



Universiteit
Leiden
The Netherlands

Algebraic filters for filtered backprojection

Plantagie, L.

Citation

Plantagie, L. (2017, April 13). *Algebraic filters for filtered backprojection*. Retrieved from <https://hdl.handle.net/1887/48289>

Version: Not Applicable (or Unknown)

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/48289>

Note: To cite this publication please use the final published version (if applicable).

Cover Page



Universiteit Leiden



The handle <http://hdl.handle.net/1887/48289> holds various files of this Leiden University dissertation

Author: Plantagie, L.

Title: Algebraic filters for filtered backprojection

Issue Date: 2017-04-13

APPROXIMATING ALGEBRAIC TOMOGRAPHY
METHODS BY FILTERED BACKPROJECTION: A
LOCAL FILTER APPROACH

Abstract – Filtered Backprojection is the most widely used reconstruction method in transmission tomography. The algorithm is computationally efficient, but requires a large number of low-noise projections acquired over the full angular range to produce accurate reconstructions. Algebraic reconstruction methods on the other hand are in general more robust with respect to noise and can incorporate the available angular range in the underlying projection model. A drawback of these methods is their higher computational cost.

In a recent article, we demonstrated that for *linear algebraic reconstruction methods*, a filter can be computed such that applying Filtered Backprojection using this filter yields reconstructions that approximate the algebraic method. In the present work, we explore a modification of this approach, where we use more than one algebraic filter in the reconstructions, each covering a different region of the reconstruction grid. We report the results of a series of experiments to determine the how well the reconstruction and approximation accuracy of this approach.

4.1 INTRODUCTION

Computed Tomography is a general technique for recovering an image of the interior of a scanned object from its projections, acquired from a range of angles. In this article, we focus on *transmission tomography*, where the source and detector of the scanning device are located on opposite sides outside the object. The source emits a penetrating beam that traverses the object. Depending on the thickness and composition of the materials on its path, part of the incoming beam is absorbed or scattered by the object. The measurements at the detector can be considered as *projections* of the object. This data forms the input for a *tomographic reconstruction algorithm*, which computes an image of the object from all available projections [1–3].

Most of the reconstruction algorithms proposed in the literature can be divided into two classes: (i) analytical reconstruction methods and (ii) algebraic reconstruction methods. The analytical reconstruction methods are based on discretizations of an inversion formula for the idealized continuous representation of the reconstruction problem [4]. The algebraic reconstruction methods are derived from a discrete approach of the reconstruction problem, which is formulated as a system of linear equations. Due to the large size of the problem, iterative methods are used to solve the problem instead of calculating a least-squares solution directly.

The best known example of an analytical reconstruction method is *Filtered Backprojection* [5]. Due to its fast computation time and robustness with respect to noise, it is often used to obtain reconstructions. Variants of FBP, such as the FDK algorithm for cone-beam reconstruction [6], are also highly popular in practice. A disadvantage of these reconstruction methods is their poor reconstruction quality when there are only few projection angles or when the angular range is limited.

Algebraic reconstruction algorithms, such as ART, SART, and SIRT [1, 7, 8], typically yield more accurate results in such limited data scenarios, as modeling the problem as a system of linear equations allows to encode exactly which information is known about the object, instead of interpolating data that has not been measured. The main drawback of iterative algebraic methods is that they are more computationally intensive, requiring a few times up to a hundred times the computation time of FBP.

An alternative route to obtaining more accurate reconstructions is to adjust the filter of FBP to specific properties of the projection data, the set of projection angles, or both [9–12]. In a recent article, we demonstrated that the operation of certain algebraic methods (those that are *linear*) can be approximated by applying FBP with a custom filter. This filter is computed once and depends on the particular algebraic method that is to be approximated. This method is known as *algebraic filter-filtered backprojection* (AF-FBP) [13]. For a particular pixel, the reconstructed image that results from applying FBP with this custom filter has exactly the same value as in the algebraic reconstruction. The AF-FBP reconstructions have been demonstrated to approximate the accuracy and robustness of the underlying algebraic method, while the computational efficiency is comparable to that of the fast FBP reconstruction method.

The AF-FBP method is based on selecting a particular pixel in advance, for which the filter is computed. This filter is subsequently applied to reconstruct all other image pixels as well. A logical question is therefore whether the AF-FBP method can be improved by computing *multiple* filters, at different locations in the image, and combining the FBP results for the set of computed filters.

In this article we examine the performance of such an approach, where the results of multiple algebraic filters are combined. We consider a filter sampling where nine filters are computed instead of a single one, each covering a tile in the reconstructed image. Two different ways of combining these nine AF-FBP reconstructions are experimentally examined. We compare the reconstruction accuracy with the standard AF-FBP method, with the algebraic SIRT method, and with FBP with the standard Ram-Lak filter.

