

Image reconstruction and correction methods in neutron and X-ray tomography*

Zoltán Kiss,[†] Lajos Rodek,[†] and Attila Kuba[†]

Abstract

Neutron and X-ray tomography are imaging techniques for getting information about the interior of objects in a non-destructive way. They reconstruct cross-sections from projection images of the object being investigated. Due to the properties of the image acquisition system, the projection images are distorted by several artifacts, and these reduce the quality of the reconstruction. In order to eliminate these harmful effects the projection images should be corrected before reconstruction. Taking projections is usually an expensive and time consuming procedure. One of our main goals has been to try to minimize the number of projections – for example, by exploiting more a priori information. A possible way of reducing the number of projections is by the application of discrete tomographic methods. In this case a special class of objects can be reconstructed, consisting of only a few homogenous materials that can be characterized by known discrete absorption values. To this end we have implemented two reconstruction methods. One is able to reconstruct objects consisting of cylinders and spheres made of homogeneous materials only. The other method is a general one in the sense that it can be used for reconstructing any shape. Simulations on phantoms and physical measurements were carried out and the results are presented here.

1 Introduction

Nowadays it is an interesting common task of physics and image processing to get information about the interior of an object without damaging it in any way. For this purpose several kinds of physical methods are deployed like X-ray, gamma, or neutron imaging. In industrial metallic specimen examination, neutron radiation is generally used. Tomography is an imaging tool for reconstructing objects from their projection images. (A brief review of the principles of tomography is given in Section 2.) However, the acquisition of neutron projections is a very time consuming

*This work was supported by the NSF grant DMS0306215 (Aspects of Discrete Tomography) and the OTKA grant T 048476 (New Aspects and Applications of Discrete Tomography in Neutron Radiography).

[†]Department of Image Processing and Computer Graphics, University of Szeged, H-6720 Szeged Árpád tér 2. E-mail: {kissz, rodek, kuba}@inf.u-szeged.hu

and expensive procedure, hence the goal is to do the reconstruction using as few projections as possible.

Discrete tomography (DT) is a special field of tomography where the object to be reconstructed consists of a small number of homogeneous regions with known absorption coefficients. For example, if the object is made of pure iron, the number of regions is 2 (iron and air) and the reconstructed function can have only two values: the absorption coefficients of iron and air. In DT we cleverly use the information that the function has a known discrete range. This is the main difference between the DT and the classical reconstruction techniques, as in the latter the function/object can in general have arbitrary (non-negative) values. (Comparison with other reconstruction techniques can be found in [7].) The knowledge of known, discrete absorption values may allow us, using DT methods, to reconstruct such objects from a small number of projections (e.g. 2-4) and/or to improve the quality of the reconstruction (more details about DT see [5]).

Two DT methods have been developed at our department. Both consider the reconstruction problem as an optimization task, but the main difference between the two methods lies in the object representation. One of them represents the object being investigated as a digital image and the other as a parameterized geometrical model. These two methods and results obtained using them in phantom/measured experiments are presented in Sections 5, 8 and 9, respectively. Both techniques have been integrated into the system called DIRECT (DIcrete REConstruction Techniques) [14], which is a framework available on-line that incorporates various DT methods being developed at our department.

Since the neutron images we get are usually distorted by several different effects, it is common practice to apply some correction methods before doing any reconstruction from the acquired projections. These effects are usually due to the incorrect settings of the image acquisition apparatus and the physical properties of the radiation used. These distorting effects and a description of the proposed correction techniques are given in Section 4, and our success in dealing with the former is outlined in Section 7.

2 Basic definitions

In this section we shall give a brief overview of the mathematical foundations of tomography, describing the methods, techniques, and algorithms used in this paper.

In non-destructive testing (NDT) several kinds of objects are imaged by some transmission rays like X-rays or neutron rays. The rays transmitted through the object are then partially absorbed. The relation between the initial (unabsorbed) and transmitted intensities, I_S and I_D respectively, can be expressed as a function that depends on the absorption coefficient (μ) of the object. Namely,

$$I_D(s, \vartheta) = I_S \cdot e^{-\int_s^D \mu(u) du} . \quad (1)$$

This equation is a basic relation in *transmission tomography*, where the cross-

sections of the object being studied are to be determined from such measurements. Mathematically the transmission tomography is modelled by the *Radon transformation*, giving the line integrals of a two-dimensional integrable function f , denoted by $\mathcal{R}f$. Formally,

$$[\mathcal{R}f](s, \vartheta) = \int_{-\infty}^{\infty} f(x, y) du, \quad (2)$$

where s and u denote the variables of the coordinate system rotated by ϑ . The function $\mathcal{R}f$ for a fixed value of ϑ is also called the ϑ -angle *projection* of f .

Let f denote the absorption coefficients of the 2D object being studied. Then the ϑ -angle projection of f can be computed from the transmission measurements after a suitable logarithmic transformation. That is

$$[\mathcal{R}f](s, \vartheta) = g(s, \vartheta) = \ln \frac{I_S}{I_D(s, \vartheta)}. \quad (3)$$

Then the *reconstruction problem* can be posed as one where the goal is to find a function f such that its projections are equal to some given function $g(s, \vartheta)$. In other words, we are looking for the inverse of \mathcal{R} , i.e., $\mathcal{R}^{-1}g$. The function f is sometimes called the *image function*, or briefly the *image*.

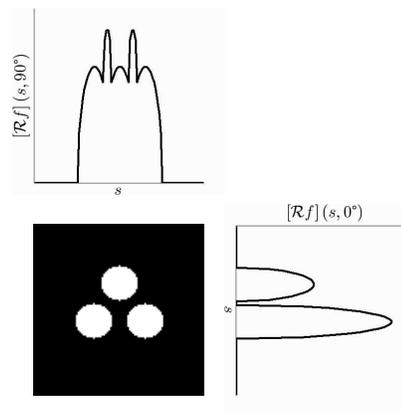


Figure 1: Horizontal and vertical projections of a binary image (black: 0, white: 1).

In this paper we are interested in a special kind of tomography called *discrete tomography*, where the range of the function to be reconstructed is a known discrete set. Such information is usually available in NDT, where the materials (and their absorption coefficients) of the object being studied are known. The simplest example of DT is when the range of the function f consists of only two values, like 0 and 1. In this case the image function is a *binary image* (see Fig. 1). A summary of the theory, algorithms, and applications of DT is given in [5].

3 Neutron and X-ray radiography

In order to get an image of an object, several kinds of radiation (gamma, neutron, X-ray) can be used. The principle of the apparatus of neutron radiography presented here is widely used nowadays, but it is employed in other imaging techniques as well (Fig. 2).

The object to be investigated is placed on a rotating table. The table can be rotated by a PC-controlled stepper motor, thus letting the beams pass through the object in different directions. The beams attenuated by the object impact on a scintillator, which then transforms the detected radiation into visible light detected by a CCD camera. Since the camera can be damaged by direct radiation, an optical mirror system conveys the light from the scintillator to the CCD camera. The images taken by the camera are temporarily stored by the camera controller, and finally a dedicated PC reads out the raw image data from it. A more thorough description of the imaging apparatus can be found in [1].

In this paper we will assume that the radiation source emits parallel beams. This assumption is not unrealistic since if the object being investigated is far enough from the reactor core, the transmitted beams will be almost parallel, so no significant geometric distortions will be introduced into the projections.

The remaining part of this section will discuss various artifacts often encountered during image acquisition.

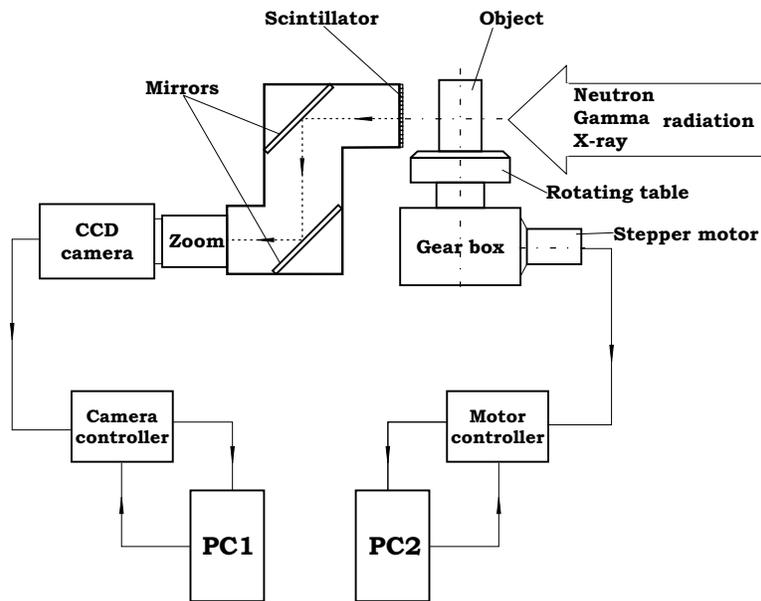


Figure 2: Apparatus for collecting projections.

