A Linear Programming Approach to Limited Angle 3D Reconstruction from DSA Projections

S. Weber[1][a], T. Schüle[a,c], C. Schnörr[a], J. Hornegger[b]

[a]University of Mannheim, Dept. M&CS, CVGPR-Group
D-68131 Mannheim, Germany
www.cvgpr.uni-mannheim.de
{wstefan,schuele,schnoerr}@uni-mannheim.de

[b]Friedrich-Alexander University Erlangen-Nürnberg
Dept. CS, Chair of Pattern Recognition
www5.informatik.uni-erlangen.de
Joachim.Hornegger@informatik.uni-erlangen.de

[c] Siemens Medical Solutions, Inc.
Forchheim, Germany
www.medical.siemens.com

[1]Correspondence: Phone +49 621 181-3495, Fax -2744

Abstract

Objectives: We investigate the feasibility of binary-valued 3D tomographic reconstruction using only a small number of projections acquired over a limited range of angles.

Methods: Regularization of this strongly ill-posed problem is achieved by (i) confining the reconstruction to binary vessel / non-vessel decisions, and (ii) by minimizing a global functional involving a smoothness prior.

Results: Our approach successfully reconstructs volumetric vessel structures from 3 projections taken within $90°$. The percentage of reconstructed voxels differing from ground truth is below 1%.

Conclusion: We demonstrate that for particular applications – like Digital Substraction Angiography – 3D reconstructions are possible where conventional methods must fail, due to a severly limited imaging geometry. This could play an important role for dose reduction and 3D reconstruction using non-conventional technical setups.

MeSH–Keywords: X-Ray Tomography, Computer-Assisted Image Analysis, Image Reconstruction, Computer-Generated 3D Imaging, Digital Subtraction Angiography

1 Introduction

1.1 Motivation

The process of reconstructing the 3D density distribution within the human body from multiple X-ray projections is well understood [1]. Today, filtered back-projection is the fundamental algorithm for Computerized Tomography. This algorithm, however, has its limitations in that a large amount of input data is required for the reconstruction to become feasible. More specifically, a necessary condition for its success is the rotation of the X-ray tube of at least 180 degrees plus fan angle and the acquisition of a large number of projections [1].

There are prospective applications of 3D imaging where the technical setup does not allow for 180 degree rotations and, therefore, filtered back-projection cannot be applied. For instance, the reconstruction of the coronary vessels of the moving heart using the Feldkamp algorithm requires much more data than can be captured by C-arm systems during interventions. More generally, it is reasonable to investigate situations where the dose rate to which patients are exposed can be considerably reduced. Unfortunatly, violating the constraints for data acquisition mentioned above renders 3D-reconstruction almost impossible [1].

A particular situation concerns *Digital Subtraction Angiography (DSA)* where vessels are filled with a constrast agent and the background is almost homogeneous. As a consequence, we shift our focus from 3D-reconstruction of a *continuous* density functions to the 3D-reconstruction of a *binary-valued* functions where each voxel indicates the presence or absence of a vessel. The natural question then is whether the considerably reduced degrees of freedom of the function to be reconstructed can compensate for the lack of input data when recording *few projection data only over a limited range of angles*. The present paper addresses this question.

1.2 Related Work

Our approach is based on recent research in the field of *Discrete Tomography* [2] which, historically, originated from several branches of mathematics like, for example, the combinatorial problem to determine binary matrices from its row and column sums (see the survey [3]). Due to the fact that reconstructions are required to be *discrete-valued*, the reconstruction task inevitably leads to *combinatorial optimization problems* which, accordingly, require a quite different treatment than conventional 3D tomography. Meanwhile, however, progress is not only driven by

challenging theoretical problems [4, 5] but also by real-world applications where discrete tomography might play an essential role (cf. [2, chapters 15–21]).

Despite the fact that 3D binary reconstructions from few projections may be arbitrarily worse *in theory* [5], numerical experiments [6] based on determining reconstructions being consistent with given projection data by *linear programming (LP)* yielded encouraging results. Furthermore, various objective functions were investigated in [7] in order to obtain "maximal" consistent reconstructions[2], among other things, and suboptimal solutions were computed within a greedy framework. Our work supplements the LP approach for computing feasible reconstructions suggested in [6] with a particular objective function proposed in [7] and, additionally, includes a smoothness prior enforcing spatially homogeneous reconstructions. The latter is a natural but non-trivial extension (in the present context) and under the notion *regularization* well-known from many variational approaches in image processing and computational vision (cf. , e.g., [9, 10]). In connection with discrete-valued tomographic reconstruction, our LP-based regularization approach is novel and further motivated from the viewpoint of combinatorial optimization by favourable approximation properties proved in [8].