This chapter is structured as follows. In Section 4.2 we briefly discuss FBP and review the standard AF-FBP method. In Section 4.3 we describe a variation on the standard AF-FBP. The experiments that we will perform are described in Section 4.4. In Section 4.5 the results of these experiments are presented. Section 4.6 contains the discussion, and the conclusions are drawn in Section 4.7.

4.2 INTRODUCTION TO AF-FBP

In this section we briefly introduce the key concepts of the AF-FBP method. We refer to [13] for a more detailed description. We consider here the standard case of reconstructing a 2D image from its parallel beam projections. The unknown image can be considered as a map $\mathbb{R}^2 \rightarrow \mathbb{R}$, which assigns a *grey level* to each image coordinate (x, y) . For $\theta \in [0, \pi)$ and $t \in \mathbb{R}$, the *line projection* $p(\theta, t)$ of f is defined by

$$p(\theta, t) = \iint_{\mathbb{R}^2} f(x, y) \delta(x \cos \theta + y \sin \theta - t) dx dy. \quad (4.1)$$

In practice, line projections are measured for a finite set of projection angles θ and at a finite set of detector positions t . Denote the finite set of angles for which projections are available by $\Theta = \{\theta_1, \dots, \theta_k\}$ and the finite set of detector positions for which a projected line has been measured by $T = \{t_1, \dots, t_\ell\}$. The *tomographic reconstruction problem* then consists of recovering f from its line projections for all $(\theta, t) \in \Theta \times T$. Typically, only an approximate solution can be obtained in practice.

We recall that the Filtered Backprojection algorithm can be described by

$$\tilde{f}(x, y) = \sum_{\theta \in \Theta} \sum_{\tau \in T} p(\theta, \tau) g(\theta, \tau - x \cos \theta - y \sin \theta), \quad (4.2)$$

where $\tilde{f}(x, y)$ denotes the reconstructed image and $g(\theta, t)$ is a pre-defined *filter*, such as the common Ram-Lak filter. Usually this filter does not depend on the projection angle θ .

In *algebraic* reconstruction algorithms, the tomographic reconstruction problem is represented by a system of linear equations of the form

$$\mathbf{W}\mathbf{x} = \mathbf{p},$$

were the matrix $\mathbf{W} = (w_{ij}) \in \mathbb{R}^{m \times n}$ denotes the discretized projection operator, the vector $\mathbf{x} \in \mathbb{R}^n$ corresponds to the unknown image and the vector $\mathbf{p} \in \mathbb{R}^m$ contains the measured projection data. The entry w_{ij} determines the weight of the contribution of pixel j to measurement i , which usually represents the length of the intersection bet-

ween the pixel and the projected line. As the equation system is usually inconsistent, an exact solution typically does not exist. For such cases, a range of numerical algorithms exist that minimize the residual $\|\mathbf{W}\mathbf{x} - \mathbf{p}\|$ with respect to some vector norm.

A reconstruction algorithm can be considered as an operator $\mathcal{S} : \mathbb{R}^m \rightarrow \mathbb{R}^n$, which maps the vector \mathbf{p} of measurements to the reconstructed image \mathbf{x} . We call a reconstruction algorithm *linear* if \mathcal{S} is a linear operator. In such cases, the algorithm can be described by a matrix $\mathbf{S} \in \mathbb{R}^{n \times m}$. Several algebraic reconstruction algorithms, including the well-known ART, SART and SIRT algorithms [1], are linear methods.

We now focus on a single pixel $c \in \{1, \dots, n\}$ of the reconstructed image. It has coordinates $(x_c, y_c) \in \mathbb{R}^2$. Denote the c th row of \mathbf{S} by $\mathbf{s}^{(c)}$. Each entry of $\mathbf{s}^{(c)}$ corresponds to an entry in the vector \mathbf{p} of measured projection data, and therefore to a pair $(\theta, \tau) \in \Theta \times \mathcal{T}$, which we denote by $s_{\theta\tau}^{(c)}$. Let $\mathbf{u} = \mathbf{S}\mathbf{p}$. We then have

$$u_c = \sum_{\theta \in \Theta} \sum_{\tau \in \mathcal{T}} p_{\theta\tau} s_{\theta\tau}^{(c)}, \quad (4.3)$$

which is very similar to the filtered backprojection formula of Eq. (4.2). The vector $\mathbf{s}^{(c)}$ can be computed by computing a series of reconstructions using a basis of unit vectors as input. For a given $\theta \in \Theta$, define $t_c^{(\theta)} = x_c \cos \theta + y_c \sin \theta$. Now define the function $h^{(c)}(\theta, \cdot)$ by the following formula:

$$h^{(c)}(\theta, \tau) = s_{\theta(\tau + t_c^{(\theta)})}^{(c)}, \quad \text{for } \tau \in \mathcal{T} - t_c^{(\theta)}. \quad (4.4)$$