3.1 Imaging artifacts

In neutron as well as in X-ray radiography the projection images can be distorted by several artifacts due to non-perfect imaging, errors of measurement, and errors in the model. These distortions should be corrected as best one can before or during the reconstruction, since reconstruction is a noise amplifying procedure.

Some of the distortions are due to the properties of the image acquisition system. For example, if the detector system is not uniformly sensitive in the whole field of view of projections, certain areas may be brighter, while others may be darker. This *nonuniformity* may cause ring artifacts in the reconstruction.

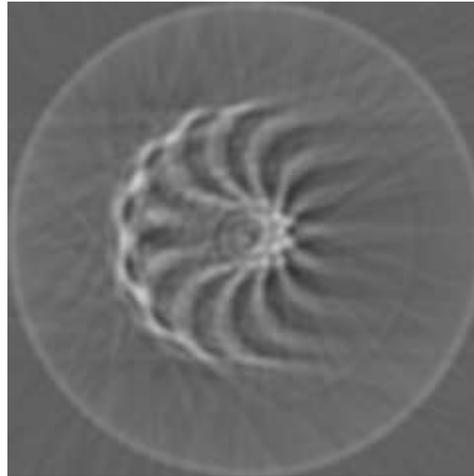


Figure 3: One of the reconstructed cross-sections of a Vidicon tube without applying the correction step.

A further problem might be when the *intensity* of the rays or the *sensitivity* of the camera changes during the acquisition period (cold or warm camera electronics may cause such effects, for instance). As a consequence of changing sensitivity, brighter or darker projections may be acquired and such images with artifacts may be reconstructed (see Fig. 3).

Another source of artifacts might be when the projections are taken not exactly as they should be in their necessary positions. For example, when the projection of the axis of rotation is not exactly in the centerline of the projection images. This *center of rotation* problem may blur the contours in the reconstructed images.

In our experience it is common for the projection images to randomly contain *white isolated points* owing to some problems with the detector system. For example, some pixels may be burnt out in the detector plane.

4 Pre-processing

The input data of the reconstruction procedure is the set of projection images taken from different directions. In neutron tomography these images are not suitable for direct reconstruction, as is shown in Section 3.1. Hence different pre-processing steps are needed prior to the reconstruction procedure. The pre-processing steps in our system are the following.

1. *Logarithmic transformation* In order to get an approximate values of the line integrals of the object along the transmitting rays, we first have to perform a logarithmic transformation on the measured intensity values, as given in (3).
2. *Cropping.* The projections of the object being investigated often cover only a small part of the whole acquired images, so the relevant part is selected and cropped from all projections (see Fig. 9). Only the cropped projections with a smaller image size are used for the further pre-processing steps and reconstruction. The reconstruction from cropped projections requires less memory and computational time.
3. *Motion correction.* It can also happen that the settings of the projection images were not perfect and that some of the images were not taken from the right position. A consequence might be that the images are the rotated or translated versions of the correct ones. In general, such distortions cannot be corrected if we do not know anything about the movement during acquisition or we have no information about the object to be reconstructed. Otherwise, there is some chance of correcting these artifacts. For example, if the data acquisition can be repeated (with the same or with some other object) and the same setting errors occurred, then the translations and rotations can be estimated and performed as a pre-processing step (e.g. ‘center of rotation correction’ [8]). In another case, if we know that the object is circularly symmetric (e.g. a circle or ring) then the projections from any direction are very similar and in this way their right position can be determined by finding the suitable geometric transformations between images, which can be done by a kind of registration [3]. Actually, two methods have been implemented in our system. Both correction methods can be separated into two sub-steps:
 - a) An estimation of parameters of the correction transformation.
 - b) The execution of the corrections transformation, obtaining a new, corrected projection sequence employed in further steps.

The difference between the two methods lies in the choice of the transformation:

- a) The first method can be applied if the positioning errors are the same during each acquisition and they can be described such that the projection of the rotation axis translates only in the projection images (along

a sine curve). In this case the necessary transformation is the translation of each projection image. Its parameter can be estimated from the projection images of a previous experiment and the actual projections can be moved to the right positions.

- b) The second method presumes that the projections of the object are very similar (e.g. the projections of a tube from directions perpendicular to the tube's axis) and by performing registration on all projections (see [13]) we can estimate the correct settings of the projections. In our case the transformation is rigid. If some of the projections are to be translated or rotated following the registration then our correction program is able to perform the necessary geometric transformation.

4. *Homogeneity correction.* Sometimes the detector plane is not uniformly sensitive in the whole field of view. This problem can be lessened if an 'empty' image is available. The empty image is acquired by imaging a homogeneous neutron flux. If the detector system is uniformly sensitive then this image is almost constant. Otherwise, the empty image shows how much correction (multiplication) is necessary, pixel by pixel on each projection, in order to obtain more constant images. The correction can be described mathematically in the following way. For each pixel i of all P_{ϑ} projections

$$R_{\vartheta}(i) = P_{\vartheta}(i) \cdot \frac{1}{P_{empty}(i)},$$

where R is the homogeneity corrected image, and P_{empty} is the empty image.

5. *Intensity correction.* It might happen that the total intensities of the projection images vary during the acquisition period. The reason could be variations in the neutron flux or the electronics of the camera. When this occurs the images should be multiplied by suitable constants such that the average intensity in the background area is roughly the same. After this correction the flickering, which is often visible while playing the projection sequence like a movie, diminishes. This correction step can be divided into two sub-steps as well:

- a) The calculation of the correction factors for each projection on a selected background area.
 b) The execution of the correction method, which yields a new corrected projection sequence for further correction steps and reconstruction.

6. *Isolated points.* The neutron projections often contain white pixels (as shown in Fig. 11(b)) and some statistical noise, which both appear as isolated points having a very different intensity value compared with its neighborhood. In order to eliminate this kind of noise in the projections we performed thresholded median filtering. For each pixel i of each P_{ϑ} projection

$$R_{\vartheta}(i) = \begin{cases} P_{\vartheta}(i) & \text{if } |P_{\vartheta}(i) - \text{med}(\text{NRH}(P_{\vartheta}, i, n))| \leq \text{thr}, \\ \text{med}(\text{NRH}(P_{\vartheta}, i, n)) & \text{otherwise,} \end{cases}$$

where R is the corrected image, $\text{med}(\cdot)$ is the median operator, thr is a suitable threshold constant, and $\text{NRH}(P, i, n)$ is a set which contains the intensity values of the n -neighborhood pixels of i in the image P . Such an n -neighborhood for $n = 8$ is depicted below.

The results of these pre-processing steps are given in Section 7.

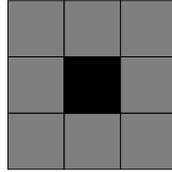


Figure 4: The 8-neighborhood of the black pixel is represented by gray pixels.

5 The reconstruction methods

Both methods essentially treat the reconstruction problem as an optimization task. Then we have to minimize the following objective functional

$$\Phi(f) = \sum_{\vartheta} \|\mathcal{R}f(\vartheta) - P_{\vartheta}\|^2 + \gamma \|f - f_0\|^2, \quad (4)$$

where P_{ϑ} denotes the input projection of angle ϑ , f is the two-dimensional image function that approximates the solution, $\mathcal{R}f(\vartheta)$ denotes the projection of the image f taken at angle ϑ , $\|\cdot\|$ is the Euclidean norm, and $\gamma \geq 0$ is the so-called regularization parameter. Lastly f_0 is the prototype function. It is an image function like f that has the same domain and range, and is similar to the expected reconstruction result.

In (4) the first term is a function, which represents the distance between the projections of the current f and the given projection data P_{ϑ} . This term expresses how well the projections of the current f approximate the given projections. It should be noted that such reconstruction problems occur quite frequently; some typical applications are described in [5], say.

As mentioned earlier, the smaller the number of projections used, the more a priori information should be exploited. Some pieces of a priori information can be incorporated into the second term of (4). One important piece of information – if the prototype of the object is known – is a formula for f_0 . The difference between the actual f and a given prototype object f_0 can be described in the second term. If the object f_0 is the zero image, then it tells us that such a solution is sought which is smooth enough. Finally, the non-negative γ regularization parameter is suitable for balancing the relative importance of the first and second terms. If γ is big, the optimization procedure returns a solution that is more similar to the prototype f_0 and less suitable for the given projection data.

Two reconstruction methods were implemented. The basic difference between them lies in their choice of object representation. The first one is a pixel-based method, while the second one is a kind of parameterized object reconstruction. Both methods reduce the reconstruction to the optimization problem described above, which is solved by applying simulated annealing (SA) [11]. Since simulated annealing is a statistical iterative optimization technique, the result is achieved through a sequence of approximating images, where the $(l + 1)^{th}$ image is constructed by modifying the l^{th} image according to a predefined modification rule. The two reconstruction methods basically differ in the representation of the object f and choice of the modification rule used.