Other related work in the field of Discrete Tomography include the MRF-based approaches by [11, 12, 13] and the binary steering technique by [14] (see also [2]). Both Matej et al. [11] and Chan et al. [12] use regularization priors as well. On the other hand, stochastic sampling (Metropolis and Gibbs sampling, respectively) is used in [11, 12] for the purpose of optimization which, when properly applied, is notoriously slow, whereas a multiscale implementation of a coordinate-wise sequential update technique (a special version of the ICM-technique) is employed in [13]. In this connection, we focus on a LP framework for which mathematically sound concepts of *convex optimization* are known [15] providing the basis for future parallel implementations.

We also stress that discrete tomography should not be confused with other reconstruction approaches like tomosynthesis, for instance. In *tomosynthesis* [16], single slices can be reconstructed from a single, limited angle movement of the X-ray source, where patient and digital detector are kept stationary. By adding and shifting the acquired images it is possible to bring a certain slice of the 3D structure in focus. The anatomy below and above the current slice is blurred and, therefore, a complete 3D volume cannot be reconstructed using this method.

---

[2]See section 2.2.1 for a precise definition.

4

1.3 Objectives

The objective of this paper is a feasibility study of 3D tomographic reconstruction using a *small number of projections* acquired over a *limited range of angles*. To this end, we combine a suitable objective function with a regularizing term within a LP approach. The study is mainly performed by comparing numerous reconstructions obtained by systematically varying the image acquisition conditions with ground-thruth. Ground-truth is defined to be the full 3D-reconstruction of real vessel structures obtained by common techniques and *many* projections over a *wide* range of angles.

We stress that our approach is applicable only in situations where the assumption of *binary-valued* volumes is justified. The corresponding application we have in view is 3D reconstruction from DSA projections. While the range of applications of our approach is clearly limited, the prospect of both exposing patients to lower dose rates, and obtaining 3D-information under conditions where conventional methods fail, motivates and justifies our work.

2 Methods

In this section, we describe the mathematical and computational methods underlying our approach. Subsection 2.1 concerns pre-processing: DSA data acquisition, discrete representation used for volumetric structures, and determination of locations which, based on given projection data, do surely not belong to any volumetric data to be reconstructed in the subsequent reconstruction step.

The reconstruction of volumetric structures amounts to solve a large-scale LP whose specific form depends on the projection data, the objective function for reconstruction, a regularization term and the relaxation of the combinatorial integer constraint. These points are described in subsection 2.2.

2.1 Preprocessing

2.1.1 DSA - Digital Subtraction Angiography

DSA is a common technique in medical diagnostics. Two images of a vascular system are taken from each direction, one with contrast agent and another one without. Afterwards, both images are subtracted from each other for each image pair so that only the distribution of the contrast agent remains. The logarithm of

each image intensity is taken before subtracting due to the exponential absorption function. Figure 1 illustrates this process.
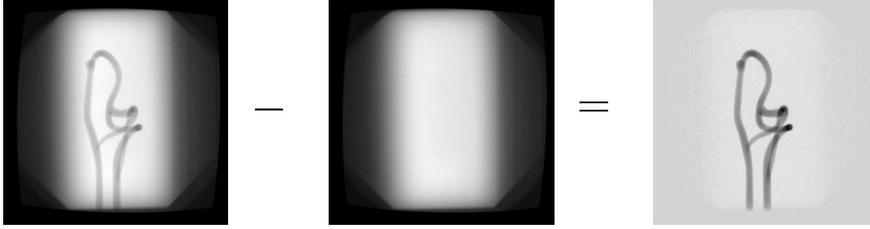
Figure 1: Process of DSA imaging. In principle, two images of a vascular system are taken with (left) and without (center) contrast agent, respectively. This is illustrated here by means of a phantom (made of knead) scanned with a C-arm tomograph (in $1°$-steps over a range of $180°$). Subtraction of both images leads to an image (right) that contains only information about the distribution of the contrast agent.

2.1.2 Discretization

Consider a function $f : \Omega \subset \mathbb{R}^N \to \mathbb{R}$, defined by $\mathbf{x} \mapsto f(\mathbf{x})$, and a linear ray $\tau \mapsto \mathbf{x}(\tau) = \tau \mathbf{m} + \mathbf{c}$ parametrized by $\tau \in \mathbb{R}$, $\mathbf{m}, \mathbf{c} \in \mathbb{R}^N$. In the continuous space the projection value b measured along this ray is given by equation (1).

$$\int_R f(\tau \mathbf{m} + \mathbf{c}) \ d\tau = b \tag{1}$$

Splitting up the integral for each pixel (voxel) traversed by the ray yields a discretization of equation (1). This is shown in Figure (2) for the 2-dimensional case where some ray (indicated by an arrow) hits the pixels assigned to the variables x_3, x_4, x_5, x_6, x_7. For all these pixels x_i, dots depict entry and leaving point of the ray, and a_i denotes the length of the corresponding intersection. In the 3-dimensional case one has to compute the entry and leaving points for "boxes" representing voxels instead of "squares" representing pixels.