We refer to $h^{(c)}$ as the *algebraic filter* for the algorithm \mathcal{S} . Note that $h^{(c)}(\theta, \cdot)$ is only defined on a discrete domain. Intermediate values can be obtained by interpolation. Also note that the value $h^{(c)}(\theta, \tau)$ can depend on the projection angle θ . Just as in FBP, the algebraic filter determines the weight of the contribution of each detector position to the reconstructed value. For the selected pixel c , using the algebraic filter in FBP will result in an identical reconstructed value as for the linear algebraic method \mathcal{S} . It was demonstrated in [13] that if this same filter is also used to reconstruct all other image pixels, using the FBP

formula, an image is obtained that approximates the result of the linear algebraic reconstruction method very well.

4.3 AF-FBP WITH MULTIPLE FILTERS

In this section we describe a modified AF-FBP method where reconstructions of several filters are combined into one final reconstruction. As described in Section 4.2, we can create a filter for FBP based on any linear algebraic reconstruction method. This filter is created by computing a series of reconstructions (one for each detector element in the set T), and storing the value of a single image pixel. The resulting filter depends on the choice of the image pixel. For the standard AF-FBP method the pixel in the center of the image grid is chosen (referred to as the central image pixel at location $(x_c, y_c) \in \mathbb{R}^2$) to compute the filter. Afterwards, this filter can be applied to any vector \mathbf{p} of projection data, resulting in a reconstruction for which the central pixel has the same value as for the algebraic method.

In algebraic methods, the positioning of the reconstruction grid imposes implicit constraints on the reconstruction problem, namely that the reconstructed image is zero outside the region of the reconstruction grid. These constraints mainly affect the reconstruction near the boundary of the reconstruction grid. In general, these implicit constraints are beneficial to the quality of reconstructions computed by algebraic methods, as they limit the support of the reconstructed images to a subregion of \mathbb{R}^2 . As a consequence, an object that is positioned in a corner of the reconstruction grid is not reconstructed identically (up to a shift) to an object that is positioned in the center of the grid. It is therefore natural to consider the question whether an algebraic filter computed for a pixel near a corner of the reconstruction grid is substantially different from the central filter, and whether it can provide more accurate reconstructions in this region. At the same time, using multiple filters for the reconstruction increases the implementation complexity of the reconstruction algorithm compared to FBP, as well as the computation time.

For the sake of simplicity, we limit our investigation here to the case where nine filters are used, obtained by dividing the image grid symmetrically in three columns and three rows. The central image pixel is also the center of one of these subregions. Since the size of

the image grid need not be divisible by three, the subregions need not be of equal size. The size of the subregion in the center of the image equals the image size divided by three. If this is not an integer, we use the smallest integer that is larger than this number. This integer is denoted by s . The pixels that are selected are located at $(x_c \pm s, y_c \pm s)$, $(x_c, y_c \pm s)$, $(x_c \pm s, y_c)$, (x_c, y_c) . For each pixel a filter is created as described in Section 4.2.

According to Eq. (4.4), computing the filter for an image pixel $j \in \{1, \dots, n\}$ with coordinates (x_j, y_j) requires the computation of the row $s^{(j)}$ of the matrix \mathbf{S} . Subsequently, for each angle θ , the part of $s^{(j)}$ that corresponds with this angle is shifted over a distance $t_j^{(\theta)} = x_j \cos \theta + y_j \sin \theta$. Zero padding is applied for the filter values that fall outside the region where this shifted filter can be defined by interpolation.

After computation of the nine filters, each of these filters can be used in a standard FBP implementation. In this way we obtain nine reconstructions of the original image. The nine reconstructions can be combined in several ways to obtain a final reconstruction of the original image. In this chapter we explore two different ways to combine the reconstructions and we compare the reconstruction accuracy of both methods with the linear algebraic method, standard AF-FBP and FBP with the standard Ram-Lak filter.

4.4 EXPERIMENTS

In this section we describe a series of simulation experiments that we performed to examine the accuracy of the modified AF-FBP method.

4.4.1 *Phantoms*

The phantoms that are used in the experiments are shown in Fig. 4.1. Fig. 4.1a is a cross-section of a cylinder head in a combustion engine, Fig. 4.1b represents a cross-section of a wind turbine blade, Fig. 4.1c is the Shepp-Logan phantom which represents a simplification of a cross-section of a human brain [14], and Fig. 4.1d is a cross-section of a mandible.