5.1 Pixel-based method

In this method the image f is represented as a *digital image*. The objects take their intensity values from a predefined known discrete set of attenuation coefficients $\{\mu_1, \mu_2, \dots, \mu_d\}$, where d is the number of homogeneous materials from which the object is made. In the case of binary images the set of attenuation coefficients is $\{0, \infty\}$, thus the pixel values can be 0 or 1. In the SA procedure the pixel values will be modified. The modification rule used is simple. Let us take a pixel of f at random and change the 0 or 1 intensity to the other intensity value. More generally, if the image is not binary then the modification rule we apply changes the intensity value μ_i to μ_k at random, where $k \neq i$ and $1 \leq k \leq d$.

The annealing schedule (i.e. setting the parameters) of SA is a cornerstone of successful optimization. (Parameters needed in SA being described in detail in [12], say). One of the most important parameters is the initial temperature. If it is too high the optimization procedure can be too slow. However, with a low temperature there is a possibility that the algorithm will not find the global minimum but will get stuck in a local one. As was pointed out in [4], the optimal initial temperature is 4 degrees centigrade in practice. In this case the choice of the initial image can be arbitrary.

Another important parameter is the annealing function, which determines the speed of the temperature reduction. If the temperature decreases too quickly, the optimization can stop again in a local minimum, but if it decreases too slowly then slow annealing process decelerates the algorithm. The experimentally chosen cooling strategy in our system is defined as follows:

$$t(m + 1) = t(m) \cdot c ,$$

where $t(m)$ returns the temperature (in degrees centigrade) in the m th iteration, and $c \in]0, 1[$ is a feasible constant taking its value from the interval $[0.85, 0.95]$.

The optimization can be accelerated if the initial image is not far from the solution. Then the initial temperature can be decreased and SA finds the optimum faster than it does when starting it from an initial temperature of 4 degrees centigrade. What the right initial temperature is should be investigated further, however.

5.2 Parameter-based reconstruction method

In this method the image function f represents a three-dimensional object consisting of geometric primitives (cylinders and spheres) determined by a small number of parameters. More precisely, we assume that the object to be reconstructed consists of a tube encompassing a solid cylinder called the interior (i.e. the inner space of the tube), which contains a known number of disjoint solid spheres or cylinders made of homogeneous materials. In the present implementation the range of f may consist of at most four distinct values. The object to be reconstructed should consist of the following parts:

- The tube;
- the interior encompassed by the tube;
- the spheres and cylinders contained in the interior;
- and the background surrounding the object, which is usually air or a vacuum.

It should be noted that our aim was to perform a so-called *truly 3D reconstruction*, as opposed to our earlier approaches [9, 10] when the 3D problem was reduced to 2D sub-problems, that is to the reconstruction of 2D cross-sectional slices of the object.

In addition to the a priori information mentioned above, other assumptions were made here. These are the following:

- The cylinders of the wall of the tube are concentric.
- The interior of the tube may contain either spheres or cylinders, but not both. Moreover, their number is known.
- The absorption coefficients are known, at least approximately. In addition, the absorption coefficient of the background is zero or nearly zero, making its effects negligible.
- The projections are taken along parallel beams using equidistant detector spacing.
- The axis of rotation of every cylindrical component is perpendicular to the projecting beams.

Every sphere in the object is uniquely determined by its radius and the coordinates of its center, which are the parameters of the sphere. Cylinders and the tube can be parameterized in a similar way. Furthermore, the absorption coefficients (μ_1, μ_2, μ_3) corresponding to the three materials (i.e. the material of the tube, that of the interior, and that of the spheres and cylinders) are also considered as parameters. (Note here that the absorption of the background material is assumed to be zero.) Such objects can then be uniquely described by a vector of parameters called a *configuration*.

Our second reconstruction method is also an iterative procedure. We start with an initial configuration and during the iterations the current configuration (i.e. the current image) is altered in order to get a hopefully better approximation of the object to be reconstructed. The current image function f is considered *admissible* if certain geometric constraints are met. These are the following:

- The cylinders of the tube remain concentric, the tube is contained by image f , and it has a non-negative wall thickness. Note that allowing a zero thickness may be useful in the case when the object is in fact a solid cylinder containing some spheres or cylinders. This property is used, for instance, in the physical experiments shown in Section 9.2.
- The radius of the interior equals the inner radius of the tube, so that the interior completely fills the inner volume of the tube.
- All distance parameters (namely radii, heights) are positive numbers.
- All the spheres and cylinders are located within the interior, and they are pairwise disjoint.

It is a relatively straightforward task to check all of these conditions during the modification of the current configuration.

It is often desirable to generate random configurations automatically for testing and simulation studies – see Section 9.1, say. Our implementation allows this by letting the user specify how the various parameters may vary, then randomly generating a new configuration based on these settings, while maintaining the above-mentioned geometric restrictions.

The optimization of the functional $\Phi(f)$ is performed in the parameter space. At each iteration step, a new configuration is built by modifying one of the parameters of the current configuration. This is done so that every parameter may be selected with the same probability (i.e. uniformly), and the amount of modification is chosen randomly (again, uniformly) from the whole set of feasible values. Those configurations which are inadmissible are always rejected. The projections $[\mathcal{R}f](s, \vartheta)$ of f are then calculated analytically. It should be remarked that the computation of $\Phi(f)$ can be speeded up by some orders of magnitude by restricting calculations to the sub-domain in which the value of $\Phi(f)$ has changed in the last iteration step.

In order to keep the number of configuration changes as low as possible it is important to find a good initial configuration. The SA algorithm can probably find the optimum in a fewer number of steps when an initial image is quite near the optimum. The initial configuration for our system is constructed as follows.

1. First, the initial locations of the tube and its interior are estimated from the projections (see Fig. 5(a)). This is done after using filtered versions of the input projections, obtained by applying a Gaussian and several averaging filtering kernels to every projection. The boundary of the projection of the tube, namely the coordinates of its centerline, its radius, and its upper and

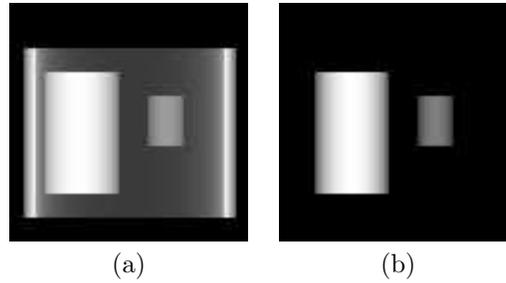


Figure 5: (a) One of the projections of the original object. (b) One of the reduced projections generated by subtracting the projection of the tube and its interior from the projection shown in (a).

lower ends can be found by defining a threshold for the background noise, and scanning the filtered projections for the leftmost, rightmost, uppermost, and lowermost values above this threshold. The boundary of the interior of the tube can be determined in a similar way. The threshold is calculated as a user-specified percentage of the maximum value of the projections, based on an assumed level of noise.

- Using the location of the tube and its interior, their projections can be subtracted from the input projections. This is shown in Fig. 5(b).

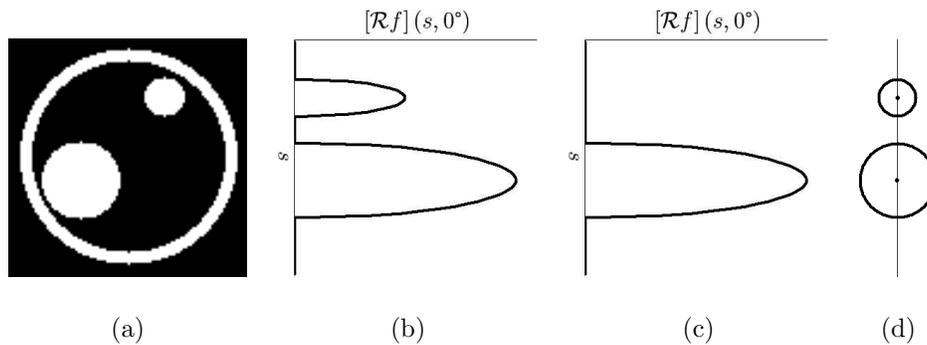


Figure 6: (a) Cross-section of the original object depicted in Fig. 5(a). (b) Cross-section of the reduced projection shown in Fig. 5(b). (c) Elimination of the projection of a disc from the projection displayed in (b). (d) Discs detected in projection (b).

- Next, the initial positions of the spheres or cylinders in the interior are to be guessed based on the reduced projections (see Fig. 6). In the present implementation this task has, for simplicity, been reduced to 2D sub-problems, but it may be possible in future to extend the process to 3D. Currently this procedure is performed in two steps:

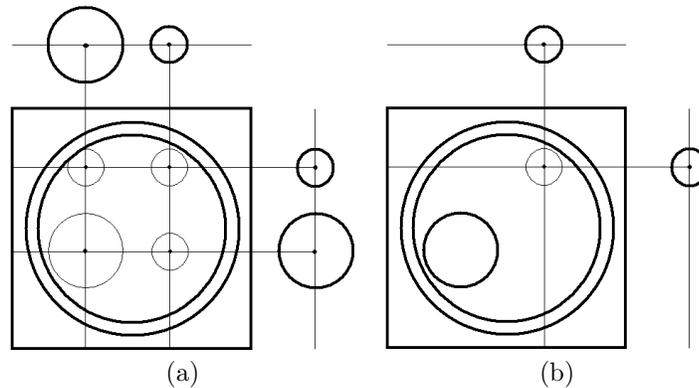


Figure 7: (a) Initial intersections of the cross-section shown in Fig. 5(a) using a horizontal and a vertical projection. (b) Deletion of intersections after choosing the disc with the largest radius.

