

Cover Page



Universiteit Leiden



The handle <http://hdl.handle.net/1887/39638> holds various files of this Leiden University dissertation.

Author: Pelt D.M.

Title: Filter-based reconstruction methods for tomography

Issue Date: 2016-05-03

2

Data-dependent filtering

2.1 Introduction

Tomographic reconstruction problems are found in many applications, such as X-ray scanners in medical imaging, or electron microscopy in materials science [Gra13]. In the standard tomographic problem, we aim to reconstruct an object from its projections, acquired for a range of angles. This problem has been studied extensively because of its practical relevance, leading to a wide range of reconstruction methods. For an overview of previous work, see [KS01; Nat01; Buz08]. Most of the current reconstruction methods can be separated into two groups: *analytical* methods and *algebraic* methods.

The basis of analytical reconstruction methods is a continuous representation of the tomographic problem. This continuous model is inverted, and the result is discretized. The resulting reconstruction methods, of which the filtered backprojection (FBP) method is the most widely used, are usually computationally efficient. Furthermore, if projection data of sufficiently high quality is available, reconstructions computed by these methods are often accurate. These two properties are among the reasons that the FBP method is very popular in practice [PSV09], along with its ease of implementation. An important drawback of analytical methods is that they are based on an approximation of a model where perfect data is available for *all* projection angles. If the available data is not perfect, either because few projections are available or because the data is noisy, the quality of analytical reconstructions will suffer from interpolation effects.

Practical considerations can lead to limited or noisy projection data in many applications of tomography. In electron tomography, for example, the electron beam

This chapter is based on:

D. M. Pelt and K. J. Batenburg. “Improving Filtered Backprojection Reconstruction by Data-Dependent Filtering”. *IEEE Transactions on Image Processing* 23.11 (2014), pp. 4750–4762.

damages the sample, leading to a hard limit on the number of projections that can be measured [MDG95]. In many other applications, there is a limit on the duration of a single scan. To decrease the scan duration, one can either acquire fewer projections or use a reduced dose per projection. In industrial tomography, process speed considerations limit the duration of each scan [Sip93].

Algebraic methods are based on a discrete representation of the tomographic problem, leading to a linear system of equations. This system is solved to obtain a reconstructed image. Since algebraic methods use a model of the actual data that is available, they usually yield more accurate reconstructions from limited data than analytical methods. Furthermore, by using specific ways of solving the linear system, it is possible to reduce the effect of noise on the reconstruction. An important drawback of algebraic methods is that they are computationally more expensive than analytical methods. The linear system that has to be solved is usually very large, and the iterative methods that are used often need a large number of iterations to converge to an acceptable solution.

In many applications of computed tomography, the computational efficiency of a reconstruction method is an important consideration. For example, in fast x-ray microtomographic experiments at synchrotrons, the speed of the post-processing pipeline has to match the high speed of data acquisition [Mok+13]. In fact, the computation efficiency of the FBP method is an important reason for why it is still commonly used instead of more advanced reconstruction methods [PSV09].

Methods that reduce the computation time of algebraic methods have been proposed by other authors. One approach is to implement algebraic methods more efficiently by using graphic processing units (GPUs) [XM05; Pan+11]. Other approaches focused on improving the convergence of algebraic methods, for example by improving the properties of the linear system [GB08]. Although these improvements reduce the computation time of algebraic methods significantly, even faster methods can be obtained by changing the algebraic methods themselves.

One such approach is taken in [BP12], where an angle-dependent FBP filter is calculated, such that the resulting FBP method approximates an algebraic method. Although the resulting method is able to approximate the algebraic method well, calculating the filter requires a large number of runs of the algebraic method, which is computationally expensive. The resulting filter can be reused for problems with identical projection geometry, but a change in geometry requires calculation of a new filter.

A filter that approximates an algebraic method is also derived in [Zen12], in which a reformulation of the SIRT algebraic method is translated to a fixed filter for the FBP method. An extension of the method for noisy projection data is given in [ZZ13]. The derived filter does not depend on the scanning geometry of the problem, and during derivation it is assumed that enough projections are available such that certain approximations are accurate. As such, the resulting method has more in common with analytical reconstruction methods than with algebraic methods.