The phantoms are defined on a square grid of 2044×2044 pixels and are contained in the inscribed circle with diameter 2044 pixels. Real world objects can in general not be represented exactly on such a




Figure 4.1: The phantoms.

pixel grid. To take this into account in our experiments, we reconstruct the phantoms on a coarser grid of 511×511 pixels. To determine the reconstruction error we then refine the reconstruction to 2044×2044 pixels by replacing every pixel in the coarse grid by a block of 4×4 pixels with the same value.

As shown by experiments in [13], the quality of the reconstructions improves when the filters are created on a larger grid of 767×767 pixels instead of 511×511 pixels. Therefore, we use this larger grid to create the filters and to create the algebraic reconstructions. With error measures described in Section 4.4.4, we examine the reconstruction quality of the modified AF-FBP method.

4.4.2 Projection data

The projection data are obtained using a parallel beam scanning geometry. The detector consists of 511 elements, each having a width of four image pixels. The Joseph's model is used to obtain the contribution of each image pixel to an individual projection ray [15]. Per detector element we use four rays, each having a width of one image pixel. This ensures that every image pixel contributes with nonzero weight to at least one ray per projection angle. The projection angles are equiangularly distributed within the full range of 180 degrees. The number of projection angles d varies during the experiments.

4.4.3 Reconstruction methods

In the experiments we use four different reconstruction methods to obtain the reconstructions. All algebraic filters are based on the iterative SIRT algorithm. In particular, the following algorithms are compared:

SIRT Simultaneous Iterative Reconstruction Method. The k th iteration $\mathbf{u}^{(k)}$ is given by Eq. (4.5), where $\mathbf{u}^{(0)} = \mathbf{0}$.

$$\mathbf{u}^{(k+1)} = (\mathbf{I}_n - \omega \mathbf{C} \mathbf{W}^T \mathbf{R} \mathbf{W}) \mathbf{u}^{(k)} + \omega \mathbf{C} \mathbf{W}^T \mathbf{R} \mathbf{p}, \quad (4.5)$$

where $\mathbf{C} = (c_{ij}) \in \mathbb{R}^{n \times n}$ and $\mathbf{R} = (r_{ij}) \in \mathbb{R}^{m \times m}$ are diagonal matrices given by $c_{jj} = \alpha(\sum_{i=1}^n |w_{ij}|)$ and $r_{ii} = \beta(\sum_{j=1}^m |w_{ij}|)$ respectively, $\mathbf{I}_n \in \mathbb{R}^{n \times n}$ denotes the identity matrix, and ω is a relaxation parameter [8, 16]. In this chapter we set $\omega = 1$, $\alpha(x) = \beta(x) = 1/x$ for $x \in \mathbb{R}$. We stop the SIRT algorithm after a fixed number of 200 iterations.

FBP-RL Filtered Backprojection method given by Eq. (4.2) using the standard Ram-Lak filter.

SIRT-FBP SF SIRT-FBP ‘Single Filter’; this is the standard SIRT-FBP method using one angle-dependent filter based on the central pixel of the reconstruction grid; see Section 4.2.

SIRT-FBP IR SIRT-FBP ‘Isolated Regions’; this is a variation on SIRT-FBP using nine pixels to calculate nine angle-dependent filters. The selection of these nine pixels is described in Section 4.3. For each filter a reconstruction is created using FBP with this filter. The subregion containing the pixel that was used to create the filter is used for the final reconstruction. These subregions do not overlap and will be referred to as isolated regions. The collection of all these isolated subregions forms a complete reconstruction.

SIRT-FBP BI SIRT-FBP ‘Bilinear Interpolation’; for this method we use the same subregions as for IR, but we perform an additional bilinear interpolation step instead of only collecting the isolated subregions.

4.4.4 Error measure

The accuracy of the reconstruction methods is examined using two different error measures. The first measure is obtained by calculating the sum of absolute pixel differences between the reconstruction and the original object. After normalization, the resulting error is denoted as the *mean reconstruction error* E_r . Since the reconstruction has a coarser grid than the original image, we cannot compare them directly. We first select in the 767×767 reconstruction grid the smaller 511×511 pixel grid with the same central pixel as the reconstruction grid. We then enlarge this grid by replacing every pixel by a block of 4×4 pixels with the same value to obtain a grid of 2044×2044 pixels. We denote this image by $\hat{\mathbf{u}} = (\hat{u}_{ij}) \in \mathbb{R}^{n^2}$ with $n = 2044$ and the original phantom by $\mathbf{v} = (v_{ij}) \in \mathbb{R}^{n^2}$. The mean reconstruction error is then defined by

$$E_r = \frac{\sum_{i,j} |\hat{u}_{ij} - v_{ij}|}{\sum_{i,j} v_{ij}}. \quad (4.6)$$

The second measure is obtained by calculating the sum of absolute pixel differences between the reconstruction $\mathbf{u} = (u_{kl})$ and the algebraic reconstruction $\mathbf{u}^{\text{ARM}} = (u_{kl}^{\text{ARM}})$ on the same grid of 511×511 pixels that was used in the mean reconstruction error before refining it. After normalization this error is denoted as the *mean ARM reconstruction error* E_r^{ARM} , defined by

$$E_r^{\text{ARM}} = \frac{\sum_{k,l} |u_{kl} - u_{kl}^{\text{ARM}}|}{\sum_{k,l} u_{kl}^{\text{ARM}}}. \quad (4.7)$$

4.4.5 Series of experiments

The reconstructions are compared in three series of experiments.