- a) First, the locations of discs are found in each cross-section separately (see Fig. 6(a)). That is, the intersections of the spheres or cylinders with the 2D cross-sections orthogonal to the axis of rotation have to be found. This is accomplished by the following greedy algorithm based on geometric considerations:
 - i. Possible projections of discs are located within the 1D projections of the cross-sections (see Fig. 6(b)) by the application of model-fitting, i.e. disc parameters are estimated. In particular, a deterministic relaxation algorithm called iterated conditional modes (ICM, [2]) is used to find the parameters of the projection of a disc which best fits the 1D reduced projections. After storing these parameters, the projection of the disc is eliminated from the already reduced projections (see Fig. 6(c)). This process is repeated until no more discs can be detected (see Fig. 6(d)).
 - ii. The centers of the discs found in the previous step are ‘projected back’ into the plane of the 2D cross-section, thus forming several pairs of intersection points and radii of the discs associated with them (see Fig. 7(a)). Those intersections whose distance from one another is below a given limit will be merged into a single intersection. Whatever the case, the radius associated with a particular intersection is calculated by taking the minimum of the radii of the corresponding discs found in the 1D projections.
 - iii. The discs actually located in the 2D cross-section are selected by a greedy strategy: The center of a candidate disc is chosen to be the intersection which was defined by the maximum number of possible projections. Should there be more than one such intersections, the one with the largest radius associated with it will be taken. The in-

tersection chosen is removed instantly from the list of intersections. We should remark here that some additional intersections may be deleted as well in order to retain consistency between the intersections and the 1D projections (see Fig. 7(b)). If this disc results in an admissible configuration, it will be added to the 2D configuration of the cross-section. Otherwise, another intersection will be chosen. This procedure is repeated until no more intersections are left.

- b) Lastly, the initial positions of the spheres or cylinders are determined from the discs found in neighboring cross-sections. Specifically, the discs detected in the previous step are treated as solid cylinders of unit height. An auxiliary 3D configuration can be built by stacking the cross-sections onto each other, that is placing these discs into a 3D coordinate system. This configuration is then examined to find candidate spheres and cylinders that will be included in the initial configuration.

6 Modelling the noise

In order for the conditions to be as realistic as possible in the simulation experiments, artificial noise was generated with a uniform distribution and then added to the projections. An example of this is shown in Fig. 8.

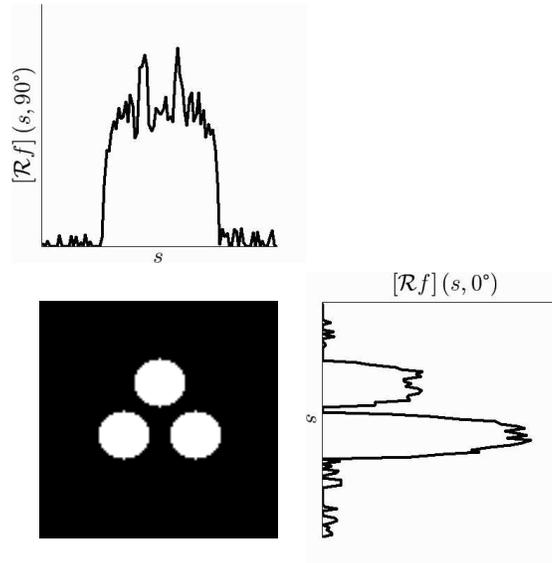


Figure 8: Horizontal and vertical projections of a binary image with 10% noise.

7 Results of correction methods

The pre-processing steps discussed in Section 4 were tried out on the projections of a VIDICON tube (see Fig. 9). There were 360 projections taken in 1° angular steps using X-ray radiation.

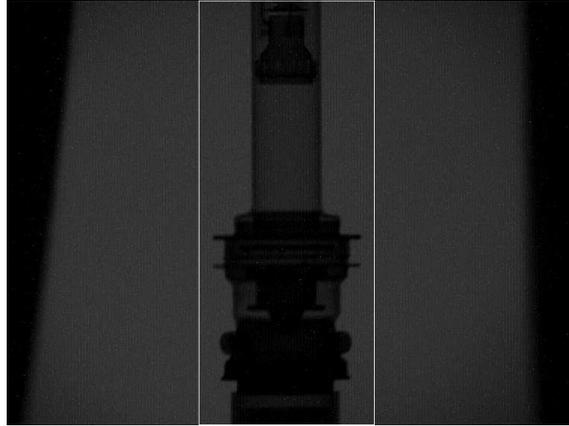


Figure 9: One of the projections of the VIDICON tube (height ≈ 10 cm, width ≈ 4 cm). The white lines indicate the cropped area.

The first step involved cropping the relevant area from the original projections when the object to be reconstructed covered only a small portion of the projections.

The second step of pre-processing was the correction of the motions during the image acquisition. As a result of the analysis of the projections (using registration) we found that in this case each projection had to be translated along a sine curve having a phase of 90° and amplitude of 2 pixels.

When an empty projection image is available as well (as was the case here), homogeneity correction can be performed. Homogeneity correction eliminates the effects of the non-uniformly sensitive detector system.

The results of the pre-processing steps (cropping, motion and uniformity corrections) can be seen in Fig. 11.

The intensity correction step of the pre-processing multiplies the projections by suitable coefficients such that the total intensity of the corrected projections is practically constant for each projection (see Fig. 10).

Finally, the isolated points correction was applied to each projection. This correction means the application of a thresholded median filter (8-neighborhood, $\text{thr} = 20$ assuming a maximum intensity level of 255). In Fig. 11(c) we see that a part of the noise was eliminated from the projections and the reconstructed cross-section became a bit more homogeneous compared with Fig. 11(f). This homogeneity feature is especially noticeable in the interior of the dashed circle.

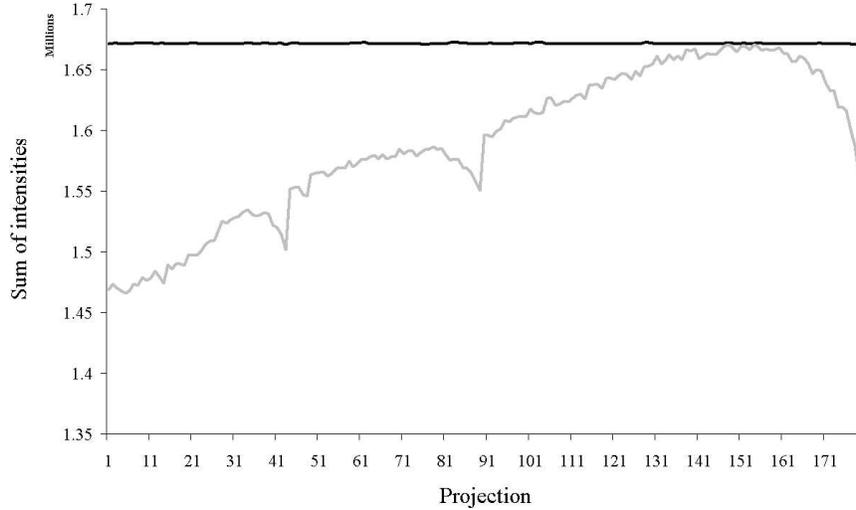


Figure 10: Total intensity values in the first 180 projections before (gray line) and after (black line) intensity correction.

8 Results of the pixel-based reconstruction method

In this section we demonstrate the efficiency of the pixel-based reconstruction method. First, software simulated phantom studies will be presented where we show the effects of changing different parameters on the reconstruction. Then some preliminary results on some neutron tomography data are described in Section 8.2.

In order to measure the accuracy (precision) of the reconstruction quantitatively, we used the relative mean error (RME for short), which is defined here as

$$\text{RME} = \frac{\sum_i |f_i^o - f_i^r|}{\sum_i f_i^o} \cdot 100\% ,$$

where f_i^o and f_i^r denote the value of the i th pixel of the original and the reconstructed image functions, respectively. It is quite clear that $\text{RME} \geq 0$ and a smaller RME value means a better reconstruction result.

8.1 Simulation studies

The simplest case is when the range of the image function to be reconstructed contains only two values, 0 and 1. The results of such reconstructions can be seen in Figs. 12 and 13.