Each pixel variable takes ranges over two values only: $x_i \in \{0, 1\}$. Its values depend on whether a pixel contributes to the measured projection b or not. The integral value, that is the contribution of the i-th pixel to the projection value b, accordingly is $a_i x_i$. For pixels x_i not hit by the ray, we set $a_i = 0$. Defining
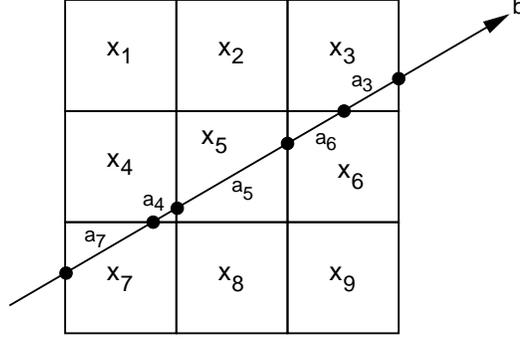
Figure 2: Illustration of the discretization of the projection (line integral) along a ray (indicated by the arrow) for a 3×3 image. Dots indicate the intersection of the ray with image pixels x_i. The lengths of these intersections are assigned to corresponding variables a_i. Pixel variables may only take two values $x_i \in \{0, 1\}$ depending on whether a pixel conributes to the measured projection value b or not. Accordingly, the equation $\sum_i a_i x_i = b$ represents all possible pixel values consistent with the observed value b. The procedure for 3D projections is exactly the same with "squares" x_i (pixels) replaced by "boxes" x_i (voxels).

$\mathbf{a} := (a_1, \ldots, a_N)^\top$ and $\mathbf{x} := (x_1, \ldots, x_N)^\top$, we can set up an equation for the projection ray,

$$\sum_{i=1}^{N} a_i x_i = \mathbf{a}^\top \mathbf{x} = b , \tag{2}$$

which describes all possible pixel configurations consistent with the observed projection b.

So far, equation (2) represents a discretization for a single projection ray only. In order to obtain a complete description of the reconstruction problem, one has to consider many projection rays from several projection directions. The corresponding measurements are collected into the vector $\mathbf{b} := (b_1, \ldots, b_M)^\top$. Analogous to (2), we obtain equations $a_j^\top x = b_j$, $j = 1, \ldots, M$ for each projection ray j. Assembling a_{ji}, $1 \le i \le N$, $1 \le j \le M$, into a matrix $\mathbf{A}$ leads to the following representation of the image measurements:

$$\mathbf{A}\mathbf{x} = \mathbf{b} \tag{3}$$

As detailed above, vector b contains all observed projection measurements, and matrix A encodes the known geometry of the projections. The *reconstruction problem* is to recover the unknown values $\mathbf{x}$ from equation (3).

2.1.3   Removing Unambiguous Variables

As Figure (2) suggests, the value $b = 0$ is measured for many projections ray which do not hit any absorbing volumetric structure. As a consequence, all variables x_i hit by any of these rays can be removed from the linear system (3) before the reconstruction process starts.

Especially for vascular systems, the reduction of unknowns obtained in this way is significant since the vessels (non-zero voxels) cover only a small fraction of the entire volume.

2.2   Regularization and Binary Reconstruction

This section describes our approach to 3D reconstruction in terms of the unknown variables $\mathbf{x}$ to be recovered from the linear system (3). We discuss the objective function for optimization as well as the relaxation of the integer constraint $x_i \in \{0, 1\}$ first (subsection 2.2.1), followed by extensions in order to include regularizing terms (subsection 2.2.2).

2.2.1   Objective Function and Problem Relaxation

It is known that the computational effort to compute a binary solution $\mathbf{x}$ to the system (3) grows exponentially with the number of unknowns if more than two projection directions are considered [5]. Therefore, the integer constraint $\mathbf{x} \in \{0, 1\}^N$ is relaxed to the *convex* constraints set

$$C := \left\{ \mathbf{x} \in \mathbb{R}^N \mid 0 \le x_i \le 1, \quad i = 1, \dots, N \right\} \tag{4}$$

and the linear system (3) is supplemented by an objective function to form a LP for 3D reconstruction.

In the literature on discrete tomography, two LP approaches are known. The first one (5), suggested by Fishburn, Schwander, Shepp, and Vanderbei [6], optimizes the dummy functional "zero" subject to the linear projection constraints (3). Thus, any interior point method for solving large scale LPs can be used for computing *some feasible point (FP)* satisfying the projection constraints.

$$(FP) \qquad \min_{\mathbf{x} \in C} \mathbf{0}^\top \mathbf{x}, \quad \text{subject to} \quad \mathbf{A}\mathbf{x} = \mathbf{b} \tag{5}$$

The second approach (6), suggested by Gritzmann, de Vries, and Wiegelmann [7], replaces the dummy functional in (5) by the term

$$\mathbf{e}^\top \mathbf{x}, \quad \mathbf{e} := (1, 1, \dots, 1)^\top,$$

which measures the size of the reconstructed volume in terms of voxels i with $x_i = 1$. Furthermore, the linear projection equations are changed to linear inequalities which copes better with slight inconsistencies of (3) due to measurement errors and discretization effects (i.e., (3) may not hold *exactly* for *binary* vectors $\mathbf{x}$). In summary, the *"best inner fit" criterion (BIF)* (6) aims at computing a maximal volume among all solutions not violating the projection constraints.