In the first series of experiments the number of projection angles is varied. We show the resulting mean reconstruction errors E_r , as well as the mean ARM reconstruction errors E_r^{ARM} , for all four phantoms and using several different interpolation methods for the filters. The

aim of this series is to compare the accuracy of the five reconstruction methods for a varying number of projections ($d \in \{16, 24, 32, 64\}$).

In the second series of experiments we apply varying amounts of Poisson noise to the projection data. FBP with a Ram-Lak filter is known to produce poor quality reconstructions if the projection data has a low signal-to-noise ratio, while SIRT provides superior results for such data. We examine the reconstruction quality of SIRT-FBP with multiple filters, compared to SIRT, FBP-RL and standard SIRT-FBP.

In the third series of experiments a larger detector is used to obtain the projections and to create the filters. As discussed in Section 4.3, a shift operation has to be applied when computing an algebraic filter for a non-central image pixel. If such an image pixel is near the boundary of the reconstruction region, then substantial zero-padding is needed for some angles on one side of the shifted filter. Enlarging the size of the detector (and hence the size of the filter), decreases the number of zeros that has to be inserted. A disadvantage is the larger computation time to create the filter. In this series of experiments we examine the effect of using these larger filters on the reconstruction quality.

4.5 RESULTS

In this section we present the results of the experiments described in the previous section. We start with discussing the algebraic filters for the different selected image pixels and show their effect on the reconstructions. Then we compare the reconstructions of SIRT-FBP with multiple filters with reconstructions of both the standard SIRT-FBP and FBP-RL.

4.5.1 Filter dependence on the selected pixel

The algebraic filters depend quite strongly on the pixel that is used to create them. Each image in Fig. 4.2 provides a visual representation of a particular row $s^{(j)}$ of the matrix \mathbf{S} , where the vertical axis corresponds with the projection angle and the horizontal axis corresponds with the position on the detector. The selected pixels p_1, \dots, p_9 are numbered row-wise, starting at the top left corner with p_1 .




Figure 4.2: Visual representations of rows of the matrix S for $d = 32$ projection angles.




Figure 4.3: Filters for $d = 32$ projection angles after linear interpolation and shifting.

Fig. 4.3 shows the corresponding filters that are the result of Eq. 4.4, after applying a linear interpolation and angle-dependent shifting of

the images in Fig. 4.2. The large homogeneous grey areas arise due to the necessary zero-padding to obtain a filter of size 511 after shifting. These areas might have a large influence on the resulting reconstructions. To examine this influence, we also include experiments for a larger filter of 767 elements. Since this gives rise to a larger system of equations $\mathbf{W}\mathbf{x} = \mathbf{p}$, the algebraic reconstructions will change as well. The corresponding filters will therefore not only contain fewer zero elements, but are different from the filters in Fig. 4.3 in the nonzero entries as well. The results are shown in Section 4.5.4.


In Fig. 4.4 the reconstructions are shown for the nine different angle-dependent filters with $d = 32$. These reconstructions are combined into a final reconstruction using either the filters corresponding to the central pixel (denoted by SF), or to the collection of isolated subregions (denoted by IR), or to a bilinear interpolation of the isolated regions (denoted by BI), as discussed in Section 4.4.3.

4.5.2 *Varying the number of projection angles*


In the first series of experiments the four phantoms are reconstructed using different numbers of projection angles. The projection data are noiseless in these experiments. Fig. 4.6 shows the mean reconstruction errors for the turbine blade phantom.

The mean reconstruction error of both standard SIRT-FBP and modified SIRT-FBP is very close to the error of SIRT. The variation SIRT-FBP with bilinear interpolation outperforms SIRT-FBP with isolated regions. This is expected behavior, since bilinear interpolation uses the values of a pixel in two or more reconstructions, which results in a smoothing of the presumed larger pixel errors near the boundary of the isolated regions. Therefore, we expect the mean reconstruction error, E_r , to be smaller than that of SIRT-FBP with isolated regions. Furthermore, we observe that standard SIRT-FBP results in reconstructions that are slightly more accurate than SIRT or SIRT-FBP with multiple filters. FBP-RL is substantially less accurate than the other reconstruction methods considered here.