First, we investigated how the number of projections influences the reconstruction (see Fig. 12). It is conspicuous that the pixel-based method was able to almost

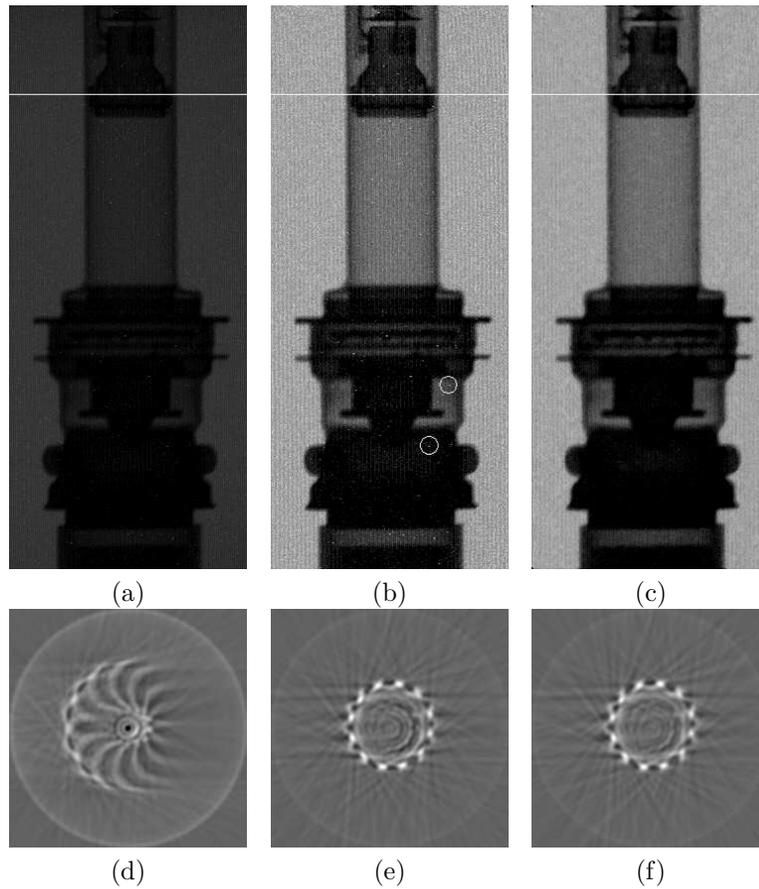


Figure 11: (a) One of the projections of a VIDICON tube after cropping having the size 241×572 . (White line indicates the cross-section shown below.) (b) The same projection after motion and intensity corrections (the white circles indicate isolated points to be corrected in the next step). (c) The same projection as in (b) after homogeneity and isolated points corrections. (d) Reconstruction (241×241) of the cross-section shown after cropping. (e) Reconstruction of the cross-section shown after motion and intensity corrections. (f) Reconstruction of the cross-section shown after homogeneity and isolated points corrections. The reconstruction was performed by the software package SNARK93 [15] (filtered back-projection, cosine filter, cut-off frequency 0.5, Lagrange interpolation).

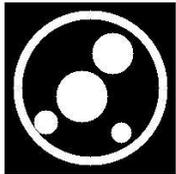
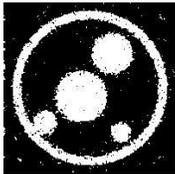
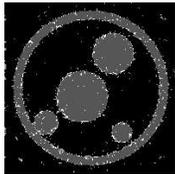
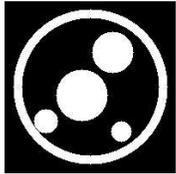
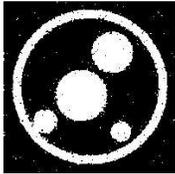
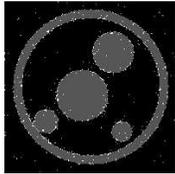
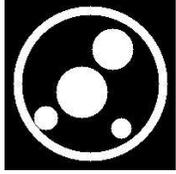
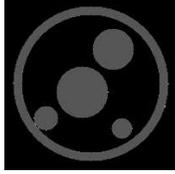
# of projs	Original	Result	Difference	RME (%)
8				12.57
10				6.98
12				0.10

Figure 12: Reconstruction of circles from different numbers of projections by the pixel-based method (0% noise, 200 measurements/projection). First column: number of projections. Second column: original object. Third column: reconstructed images. Fourth column: RME of the original and reconstructed images (black/gray: when the corresponding pixels are black/white in both the original and reconstructed images; white: when the corresponding pixel intensities are different). Fifth column: RME of the original and reconstructed images.

reproduce the original phantom with just 12 projections.

We also studied the effect of noise on the reconstruction. 0%, 10%, and 40% noise was added to the exact projections and then the reconstruction was done with the noisy projections as well (see Fig. 13). As expected, the quality of the reconstruction grew worse the higher the level of noise in the projections was chosen. The same can be seen from the RMEs. But the object, however, is still recognizable even with 40% noise.

8.2 Physical experiments

We had the chance to test our reconstruction methods with real physical data as well. One of them was the battery from a pacemaker (Fig. 14(a)). As can be seen in the image, especially in Fig. 14(b), the projections were almost noiseless and of a good quality. In order to compare the results of the classic FBP technique with our pixel-based one, we reconstructed several slices. One of them is shown in Fig.

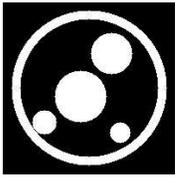
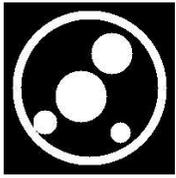
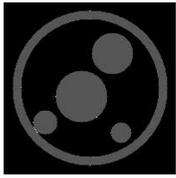
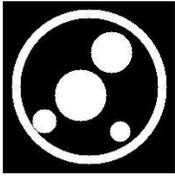
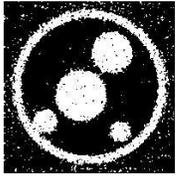
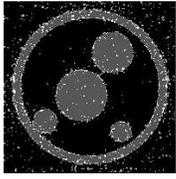
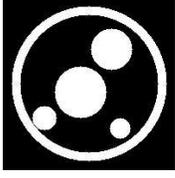
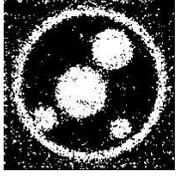
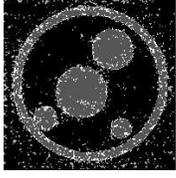
Noise (%)	Original	Result	Difference	RME (%)
0				0.12
10				20.64
40				34.63

Figure 13: Reconstruction of circles from exact and noisy projections by the pixel-based method (16 projections, 200 measurements/projection). First column: noise level. Second column: original object. Third column: reconstructed images. Fourth column: difference images. Fifth column: RME of the original and reconstructed images.

14(a), and the corresponding reconstruction result, based on 200 projections and using the FBP method of SNARK93 [15], can be seen in Fig. 15. This result can be considered quite favourable as the object is reasonably distinguishable.

In order to compare the results of the classical method we performed an FBP using 20 projections, and we ran our pixel-based technique with the same input data as well. In Fig. 14(c) the results of FBP are presented, which contain streaks due to the small number of projections used. In the case of the pixel-based method (Fig. 14(d)) these streaks are absent, but the quality of the results are still worse than the previous one. The reason might be that the DT algorithm did the reconstruction using only 21 intensity levels, but the object could not be considered one which satisfies our basic assumption: the object consists of only a few *inhomogeneous* materials. This inhomogeneity or irregular material distribution is quite visible in Fig. 15. In addition, we were not aware of the exact absorption coefficients. If the object had been made of homogeneous materials, we would have hoped for better results.

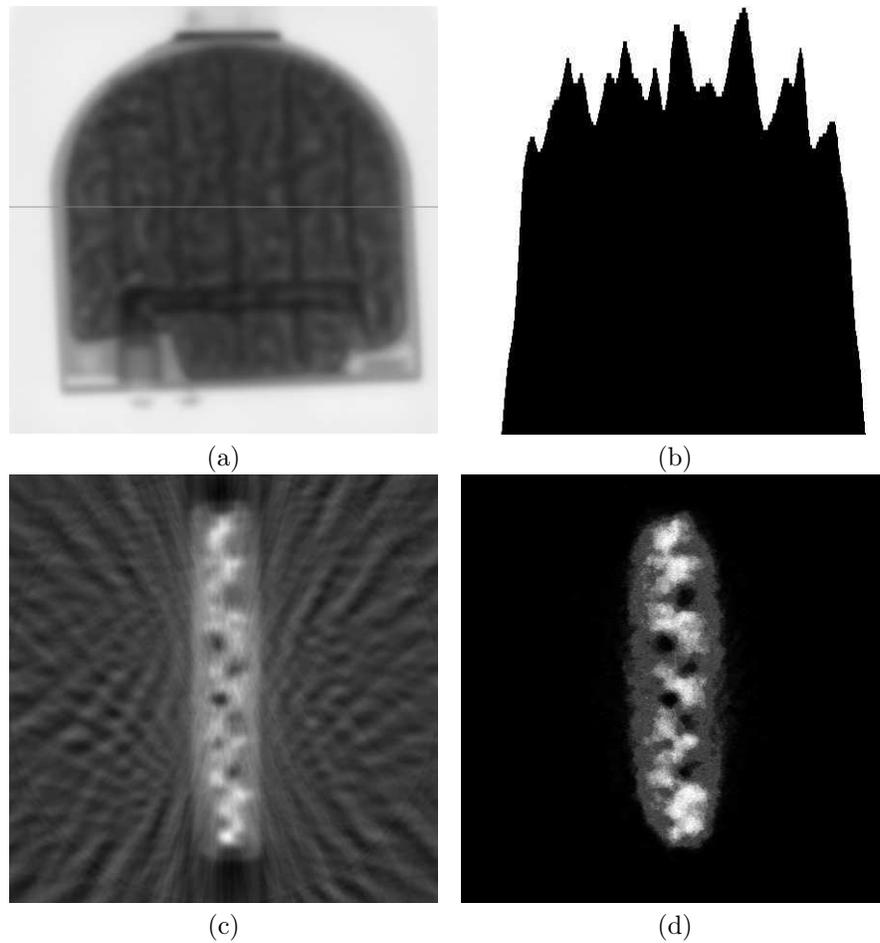


Figure 14: (a) One of the neutron projections of a pacemaker battery. (b) The bar diagram representing the intensity values of the projections in the indicated row in (a). (c) Reconstructed cross-section in the position indicated in the row in (a) (FBP method 20 projections, cosine filter, cutoff 0.5, Lagrange interpolation as done in SNARK93 [15]). (d) Image of the same cross-section reconstructed via the pixel-based method using 20 projections and 21 intensity levels.