Note that problem (6) was optimized in [7] over the difficult constraint set $\{0, 1\}^N$ by local optimization within a greedy framework. In contrast, we again relax this problem by optimizing over the larger set C defined in (4) so as to make global optimization (by LP) feasible which takes all constraints simultaneously into account. Furthermore, this provides a basis for regularization (see next section).

$$(BIF) \qquad \max_{\mathbf{x} \in C} \mathbf{e}^\top \mathbf{x} \,, \quad \text{subject to} \quad \mathbf{A}\mathbf{x} \leq \mathbf{b} \tag{6}$$

2.2.2 Regularization

Both approaches (5) and (6) do not exploit spatial coherency which is not plausible in connection with the reconstruction of volumetric structures. As a result, spatially incoherent and thus less plausible solutions may be favored by the optimization process.

A common remedy is to include smoothness priors into the optimization criterion. However, since we deal with integer solutions, this further complicates the combinatorial optimization problem. Furthermore, smoothness priors lead to quadratic functionals which cannot be incorporated by LP relaxations.

Inspired by recent progress on general metric labeling problems [8], we introduce auxiliary variables to represent the absolute deviation of adjacent entities, as an approximation of standard quadratic priors. Consequently, spatial smoothness can be measured by a linear combination of auxiliary variables, leading to an extended LP approach.

Let $\sum_{\langle j,k \rangle}$ denote the sum over all adjacent pixel pairs (4-neighborhood) and voxel pairs (4-neighborhood), respectively. Then (6) is extended to the *regularized best inner fit (R-BIF)* approach:

$$(R - BIF) \qquad \min_{\mathbf{x} \in C} -\mathbf{e}^\top \mathbf{x} + \frac{\alpha}{2} \sum_{\langle j,k \rangle} |x_j - x_k| \,, \quad \mathbf{A}\mathbf{x} \leq \mathbf{b} \tag{7}$$

2.2.3 Implementation Details

In order to arrive at a LP in standard form, we introduce additional slack variables $z_{j,k}$ for each pair of adjacent pixels/voxels $\langle j, k \rangle$. Let the vector $\mathbf{z}$ collect all variables $\{z_{j,k}\}$. Then the vector of unknowns $\mathbf{x}$ in (7) is supplemented by $\mathbf{z}$, and the linear constraints in both approaches are supplemented by two linear inequalities for each variable $z_{j,k}$:

$$(x_j - x_k) - z_{j,k} \leq 0 \,,$$
$$-(x_j - x_k) - z_{j,k} \leq 0$$

By virtue of these additional constraints, we can replace each term $|x_j - x_k|$ in (7) by $z_{j,k}$.

2.2.4 Postprocessing

The LP step results in a solution vector $\mathbf{x} \in C$ with each component $x_i \in [0, 1]$. In order to obtain a binary solution we simply used a threshold t:

$$[x_i]_t := \begin{cases} 0 & \text{if } x_i \leq t \\ 1 & \text{otherwise} \end{cases} \tag{8}$$

3 Evaluation

The first data set consists of real projection data which were taken with a C-arm system from a phantom made of knead (Figure 1). The volume resolution is $200 \times 200 \times 200$ voxel, and 1024×1024 pixel for the projection images. The corresponding $200^3 = 8^6$ unknowns were typically reduced to about $20^3 \ldots 25^3$ variables by the procedure described in Section 2.1.3.

This was our first experiment using three-dimensional data. No 3D ground truth was available. The positive results described in Section 4.1 encouraged us to set up an 3D ground-truth experiment.

In order to precisely evaluate the capability for 3D-reconstruction based on a small number of projections taken over a limited range of projection angles, we reconstructed a real vascular volumetric structure which was scanned with a C-arm system using a conventional method (filtered backprojection). This structure is shown in Figure 3, left, and was considered as 3D ground-truth.