The standard SIRT-FBP method uses only a single filter based on the central pixel. These results therefore show that the local use of filters based on pixels other than the central pixel does not reduce the mean (ARM) reconstruction error. This might be due to the fact that



(a) FBP with filter for p_1 (b) FBP with filter for p_2 (c) FBP with filter for p_3



(d) FBP with filter for p_4 (e) FBP with filter for p_5 (f) FBP with filter for p_6



(g) FBP with filter for p_7 (h) FBP with filter for p_8 (i) FBP with filter for p_9

Figure 4.4: Reconstructions of the cylinder head phantom using the filters based on the nine different image pixels.

shifting and zero-padding of these filters is required before they can be applied to the projection data, resulting in the loss of some accuracy in the approximation of the ARM due to the additional interpolation step.

SIRT-FBP with bilinear interpolation is not only more accurate than SIRT-FBP with isolated regions, it is also a better approximation of




Figure 4.5: Reconstructions of the cylinder head phantom with $d = 32$; (a) SIRT reconstruction using 200 iterations; (b) FBP with Ram-Lak filter; (c) SIRT-FBP with the filters corresponding to the central pixel; (d) SIRT-FBP combining the nine reconstructions from isolated regions; (e) SIRT-FBP combining the nine reconstructions using bilinear interpolation.

SIRT. Its mean ARM reconstruction error is only slightly larger than that of standard SIRT-FBP. The mean ARM reconstruction error of SIRT-FBP with isolated regions is larger than that of standard SIRT-FBP and SIRT-FBP with bilinear interpolation. However, it is still much smaller than that of FBP-RL.

Furthermore, we observe that the reconstruction errors are almost indifferent to the method chosen to interpolate the filter. Every row in Fig. 4.6 shows the reconstruction error for a different interpolation method. We will therefore not include the results of the different interpolation methods in the remainder of this chapter and only show the results using the linear interpolation method for the filters. These results for the other three phantoms are shown in Fig. 4.7. The left-hand side column shows the same pattern as for the turbine blade.




Figure 4.6: Mean (ARM) reconstruction error E_T (E_T^{ARM}) for the turbine blade phantom. Several interpolation methods are used to compute the filters; SF = single filter, IR = isolated regions, BI = bilinear interpolation.




Figure 4.7: Mean reconstruction error E_T of the cylinder head phantom (top row), mandible phantom (middle row), Shepp Logan phantom (bottom row). Linear interpolation is applied to compute the filters. SF = single filter, IR = isolated regions, BI = bilinear interpolation.

The mean ARM reconstruction errors for SIRT-FBP with bilinear interpolation are again smaller than for SIRT-FBP with isolated regions. In all cases, the reconstructions computed using algebraic filters are more accurate than for FBP-RL.

4.5.3 Poisson noise

In the second series of experiments we apply Poisson noise to the projection data, by simulating the Poisson distribution based on a varying number of photons used for the beam profile. The lower this photon count, the higher the noise level. The results for varying noise levels are shown in Fig. 4.8 and Fig. 4.9.

It can be observed in Fig. 4.8 that FBP-RL cannot handle noisy projection data very well. The mean reconstruction errors for high amounts of noise are large compared to those for SIRT-FBP and SIRT. The right-hand side column of both figures is a detail of the figures on the left-hand side. In this detail we observe that SIRT-FBP performs slightly better in terms of the mean reconstruction error than SIRT. This is consistent with observations in [13] for $d = 64$. It can also be observed for the other numbers of projection angles. Furthermore, assembling the modified SIRT-FBP reconstruction using isolated regions or bilinear interpolation reduces E_r for high amounts of noise. The local reconstructions using filters corresponding to the selected pixels p_1, p_3, p_7 and p_9 contain less noisy reconstructions in the area around the selected pixel. Since only this area is used while assembling the final reconstruction, this might explain the smaller E_r . For less noisy projection data, the standard SIRT-FBP method based on one angle-dependent filter performs equally well or better than the modified SIRT-FBP method, depending on the number of projection angles.

In Fig. 4.9 the results are shown for $d = 32$ and $d = 64$. The mean ARM reconstruction errors E_r^{ARM} with $d = 16$ and $d = 24$ are similar to those for $d = 32$ and therefore not included in this chapter. The reconstructions of the standard SIRT-FBP (with one angle-dependent filter) differ less from SIRT than SIRT-FBP with multiple angle-dependent filters. We do observe that bilinear interpolation results in reconstructions that are closer to the SIRT reconstructions than using isolated regions.

All results for the cylinder head are representative for the other three phantoms. The results for the other phantoms are therefore not included in this chapter.




Figure 4.8: Mean reconstruction error E_T of the cylinder head phantom for varying amounts of noise applied to the projection data. Linear interpolation is applied to compute the filters. SF = single filter, IR = isolated regions, BI = bilinear interpolation.




