanned objects, while also having sufficient variations to provide a trustful sampling of the current set of binary solutions. In future work, we intend to explore this relation further.

Acknowledgments

This research was partially supported by the TÁMOP-4.2.2/08/ 1/2008-0008 Program of the Hungarian National Development Agency, the European Union and the European Regional Development Fund under the Grant Agreement TÁMOP-4.2.1/B-09/1/ KONV-2010-0005. Part of the research was performed when PB was visiting the ASTRA Group at Vision Lab, University of Antwerp in Belgium. The research was also supported by the János Bolyai Research Scholarship of the Hungarian Academy of Sciences and the Grant PD100950 of the Hungarian Scientific Research Fund (OTKA). K.J.B. acknowledges support from the Netherlands Fund for Scientific Research (NWO), Project Number 639.072.005.

References

- [1] G.T. Herman, Fundamentals of Computerized Tomography: Image reconstruction from projections, Springer, 2009.
- [2] A.C. Kak, M. Slaney, Principles of Computerized Tomographic Imaging, SIAM, 2001.
- [3] P.A. Midgley, R.E. Dunin-Borkowski, Electron tomography and holography in materials science, Nat. Mater. 8 (4) (2009) 271–280.
- [4] E.J. Candès, J. Romberg, T. Tao, Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information, IEEE Trans. Inform. Theory 52 (2) (2006) 489–509.

- [5] G.-H. Chen, J. Tang, S. Leng, Prior image constrained compressed sensing (PICCS): a method to accurately reconstruct dynamic CT images from highly undersampled projection data sets, Med. Phys. 35 (2) (2008) 660–663.
- [6] D. Donoho, Compressed sensing, IEEE Trans. Inform. Theory 52 (4) (2006) 1289–1306.
- [7] E.Y. Sidky, M.A. Anastasio, X. Pan, Image reconstruction exploiting object sparsity in boundary-enhanced X-ray phase-contrast tomography, Opt. Express 18 (10) (2010) 10404–10422.
- [8] E.Y. Sidky, C.M. Kao, X. Pan, Accurate image reconstruction from few-views and limited-angle data in divergent-beam CT, J. X-ray Sci. Technol. 14 (2006) 119– 139.
- [9] G.T. Herman, A. Kuba (Eds.), Discrete Tomography: Foundations, Algorithms and Applications, Birkhäuser, Boston, 1999.
- [10] G.T. Herman, A. Kuba (Eds.), Advances in Discrete Tomography and its Applications, Birkhäuser, Boston, 2007.
- [11] K.J. Batenburg, J. Sijbers, DART: a practical reconstruction algorithm for discrete tomography, IEEE Trans. Image Process. 20 (9) (2011) 2542–2553.
- [12] T. Schüle, C. Schnörr, S. Weber, J. Hornegger, Discrete tomography by convexconcave regularization and D.C. programming, Discrete Appl. Math. 151 (2005) 229–243.
- [13] L. Varga, P. Balázs, A. Nagy, Direction-dependency of binary tomographic reconstruction algorithms, Graph. Models (73) (2011) 365–375.
- [14] L. Varga, P. Balázs, A. Nagy, Projection selection dependency in binary tomography, Acta Cybernet. 20 (1) (2011) 167–187.
- [15] Z. Zheng, K. Mueller, Identifying sets of favorable projections for few-view low-dose cone-beam CT scanning, in: 11th International Meeting on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine, 2011, pp. 314–317.
- [16] R.J. Gardner, P. Gritzmann, D. Prangenberg, On the computational complexity of reconstructing lattice sets from their X-rays, Discrete Math. 202 (1999) 45– 71.
- [17] K.J. Batenburg, W. Fortes, L. Hajdu, R. Tijdeman, Bounds on the difference between reconstructions in binary tomography, in: Proceedings of the 16th IAPR International Conference on Discrete Geometry for Computer Imagery, DGCl'11, Springer-Verlag, Berlin, Heidelberg, 2011, pp. 369–380. http:// dl.acm.org/citation.cfm?id=1987119.1987156.
- [18] J. Gregor, T. Benson, Computational analysis and improvement of SIRT, IEEE Trans. Med. Imag. 27 (7) (2008) 918–924.
- [19] Y. Saad, Iterative Methods for Sparse Linear Systems, SIAM, Philadelphia, PA, USA, 2003.