The volume resolution is $256 \times 256 \times 240$ voxel, and 512×512 pixel for each projection. The diameters of the vessel structures range from $2.5 \ldots 8.5$ voxel

(see Fig. 3). The structure shown in Fig. 3 consists of 24466 voxels, whereas the number of unknowns after preprocessing (section 2.1.3) was about 3 times this number.
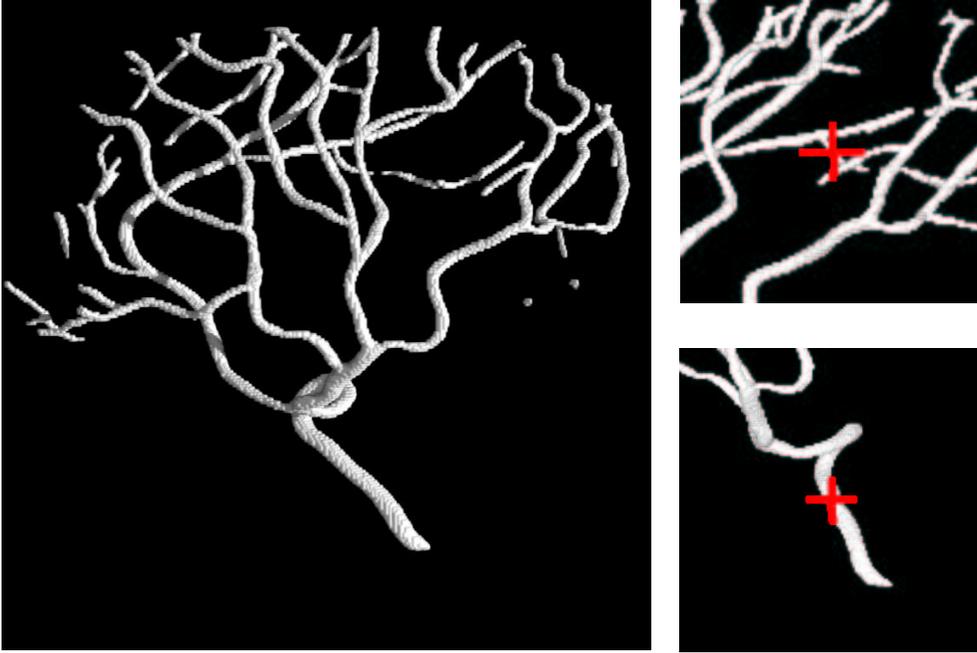
Figure 3: **Left:** A real 3D vessel structure reconstructed by conventional tomography. This data set was taken as ground truth for various reconstructions, each computed from three projection directions only within the limited range of $90°$. **Right:** Close-ups of the structure. The diameters at the locations marked with crosses are 8.5 voxel (bottom) and 2.5 voxel (top), respectively.

Figure 4 shows the imaging geometry used for the evaluation. In case of the knead phantom we used 5 projections over a range of $90°$. The reconstruction results were judged by visual inspection. This experiment confirmed that our regularized LP approach (R-BIF) significantly improved 3D-reconstruction.

Concerning the vascular structure shown in Figure 3, we constructed projection data as follows (see Fig. 4). Over a range of $90°$, only three projections only were computed. In this way, we produced 19 different data sets by varying the offset from $0°$ to $180°$ in $10°$-steps. For each of these 3-projection data sets a separate 3D-reconstruction was computed and compared with the true data set.

We used a threshold $t = 0.5$ for all experiments concerning the vessel structure. For the knead phantom, a value of $t = 0.25$ gave visually slightly more pleasant results. In our future work, we will adopt probabilistic methods to deter-

mine this value automatically.

Parameter $\alpha = 1/6$ $(= 1/$ number of neighbors) was chosen for all experiments.
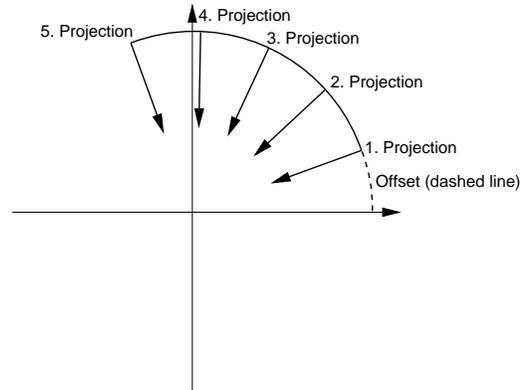
Figure 4: Illustration of the imaging geometry with a small number of projections – five projections for the knead phantom, and three projections for the vessel structure – and a range of projection directions limited to $90°$. Different data sets were produced by varying the offset from $0° \ldots 180°$. While varying the offset, the spacing of directions was kept constant.

4 Numerical Results

4.1 3D Knead Phantom

Figure 5 shows the 3D-reconstructions of the knead phantom computed with the approaches (6) and (7), respectively. The lower panel illustrates the favourable effect of the regularization employed in (7): in contrast to the result computed with the approach (6), isolated and noisy voxel configurations are avoided during the reconstruction based on (7), which results in spatially homogeneous volume structures.

4.2 3D Blood Vessel Structure

Figure 6 shows the original 3D data set (top, left), the undetermined variables after the preprocessing phase described in section 2.1.3 (top, right), and in the lower panel the results computed by (6) and (7), respectively.

Figure 5: **Left column:** 3D reconstruction of the phantom using the approach (BIF), (6). **Right column:** 3D reconstruction of the phantom using the approach $(R - BIF)$, (7).