Figure 4.9: Mean ARM reconstruction error E_{τ}^{ARM} of the cylinder head phantom for varying amounts of noise applied to the projection data. The rows correspond to different numbers of projection angles d . Linear interpolation is applied to compute the filters. SF = single filter, IR = isolated regions, BI = bilinear interpolation.

4.5.4 Varying the size of the detector

In the third series of experiments we increase the size of the detector from 511 elements to 767 elements. The filters that cover the reconstruction grid of 511×511 pixels need less zero-padding. Several filters are shown in Fig. 4.10.

The reconstruction errors using these larger filters are shown in Fig. 4.11. We compare them with the results for the smaller filters in Fig. 4.6. We observe that both E_{τ} and E_{τ}^{ARM} of SIRT-FBP are less affected by the method chosen to assemble the reconstruction than for the smaller filters with a support of 511 elements. The value of E_{τ} is comparable in both cases. The right-hand side column of the figure shows that the reconstructions using a detector of 767 elements are more similar to

(a) Filter for pixel p_1 (b) Filter for pixel p_5 (c) Filter for pixel p_6 (d) Filter for pixel p_8

Figure 4.10: Larger filters with 767 elements per angle for $d = 32$ projection angles.

the SIRT reconstructions than those using a smaller detector of 511 elements, although the differences are small.

The results for different interpolation methods used in the computation of the filter and for different phantoms are similar to those discussed above and are omitted here.

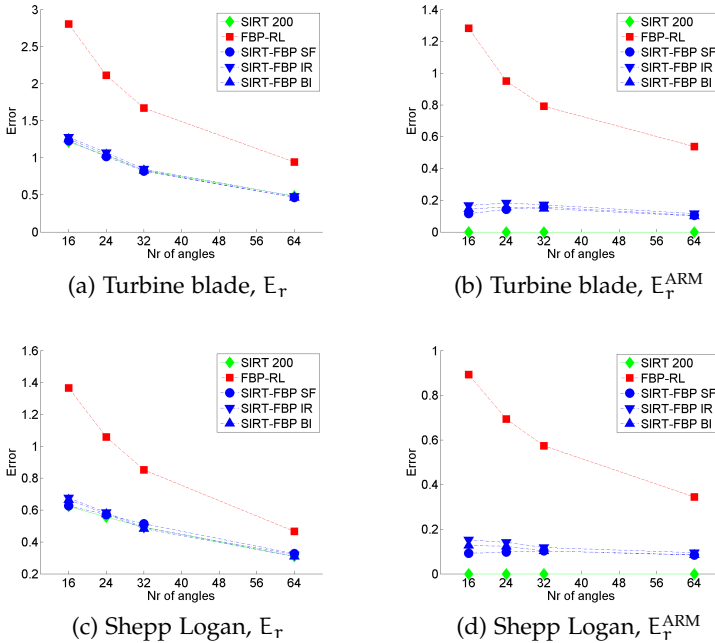


Figure 4.11: Mean (ARM) reconstruction error E_T (E_T^{ARM}) for filters with 767 elements per angle, for varying number of projection angles. Linear interpolation is applied to compute the filters. SF = single filter, IR = isolated regions, BI = bilinear interpolation.

4.6 DISCUSSION

Our experimental results suggest that the standard SIRT-FBP method approximates SIRT better than SIRT-FBP with nine algebraic filters. An interesting result is obtained in the case where we apply high amounts of Poisson noise to the projection data and compare the reconstructions with the phantom instead of with the SIRT reconstruction. In this case, the reconstructions of SIRT-FBP with bilinear interpolation are more accurate than the standard SIRT-FBP reconstructions and the FBP-RL reconstructions. Although it is not the aim of the AF-FBP method to improve the accuracy of the reconstructions, since there we want to approximate the ARM as best as possible, it opens new opportunities for further research.

The value of the pixel that was used to create the filter for SIRT-FBP is the same in both the SIRT reconstruction and the SIRT-FBP reconstruction. We can therefore view SIRT as an SIRT-FBP with an algebraic filter calculated for every image pixel. Hence we have compared the reconstruction accuracy of SIRT-FBP with one, with nine and with 511×511 algebraic filters. The results show that for high amounts of Poisson noise the reconstruction accuracy of SIRT-FBP with nine algebraic filters outperforms that of SIRT-FBP with both one algebraic filter and with 511×511 algebraic filters. Assuming that the reconstruction accuracy of SIRT-FBP depends smoothly on the number of algebraic filters, there should be an optimal number of algebraic filters that maximizes the reconstruction accuracy of SIRT-FBP. Also the position of the selected pixels in the image grid can influence this accuracy and should be examined in more detail.