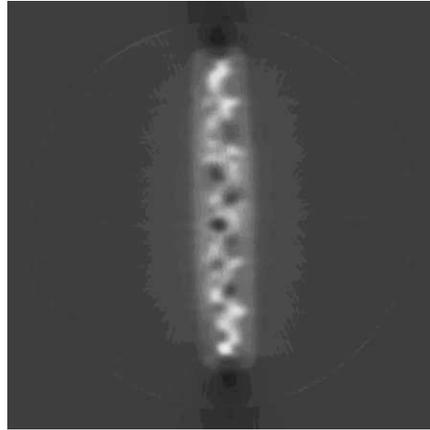


Figure 15: FBP reconstruction (as done in SNARK93 [15]) of the cross-section shown in Fig. 14(a) using 200 projections.

From the point of view of timing the technique cannot be said to be fast. Its speed greatly depends on the input data. In accordance with our experiences the technique terminates within 60 seconds in case of noiseless phantom images on a 3 GHz Intel Pentium 4, where the size of the images is at most 400×400 pixels using less than 20 projections. However, a reconstruction using real projections can cost up to 5–10 minutes.

9 Results of the parameter-based reconstruction method

In this section we will present the results of the parameter-based reconstruction method. First, simulation studies will be discussed, focussing on the effects of various parameters on the reconstruction. These parameters include the geometric complexity of the object, the amount of noise in the projections, and the number of projections. Afterwards, some results of the physical experiments will be given.

In order to assess the accuracy of this method quantitatively, the RME measure defined in Section 8 was used throughout the experiments with a slight modification: the image function was converted to a $\{0, 1\}$ binary image before calculations. It should be mentioned that a more precise way of measuring would be to compute RME analytically instead of using the digitized image, which is one of our future plans. The models of objects are visualized using the Virtual Reality Modeling Language (VRML97 [16]).

9.1 Simulation experiments

Software experiments were performed in order to investigate the effects of key reconstruction parameters: the geometric complexity of the object to be reconstructed, noise level, and the number of projections. Another important thing we need to know is how the value of the objective function $\Phi(f)$ changes during iterations. A typical plot of $\Phi(f)$ is shown in Fig. 16, produced during the reconstruction of 3 spheres from 4 noisy projections. As expected, the objective function decreases rapidly at the beginning of the optimization, and more slowly near the global minimum.

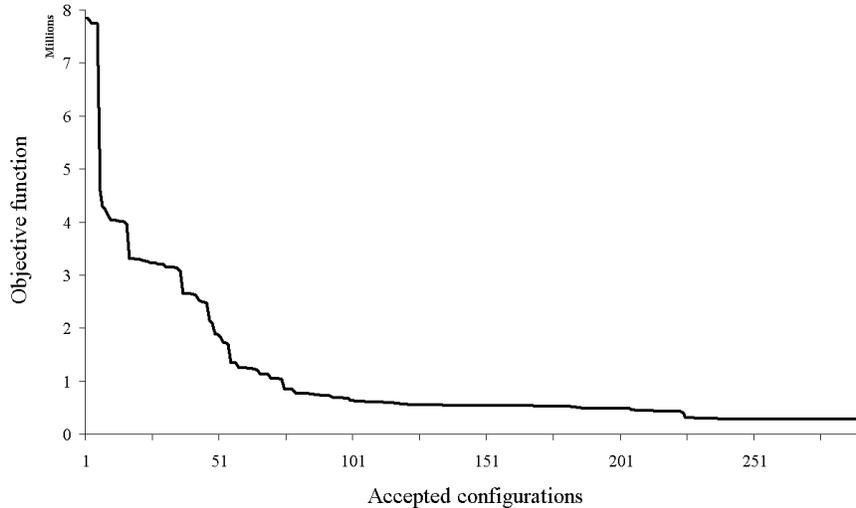


Figure 16: The value of the objective function $\Phi(f)$ as a function of the number of accepted configurations. (The total number of accepted configurations was 291. The reconstruction of 3 spheres using 4 projections degraded with 10% noise, with 100×100 measurements/projection.)

The aim of the first experiment was to study the influence of the geometric complexity of the object (see Fig. 17) using only 2 projections degraded with 10% noise. As can be seen in Fig. 18, it is hard to get an acceptable result with just 5 spheres. The algorithm, however, can successfully reconstruct fewer spheres with good precision.

We also looked at the effects of noise on our reconstruction using the parametric object shown in Fig. 19. The noise level was again set to 0%, 10%, or 40% as we did in the simulation experiments using the pixel-based reconstruction method. It is remarkable that the parameter-based reconstruction method still works even with 40% noise (see Fig. 20).

In the last simulation experiment we analyzed the impact of the number of projections on the results (see Fig. 21). Since real physical measurements are usually

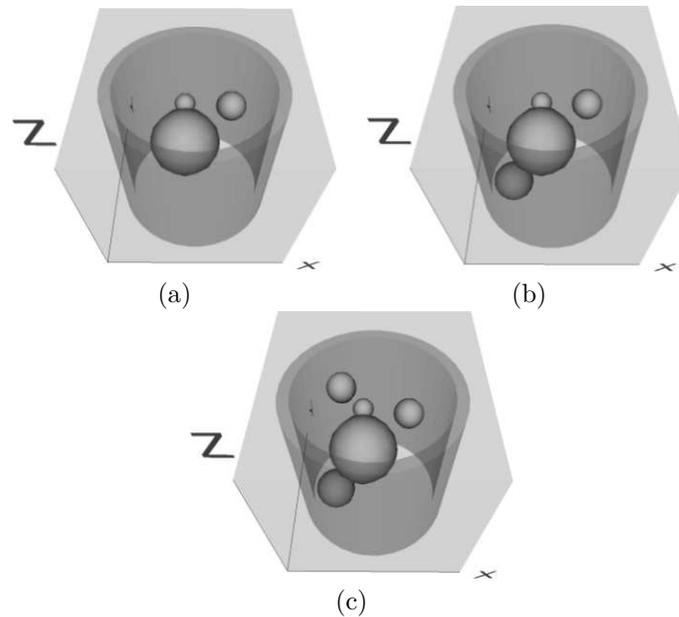


Figure 17: Perspective view of 3D parametric phantoms consisting of (a) 3, (b) 4, and (c) 5 spheres in a tube.

distorted by some noise, the same experiment was performed using noiseless projections as well as adding 10% or 40% noise to the projections. It is quite apparent here that the reconstruction of these parameterized objects does not depend much on the number of projections used. As Fig. 21 indicates, it is usually sufficient to use just two projections for simpler objects (e.g. three spheres). Actually, having more projections does not necessarily improve the precision of the result by much. At the same time, the algorithm is more sensitive to the geometrical complexity of the object to be reconstructed, as is demonstrated by the third row in Fig. 18.

9.2 Physical experiments

In addition to the software-generated data described above, the effectiveness of the parameter-based reconstruction method was also evaluated using physically measured data. However, before proceeding with a discussion of the results, two remarks should be made. First, due to the limitations of the imaging system and the measurement errors, the projections were distorted and quite noisy. To help the algorithm produce the best results possible, it was necessary to perform the pre-processing steps described in Section 4. Second, as the exact values of the absorption coefficients were unknown, they had to be estimated here.