A different approach, specific to tomosynthesis, is proposed in [Nie+12]. Instead of calculating a reconstruction image directly, Nielsen et al. calculate a filter *matrix*, which is multiplied with the projection data. The result is backprojected to produce

the final reconstruction. Nielsen et al. show that their filter matrix can be formed efficiently in the case of tomosynthesis, but a complex method is needed to obtain this efficiency. Similar to [BP12], a change in geometry requires calculation of a new filter. Other methods for tomosynthesis use algebraic reconstructions of certain test objects to create filters [Kun+07; Lud+08].

In this chapter, we introduce a new reconstruction method, the *minimum residual filtered backprojection* method (MR-FBP), that combines ideas from both the analytical and algebraic approach, resulting in a method with a data-dependent filter. The method is based on an algebraic model of the tomographic problem, resulting in a method that can reconstruct problems with limited data more accurately than analytical methods. The linear system that has to be minimized, however, is based on filtered backprojection. Therefore, the system is much smaller than the ones used in algebraic methods or other approaches, making the method computationally efficient. Furthermore, we are able to use filtered backprojection to form our linear system, leading to a simple and efficient implementation.

This chapter is structured as follows. We formally define the tomographic reconstruction problem and analytical and algebraic reconstruction methods in Section 2.2. In Section 2.3, we introduce and explain the key contribution of this chapter: the minimum residual filtered backprojection method. Considerations concerning its implementation are discussed in Section 2.4. An extension of the method is given in Section 2.5, where additional constraints are added to its linear system to improve reconstruction quality. The experiments we performed to test the new method are explained in Section 2.6. Results, where we compare MR-FBP with popular reconstruction methods, are given in Section 2.7, along with a discussion on the interpretation of the results. Finally, we conclude the chapter in Section 2.8, where we give a summary and some final remarks.

2.2 Notation and concepts

In this section, we will explain the mathematical notation used throughout the chapter, and introduce all relevant concepts. We begin by formally defining the tomographic reconstruction problem. Filtered backprojection and algebraic methods are explained, and their mathematical definitions are given.

Problem definition

We consider the problem of reconstructing a two-dimensional object from its parallel-beam projections, with a single rotation axis. The approach we introduce here can be adapted to other geometries as well, such as fan-beam or cone-beam projections. The unknown object is modeled as a finite and integrable two-dimensional function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ with bounded support.

Define a line $l_{\theta,t}$ by its characteristic equation $t = x \cos \theta + y \sin \theta$. The line integral




Figure 2.1: The two-dimensional tomography model used in this chapter. Parallel lines, rotated by angle θ , pass through the object f . A line $l_{\theta, t}$ has the characteristic equation $t = x \cos \theta + y \sin \theta$, and a projection $P_{\theta}(t)$ of f is given by the line integral of f over the line $l_{\theta, t}$.

$P_{\theta}(t)$ of f over a single line $l_{\theta, t}$ is given by:

$$P_{\theta}(t) = \int_{l_{\theta, t}} f(x, y) ds \quad (2.1)$$

$$= \iint_{\mathbb{R}^2} f(x, y) \delta(x \cos \theta + y \sin \theta - t) dx dy \quad (2.2)$$

The *tomographic reconstruction problem* is concerned with the reconstruction of the unknown object f from its measured projections $P_{\theta}(t)$ for different combinations of θ and t . This projection geometry is shown graphically in Fig. 2.1.

In practice, only a finite set of projections P_{θ} are measured, one for each combination of projection angle $\theta \in \Theta = \{\theta_0, \theta_1, \dots, \theta_{N_{\theta}-1}\}$ and detector $k \in \{0, 1, \dots, N_d - 1\}$, where N_{θ} is the number of projection angles, and N_d the number of detectors. Relative to the central detector, the position of a detector k is given by τ_k :

$$\tau_k = s \left(k - \frac{N_d - 1}{2} \right), \quad (2.3)$$

where s is the width of a detector. The entire set of measured detector positions is given by $T = \{\tau_0, \tau_1, \dots, \tau_{N_d-1}\}$. Using the measured projection data, the unknown object is reconstructed on an $N \times N$ grid of square pixels. We assume, without loss of generality, that the width of each pixel is equal to 1. Often, the number of pixels in each row of the reconstruction grid is taken equal to the number of detectors.