Note that during preprocessing phase a large number of variables (attached to each voxel) were removed, due to the volumetric sparseness of vessel structures. Relative to that, however, a considerable amount of ambiguous variables remains (Fig. 6, top right). The result shown in Fig. 6, bottom left, illustrates that the computation for solving (6) reconstructs the desired vessel structure immersed into a set of isolated voxels which do not belong to the vessel structure. Additionally invoking spatial regularization by virtue of (7) almost completely suppresses this set during the reconstruction (Fig. 6, bottom right). Figure 7 depicts close-ups of the three data sets for better visual inspection.

We confirmed and evaluated this result quantitatively by systematically repeating the reconstruction while varying the offset of the projection range (see Fig. 4, and the corresponding caption). Figure 8 shows the errors of the 3D reconstructions for all 19 data sets constructed in this way, in terms of the percentage of

voxels differing from ground truth. The evaluation revealed errors of $6\% - 10\%$ for reconstructions based on (6), and errors below 1% for reconstructions based on (7).

5 Discussion

The results depicted in Figs. 5, 6, and 6 show

1. that regularization significantly improves 3D reconstructions in terms of spatial coherency of volumetric structures, and

2. that our approximation (7) of a standard smoothness prior as part of a LP approach to discrete tomography works (see also section 2.2.3).

The 3D reconstruction errors obtained during our evaluation (Fig. 8, lower panel) strongly suggest that a real volumetric vessel structure *can* be reconstructed from *few* projections over a *limited* range of angles only, despite the severe ill-posedness from the mathematical point of view. This ill-posedness is clearly visible in terms of the variation of the reconstruction error as a function of the absolute orientation of the angular projection range (Fig. 8, upper panel). The more volumetric structures occlude each other, the larger the reconstruction error is. While this dependency is visible in the lower panel of Fig. 8 as well, the error has been reduced by a factor of about 10, thus yielding reconstructions close to ground truth (absolute error $< 1\%$) uniformly for all experiments.

Let us add two further comments in order to put this positive empirical results into perspective:

- At present, we have only a computational approach and empirical results, but no proof specifying conditions under which such accurate reconstructions are possible. Our results merely show that there exist practically relevant situations where ill-posedness due to constrained imaging conditions can be compensated by restricting the degree of freedoms of the functions to be reconstructed (in our case: binary valued functions). On the other hand, we are not aware of any other work showing similar empirical results.

- We deliberately excluded all sources of errors due to imaging, preprocessing, discretization etc., by computing (not measuring!) the projections of the vessel structure from 3D ground truth. This allowed us to focus directly on the reconstruction problem and to draw the conclusions stated under the previous item.

14

6 Conclusions and Further Work

We empirically explored the 3D tomographic reconstruction problem under imaging conditions where standard methods must fail. Our work paves the way for exploring such situations under clinical conditions in the near future.

In this context, we will focus next on Digital Substraction Angiography and noise effects due to imaging (e.g., calibration), preprocessing, and discretization. Furthermore, we will enhance our smoothness prior by learning relevant structures from large samples of reconstructed structures. Last but not least, we will focus on the open theoretical issues from the viewpoint of regularization of ill-posed problems.

7 Acknowledgments

References

[1] F. Natterer and F. Wübbeling. *Mathematical Methods in Image Reconstruction*. SIAM, Philadelphia, 2001.

[2] G. Herman and A. Kuba, editors. *Discrete Tomography: Foundations, Algorithms, and Applications*. Springer, 1999.

[3] A. Kuba and G.T. Herman. Discrete tomography: A historical overview. In G.T. Herman and A. Kuba, editors, *Discrete Tomography*, pages 3–34. Birkhäuser, 1999.

[4] R.J. Gardner and P. Gritzmann. Discrete tomography: Determination of finite sets by x-rays. *Trans. Amer. Math. Soc.*, 349(6):2271–2295, 1997.

[5] P. Gritzmann, D. Prangenberg, S. de Vries, and M. Wiegelmann. Success and failure of certain reconstruction and uniqueness algorithms in discrete tomography. *Int. J. Imag. Syst. Technol.*, 9:101–109, 1998.

[6] P. Fishburn, P. Schwander, L. Shepp, and R. Vanderbei. The discrete radon transform and its approximate inversion via linear programming. *Discr. Appl. Math.*, 75:39–61, 1997.

[7] P. Gritzmann, S. de Vries, and M. Wiegelmann. Approximating binary images from discrete X-rays. *SIAM J. Optimization*, 11(2):522–546, 2000.

[8] J.M. Kleinberg and E. Tardos. Approximation algorithms for classification problems with pairwise relationships: Metric labeling and Markov random fields. In *IEEE Symp. Foundations of Comp. Science*, pages 14–23, 1999.

[9] M. Bertero, T. Poggio, and V. Torre. Ill-posed problems in early vision. *Proc. IEEE*, 76:869–889, 1988.