The first physical experiment was performed on a phantom object called the reference cylinder, a diagram of it being shown in Fig. 22. The object is a solid

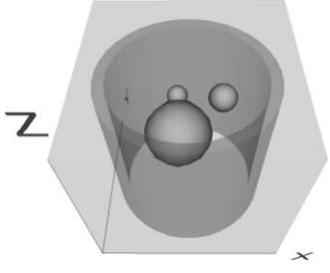
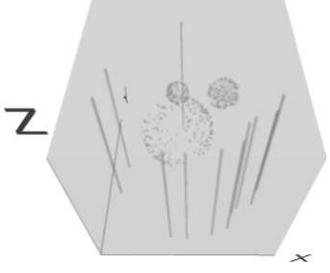
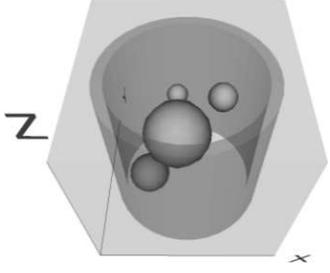
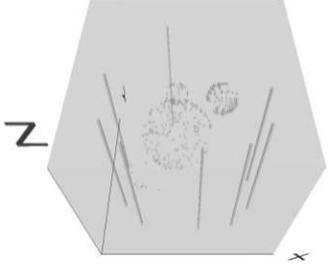
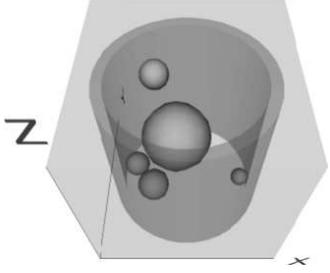
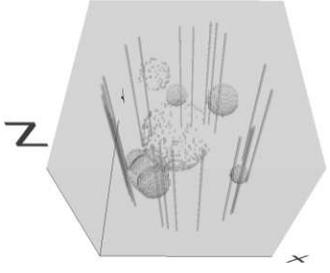
# of sph.	Result	Difference	RME (%)
3			0.97
4			0.66
5			6.27

Figure 18: Reconstruction by the parameter-based method using a different number of spheres (original object: see Fig. 17; parameters: 10% noise, 2 projections, 100×100 measurements/projection). First column: number of spheres. Second column: reconstructed object. Third column: difference between the reconstructed and original object (only mismatching voxels are painted). Fourth column: RME of the original and reconstructed image.

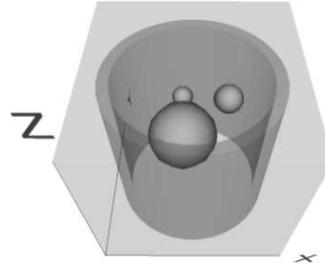


Figure 19: Perspective view of the 3D parametric phantom object consisting of 3 spheres in a tube.

cylinder made of Plexiglas, and it contains three holes of different diameter and various depths in an asymmetric arrangement. The lower part of the hole with the largest depth was filled with aluminum screws. This region is clearly visible on the projection images, which were taken using an X-ray source. One of these projections is shown in Fig. 23(a) after applying every pre-processing step except that for the filtering of isolated points.

Since our model assumes that the cylindrical holes are filled with the same material, the lower half of the projections had to be thrown away (see Fig. 23(b)). The size of the projection images was 155×212 pixels before and 155×113 pixels after cropping, respectively. It is evident that the remaining part of the object is suitable for discrete tomographic reconstruction, since it consists of two homogeneous regions containing air and Plexiglas (background and holes, respectively).

The model reconstructed using 4 projections is illustrated in Fig. 23(c) to Fig. 23(e). Since the exact structure of the object was known, it was possible to create an ‘original’ model, and thus to measure the precision of the reconstruction. The difference between the original and the reconstructed model is shown in Fig. 23(f), with an RME of 2.45%. This relatively high value is due to two facts: the projections were fairly small and noisy, and the exact value of the absorption coefficient of Plexiglas was unknown.

A phantom object very similar to the one mentioned above was used in the second physical experiment. The structure of this object was identical to that shown in Fig. 22, but the solid cylinder here was made of aluminum. The hole with the biggest diameter was partly filled with acetone and the two others contained water.

At first sight this object seems unsuitable for our parametric object model, since the bores consist of three homogeneous regions: acetone, water, and air. It turned out, however, that the absorption coefficients of the acetone and of the water are almost equal for the neutron rays used for image acquisition. The size of the projection images was 365×400 pixels, and one of them can be seen in Fig. 24(a) after applying the same pre-processing steps as in the first experiment. Moreover, those sections of the bores which contained air are hardly recognizable. Hence the object can be considered to be composed of three approximately homogeneous materials: aluminum, fluid, and air (background).

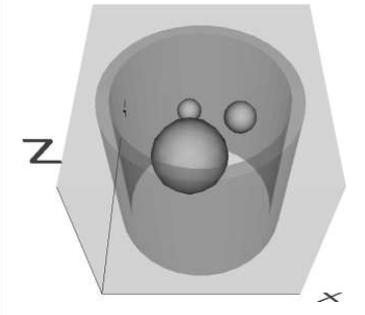
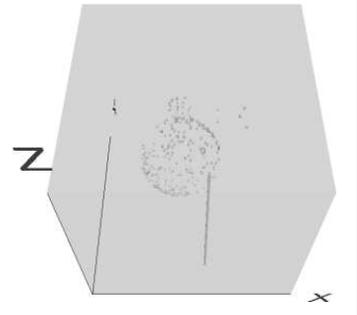
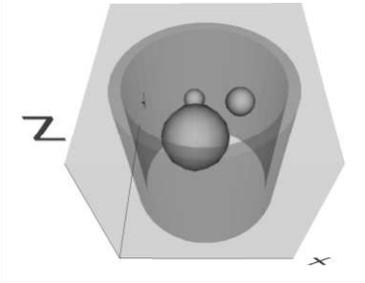
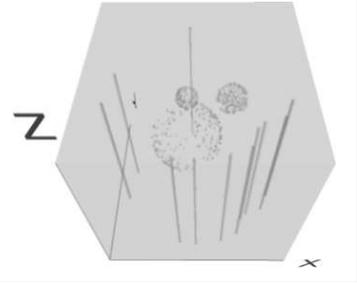
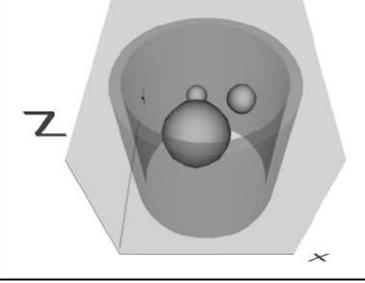
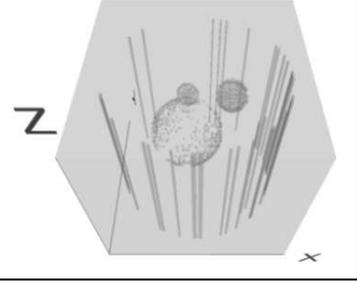
Noise (%)	Result	Difference	RME (%)
0			0.22
10			0.97
40			2.94

Figure 20: Reconstruction by parameter-based method from noise free and noisy projections (original object: see Fig. 19; parameters: 2 projections and 100×100 measurements/projection). First column: noise level. Second column: reconstructed object. Third column: difference between the reconstructed and original object. Fourth column: RME of the original and reconstructed image.

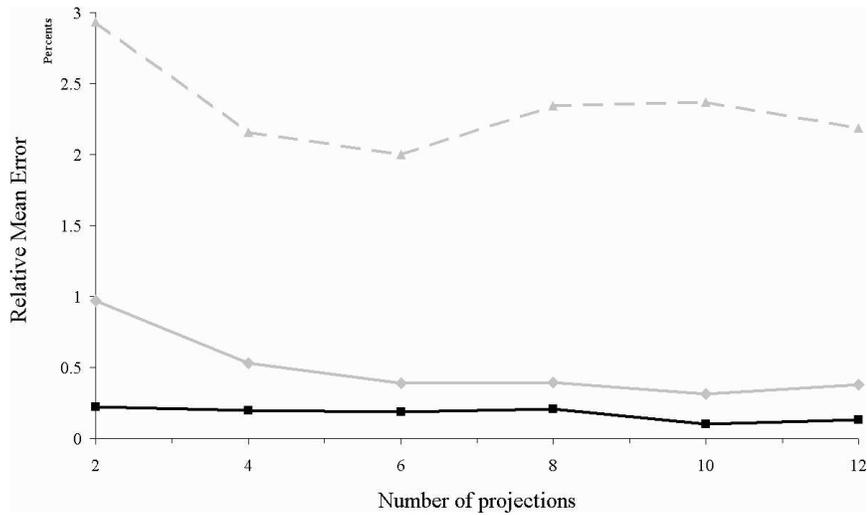


Figure 21: The RME value versus the number of projections without noise (black curve), with 10% (solid gray curve) and with 40% noise (dashed gray curve), respectively. (Reconstruction of 3 spheres using 2, 4, 6, 8, 10 and 12 projections, and 100×100 measurements/projection.)

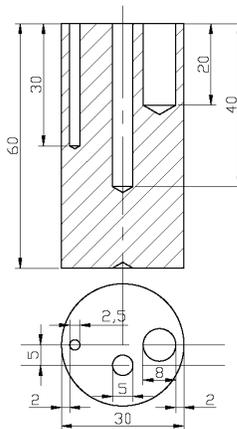


Figure 22: Diagram of the phantom object used in the experiments (dimensions are shown in mm.)