4.7 CONCLUSIONS

In this chapter we have investigated an approach to create algebraic filters (AFs) that can be used in the Filtered Backprojection method. Reconstructions of the standard AF-FBP method can be computed using the same computation time as FBP, while they resemble the reconstruction properties of the linear algebraic reconstruction method that was used to create the filter. Compared to the original AF-FBP method, which is based on a single filter, the approach proposed in this chapter is based on computing multiple filters, each covering a region of the reconstruction grid. For each (angle-dependent) filter we calculate an FBP reconstruction. Parts of these reconstructions are combined into a final reconstruction by either assembling separate regions or by using bilinear interpolation for neighboring regions from each reconstruction. In a series of simulation experiments, we examined the reconstruction accuracy of these variations by comparing the reconstructions with both the phantoms and the corresponding algebraic reconstructions, based on the SIRT algorithm.

For the sake of clarity, we focused on a particular scenario where nine filters are created, each based on a different pixel in the reconstruction grid. The resulting reconstructions show little dependence on the different interpolation methods that can be applied to compute the filters. Since in the creation of some filters substantial zero-padding

is required, we also created filters with an extended support to reduce this zero-padding. This results in only a small effect on the quality of the corresponding reconstructions, while it increases the computation time of the filters.

Assembling the final reconstruction using bilinear interpolation results in general in a more accurate reconstruction than combining separate regions. The differences between standard AF-FBP and AF-FBP based on multiple filters are small and in many situations the standard AF-FBP method outperforms the suggested variations. There are situations, such as projection data with few angles and a high amount of noise, where modified AF-FBP using bilinear interpolation outperforms standard AF-FBP when we compare the reconstructions with the original phantoms. In general however we observe that using more than one angle-dependent filter does not automatically lead to a better approximation of the corresponding algebraic method, or in a more accurate reconstruction method.

Our findings suggest that for the AF-FBP method, the central pixel is actually a very good choice for computing the filter, even if this filter is then applied to reconstruct image pixels in the outer regions of the reconstruction grid.

BIBLIOGRAPHY

- [1] A. C. Kak and M. Slaney. *Principles of Computerized Tomographic Imaging*. Philadelphia: SIAM, 2001.
- [2] G. T. Herman. *Fundamentals of Computerized Tomography: Image Reconstruction from Projections*. Berlin: Springer, 2009.
- [3] T. M. Buzug. *Computed Tomography: From Photon Statistics to Modern Cone-Beam CT*. Berlin: Springer, 2008.
- [4] A. Markoe. *Analytic Tomography*. New York: Cambridge University Press, 2006.
- [5] X. Pan, E. Y. Sidky, and M. Vannier. Why do commercial CT scanners still employ traditional, filtered back-projection for image reconstruction? *Inv. Problems* 2009; 25(12): 123009.
- [6] L. A. Feldkamp, L. C. Davis, and J. W. Kress. Practical cone-beam algorithm. *J. Opt. Soc. Am.* 1984; 1(A6): 612–619.
- [7] A. H. Andersen and A. C. Kak. Simultaneous Algebraic Reconstruction Technique (SART): a superior implementation of the ART algorithm. *Ultrasound. Imag.* 1984; 6(1): 81–94.
- [8] J. Gregor and T. Benson. Computational analysis and improvement of SIRT. *IEEE Trans. Med. Imag.* 2008; 27(7): 918–924.
- [9] A. K. Louis and Th. Schuster. A Novel Filter Design Technique in 2D - Computerized Tomography. *Inverse Problems* 1995.
- [10] H. Kunze et al. Filter determination for Tomosynthesis aided by iterative reconstruction techniques. *Proc. of Fully3D* 2007: 309–312.
- [11] J. Ludwig et al. A Novel Approach for Filtered Backprojection in Tomosynthesis Based on Filter Kernels Determined by Iterative Reconstruction Techniques. *Proc. of IWDM* 2008; 5116: 612–620.
- [12] J. Hsieh Y. Wei G. Wang. An intuitive discussion on the ideal ramp filter in computed tomography (I). *J. Comput. Math. Appl.* 2005; 49(5–6): 731–740.
- [13] K. J. Batenburg and L. Plantagie. Fast Approximation of Algebraic Reconstruction Methods for Tomography. *IEEE Trans. Image Process.* 2012; 21(8): 3648–3658.

- [14] L. A. Shepp and B. F. Logan. The Fourier reconstruction of a head section. *IEEE Trans. Nucl. Science* 1974; 21(3): 21–43.
- [15] P.M. Joseph. An improved algorithm for reprojecting rays through pixel images. *IEEE Trans. Med. Imag.* 1982; 1(3): 192–196.
- [16] P. Gilbert. Iterative methods for the three-dimensional reconstruction of an object from projections. *J. Theor. Biol.* 1972; 36(1): 105–117.