[10] Special issue on variational and level set methods in computer vision. Int. J. Computer Vision, Dec. 2002.

[11] S. Matej, G.T. Herman, and A. Vardi. Binary tomography on the hexagonal grid using gibbs priors. *Int. J. Imag. Syst. Technol.*, 9:126–131, 1998.

[12] M.T. Chan, G.T. Herman, and E. Levitan. Bayesian image reconstruction using image-modeling gibbs priors. *Int. J. Imag. Syst. Technol.*, 9:85–98, 1998.

[13] T. Frese, C.A. Bouman, and K. Sauer. Multiscale bayesian methods for discrete tomography. In G.T. Herman and A. Kuba, editors, *Discrete Tomography*, pages 237–264. Birkhäuser, 1999.

[14] Y. Censor. Binary steering in discrete tomography reconstruction with sequential and simultaneous iterative algorithms. *Lin. Algebra and its Appl.*, 339:111–124, 2001.

[15] Y. Censor and S.A. Zenios. *Parallel Optimization: Theory, Algorithms, and Applications*. Oxford Univ. Press, New York, 1998.

[16] Lauritsch G, Haerer W. *Theoretical framework for filtered back projection in tomosynthesis*. In Proceedings SPIE Conference In Image Processing 3338 San Diego (1998), 1127.