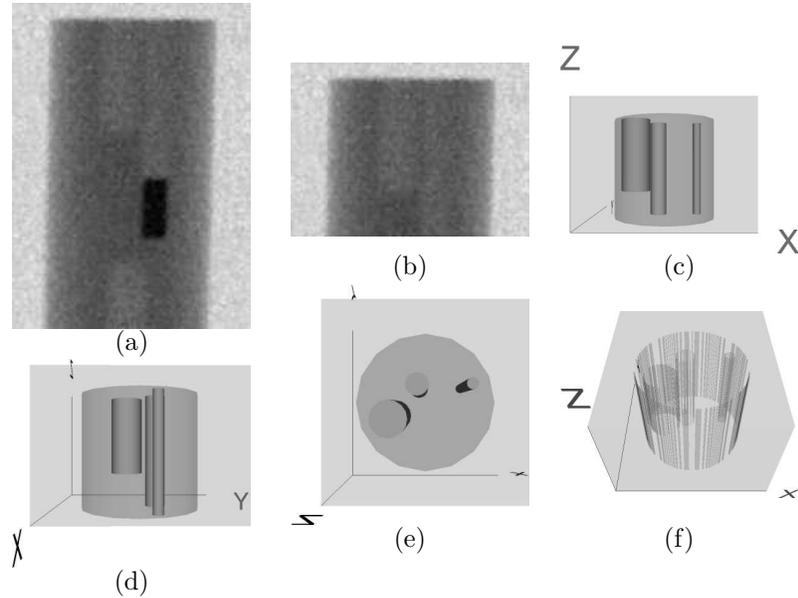


Figure 23: Projection and reconstruction results of the Plexiglas object given in Fig. 22. (a) One of the original projection images. (b) One of the cropped projection images. (c) 0° view of the reconstructed model. (d) 90° view of the reconstructed model. (e) Top-down view of the reconstructed model. (f) Difference between the reconstructed model and the original one (RME = 2.45%).

The model reconstructed using 4 projections is shown in Fig. 24(b), (c) and (d). As the structure of the object is identical to the one used in the previous experiment, an ‘original’ model can be created again, thus allowing one to test the precision of the reconstruction. The difference between the original and the reconstructed model is shown in Fig. 24(e). Even though the precise radius of the bore with the smallest diameter in the reconstructed result is somewhat larger than the real one, the RME is only 1.01%. This is because the RME is calculated using the digitized image function, and the number of voxels is much larger than in the former case. The discrepancies between the original and the resulting models are due to the following: the axis of rotation was precessing, the exact values of the absorption coefficients were unknown and, despite our earlier assumption, the absorption coefficients of water and of acetone were quite different.

It is notable that the parameter-based method performs somewhat faster than its pixel-based counterpart. In particular, typical running times are a few (5–10) seconds on a 3 GHz Intel Pentium 4, using 4 projections and assuming that the size of the image function to be reconstructed is at most $100 \times 100 \times 100$ voxels. On the other hand, when the image function is as large as that in the last physical experiment (i.e. $365 \times 365 \times 400$ voxels), it may take even 3–5 minutes to complete the reconstruction process.

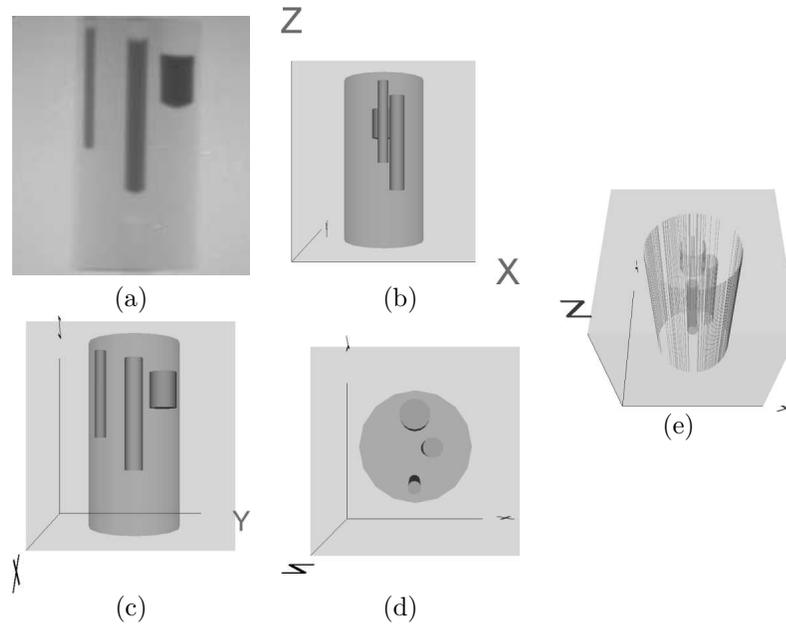


Figure 24: Projection and reconstruction results of the aluminum object given in Fig. 22. (a) One of the original projection images. (b) 0° view of the reconstructed model. (c) 90° view of the reconstructed model. (d) Top-down view of the reconstructed model. (e) Difference between the reconstructed model and the original one (RME = 1.01%).

10 Future plans

Though the techniques presented here produced promising and even acceptable results, we are planning to improve several things in the future. It would be desirable to make pre-processing as automated as possible in both methods, and for the absorption coefficients to be estimated automatically – for instance by using some kind of calibration procedure. For the pixel-based method, the most useful improvements would be to make it less sensitive to noise and to reduce the number of projections necessary for reconstruction. We would also like to make this approach work better for objects containing three or more materials. Further research could be done for situations where the projections are not a parallel beam but fan-beam, say. Finally, some practical extensions of the parameter-based algorithm might allow five or more materials, using more complex prior knowledge of the object, computing the RME analytically, and employing another kind of parametric object representation like a deformable model.

Acknowledgements

This work was supported by the NSF grant DMS0306215 (Aspects of Discrete Tomography) and the OTKA grant T 048476 (New Aspects and Applications of Discrete Tomography in Neutron Radiography). The authors wish to thank Márton Balaskó (KFKI Atomic Energy Research Institute, Budapest) and Prof. W. Treimer (Hahn Meitner Institut, Berlin) for providing us with the neutron and X-ray images of the objects used in the experiments. We would also like to thank David Curley for checking this paper from a linguistic point of view.

References

- [1] M. Balaskó, A. Kuba, A. Nagy, Z. Kiss, L. Rodek, and L. Ruskó, *Neutron-, gamma- and X-ray three-dimensional computed tomography at the Budapest research reactor site*, Nucl. Instrum. Methods Phys. Res. A **542A** (2005), 22–27.
- [2] J. Besag, *On the statistical analysis of dirty pictures*, J. Roy. Statist. Soc. Ser. B **48** (1986), 259–302.
- [3] L. G. Brown, *A survey of image registration techniques*, ACM Computing Surveys **24** (1992), 325–376.
- [4] S. Geman and D. Geman, *Stochastic Relaxation, Gibbs distributions, and the Bayesian Restoration of Images*, IEEE Trans. Pattern Anal. Machine Intell. **6** (1984), 721–741.
- [5] G. T. Herman and A. Kuba (Eds.), *Discrete Tomography: Foundations, Algorithms, and Applications*, Birkhäuser, Boston, MA (1999).
- [6] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, *Optimization by simulated annealing*, Science **220** (1983), 671–680.
- [7] Z. Kiss, L. Rodek, A. Nagy, A. Kuba, and M. Balaskó, *Reconstruction of pixel-based and geometric objects by discrete tomography. Simulation and physical experiments*, Electronic Notes in Discrete Mathematics **20** (2005), 475–491.
- [8] A. Kuba, L. Rodek, Z. Kiss, L. Ruskó, A. Nagy, and M. Balaskó, *Discrete tomography in neutron radiography*, Nucl. Instrum. Methods Phys. Res. A **542A** (2005), 376–382.
- [9] A. Kuba, L. Ruskó, L. Rodek, and Z. Kiss, *Application of Discrete Tomography in Neutron Imaging*, Proc. of 7th World Conference on Neutron Radiography, Rome, Italy (2002), 361–371.
- [10] A. Kuba, L. Ruskó, L. Rodek, and Z. Kiss, *Preliminary studies of discrete tomography in neutron imaging*, IEEE Trans. Nucl. Sci. **52** (2005), 380–385.

- [11] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, *Equation of state calculations by fast computing machines*, J. Chem. Phys. **21** (1953), 1087–1092.
- [12] N. Robert, F. Peyrin, and M. J. Yaffe, *Binary Vascular Reconstruction from a Limited Number of Cone Beam Projections*, Med. Phys. **21** (1994), 1839–1850.
- [13] A. Tanács and A. Kuba, *Evaluation of a Fully Automatic Medical Image Registration Algorithm Based on Mutual Information*, Acta Cybernet. **16** (2003), 327–336.
- [14] <http://www.inf.u-szeged.hu/~direct/>, Homepage of DIRECT framework, *A toolkit for testing and comparing 2D/3D reconstruction methods of discrete tomography*.
- [15] <http://www.cs.gc.cuny.edu/~gherman/snark2001.html>, Homepage of the SNARK93 software system, *A Programming System for 2D Image Reconstruction from Projections*.
- [16] <http://www.web3d.org/x3d/vrml/>, Homepage of VRML97, *Virtual Reality Modeling Language*.

Received July, 2005