Filtered backprojection

One approach to solving the tomographic reconstruction problem is to take Eq. (2.2), and try to find an expression for $f(x, y)$ from this equation. The *filtered backprojection* method (FBP) is a result of this approach, and starts with convolving the projections




Figure 2.2: The Ram-Lak (ramp) filter for the FBP method, in real space. This filter is a discrete approximation of the optimal filter.

with a filter $h : \mathbb{R} \rightarrow \mathbb{R}$:

$$q_\theta(t) = \int_{-\infty}^{\infty} h(\tau) P_\theta(t - \tau) d\tau \quad (2.4)$$

This convolution can be also be performed in Fourier space, where $\hat{P}$ and $\hat{h}$ denote the Fourier transforms of P and h respectively:

$$q_\theta(t) = \int_{-\infty}^{\infty} \hat{h}(u) \hat{P}_\theta(u) e^{2\pi i u t} du \quad (2.5)$$

One can show [KS01] that we obtain an expression for $f(x, y)$ if we take $\hat{h}(u) = |u|$:

$$f(x, y) = \int_0^\pi q_\theta(x \cos \theta + y \sin \theta) d\theta \quad (2.6)$$

In practice, it is not possible to use Eq. (2.6) to reconstruct the object, since it requires $P_\theta(t)$ to be known for *all* angles $\theta \in [0, \pi)$ and $t \in \mathbb{R}$. Instead, we only know $P_\theta(t)$ for the measured angles Θ and detector positions T . To be able to use these discrete measurements, Eq. (2.6) has to be discretized, after which the filtered backprojection method is obtained:

$$f(x, y) \approx FBP_h(x, y) = \sum_{\theta_d \in \Theta} \sum_{\tau_p \in T} h(\tau_p) P_{\theta_d}(t - \tau_p) \quad (2.7)$$

where $t = x \cos \theta_d + y \sin \theta_d$. Since $t - \tau_p$ is usually not equal to one of the measured detector positions, some interpolation is needed to find the value of $P_{\theta_d}(t - \tau_p)$. Linear interpolation is often used, since projection data is usually reasonably smooth.

The filter h is only needed for discrete positions $\tau_p \in T$, and is therefore usually specified as a vector $\mathbf{h}$. Several discrete approximations of the optimal filter $\hat{h}(u) = |u|$ are used in practice, such as the Ram-Lak (ramp), Shepp-Logan, and Hann filters [Far+97; WWH05]. One of the most popular filters is the Ram-Lak filter, where we take the optimal $\hat{h}(u) = |u|$, and set $\hat{h}(u) = 0$ when $u > u_c$ for some u_c . This filter is shown, in real space, in Fig. 2.2.

The filtered backprojection method can be interpreted as a two-step process. First, the projection data is filtered by convolving it with filter $\mathbf{h}$. Afterwards, the result is




Figure 2.3: Three reconstructions of the Shepp-Logan head phantom (a), showing artifacts that occur with imperfect data. In (b) and (c), FBP was used to reconstruct the images. In (b), data from only 16 projections were used, and severe streak artifacts are present in the result. In (c), data from 1024 projections were used, but a large amount of Poisson noise was present, resulting in severe noise in the resulting image. In (d), SIRT, an algebraic method, was used to reconstruct the image using the same noisy projection data as in (c). The algebraic reconstruction image (d) has less noise compared to the FBP reconstruction (c).

backprojected to obtain the reconstructed image. One of the advantages of FBP is that it is fast compared to other methods: the filtering step can be performed efficiently in Fourier space in $\mathcal{O}(N_\theta N_d \log N_d)$ time, and only one backprojection is needed, which can be performed in $\mathcal{O}(N_\theta N^2)$ time.

The quality of an FBP reconstruction depends on how well the discretized equation Eq. (2.7) approximates the continuous equation Eq. (2.6). If data for many projection angles (say, several hundred) are known, an FBP reconstruction is often highly accurate. However, when the number of projections is small compared to the size of the image, the approximation is not very accurate, and severe artifacts can appear in the reconstructed image. Furthermore, noise in the projection data can cause artifacts in the reconstruction as well. FBP with the Ram-Lak filter is especially sensitive to noise, since high-frequency components of the projection data are amplified by the filter. The artifacts can make subsequent analysis of the reconstruction very difficult. Examples of artifacts in FBP reconstructions of imperfect data are shown in Fig. 2.3.