[17] Y. Censor, D. Gordon, and R. Gordon. Component averaging: An efficient iterative parallel algorithm for large and sparse unstructured problems. *Parallel Computing*, 27:777–808, 2001.
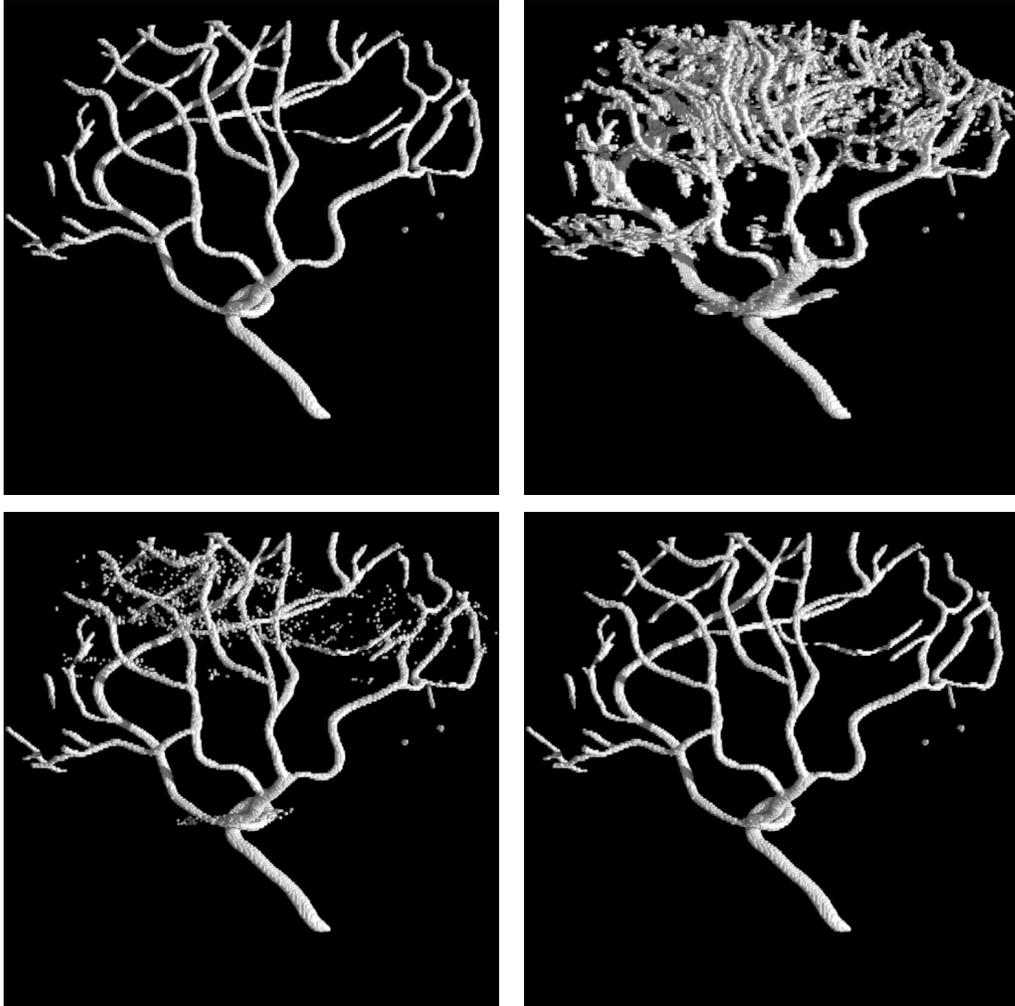
Figure 6: **Top, left:** Ground truth. **Top right:** Locations corresponding to the unknown variables after the preprocessing phase decribed in section 2.1.3. **Bottom, left:** 3D-reconstruction using (BIF), (6). **Bottom, right:** 3D-reconstruction using $(R - BIF)$, (7).
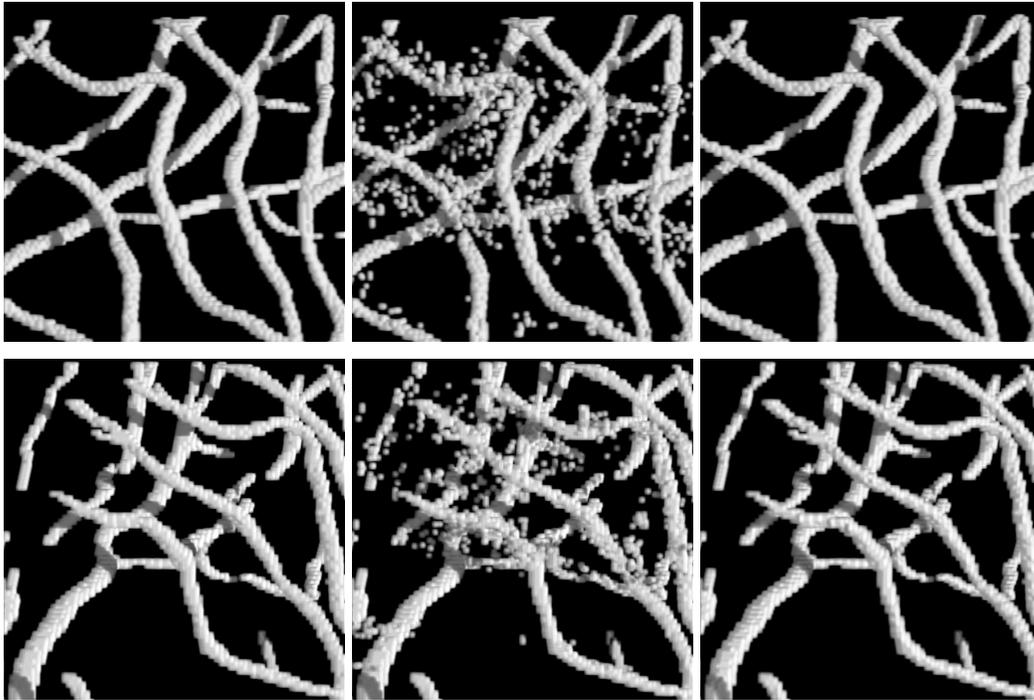
Figure 7: **Top:** Close-ups of the results shown in Figure 6. From left to right: Ground truth, 3D-reconstruction using (BIF), (6), 3D-reconstruction using $(R - BIF)$, (7). **Bottom:** The same for another volume section.
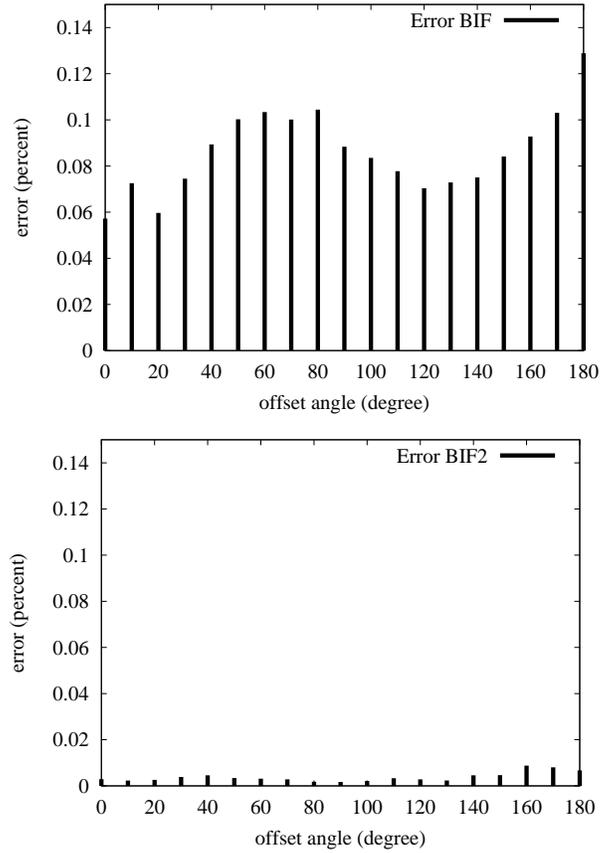
Figure 8: Reconstruction errors in terms of the precentage of voxels differing from ground truth for the vessel structure. Each bar indicates the error for the experiment with the corresponding offset in degrees (cf. Figure 4). **Top:** Errors for the 3D-reconstruction using (BIF), (6). **Bottom:** Errors for the 3D-reconstruction using $(R - BIF)$, (7).