Algebraic methods

A different approach to solving the tomographic problem is to use a discrete representation of the problem. Here, we represent the discrete projection data as a single vector $\mathbf{p} \in \mathbb{R}^{N_\theta N_d}$, and represent the unknown image as a vector $\mathbf{x} \in \mathbb{R}^{N^2}$. The *projection matrix* $\mathbf{W}$ has $N_\theta N_d$ rows and N^2 columns, with element w_{ij} specifying the contribution of pixel j to detector i . We refer to the product of $\mathbf{W}$ with an image $\mathbf{x}$ as the *forward projection* of $\mathbf{x}$. Similarly, the product of $\mathbf{W}^T$ with projection data $\mathbf{p}$ is referred to as the *backprojection* of $\mathbf{p}$. If we look at the definition of the discrete FBP method (Eq. (2.7)), we see that the backprojection in the FBP method is identical to multiplication of the filtered sinogram with $\mathbf{W}^T$.

Algebraic methods are usually designed to minimize the difference between the measured projection data $\mathbf{p}$ and the forward projection of the reconstruction image, $\mathbf{W}\mathbf{x}$, with respect to a certain vector norm. In the case of the ℓ_2 -norm, the algebraic

solution image $\mathbf{x}_{alg}$ is defined as:

$$\mathbf{x}_{alg} = \underset{\mathbf{x}}{\operatorname{argmin}} \|\mathbf{p} - \mathbf{W}\mathbf{x}\|_2^2 \quad (2.8)$$

The algebraic solution $\mathbf{x}_{alg}$ can be found by solving the linear system $\mathbf{W}\mathbf{x} = \mathbf{p}$ in a least squares sense.

The algebraic linear system is typically too large to be solved directly. Therefore, an iterative optimization method is normally used, which can often exploit the sparse structure of the projection matrix to improve computational and memory requirements. Different iterative methods can be used, leading to various algebraic reconstruction methods. One example is SIRT [KS01], belonging to the class of Landweber iteration methods [Lan51], which uses a specific Krylov subspace method to minimize the projection error. A different method is CGLS [Bjö96], which uses a conjugate gradient method.

The advantage of using an algebraic method compared to analytical methods is that the projection matrix $\mathbf{W}$ can be adapted to the actual geometry that was used during scanning. Therefore, these methods use a model of the actual data that is available, instead of assuming perfect data, as in analytical methods. Another advantage of algebraic methods is that additional constraints can be imposed on the reconstructed image $\mathbf{x}$, which can be used to improve reconstructions by exploiting prior knowledge. For example, total variation minimization based methods use algebraic methods to minimize the projection error, with an additional constraint that the ℓ_1 -norm of the gradient of $\mathbf{x}$ should be minimal as well [SP08].

The main disadvantage of algebraic methods compared to analytical methods is their computation time. Because of the large system size, and the number of iterations that are needed to solve them, the time to reconstruct an image using an algebraic method is often several orders of magnitude larger than filtered backprojection, even when optimized for graphic processor units (GPUs) [XM05]. In the next section, we introduce a new reconstruction method that uses ideas from algebraic methods to improve filtered backprojection, leading to a method that is both fast and accurate.

2.3 Minimum residual filtered backprojection

We will now present the key contribution of this chapter: the minimum residual filtered backprojection method (MR-FBP). We start by noting that the FBP method is a linear operation on the projection data. In other words, the operation of the FBP algorithm can be modeled as a linear operator $\mathcal{M} : \mathbb{R}^{N_d N_\theta} \rightarrow \mathbb{R}^{N^2}$ applied to the projection data $\mathbf{p}$, which can be written as a $N^2 \times N_d N_\theta$ matrix $\mathbf{M}_h$:

$$\text{FBP}_h(\mathbf{p}) = \mathbf{M}_h \mathbf{p} \quad (2.9)$$

As explained in Section 2.2, FBP consists of a convolution of $\mathbf{p}$ with filter $\mathbf{h}$, followed by a backprojection of the result:

$$\mathbf{M}_h \mathbf{p} = \mathbf{W}^T \mathbf{C}_h \mathbf{p} \quad (2.10)$$

where $\mathbf{C}_h \mathbf{p}$ is the convolution of $\mathbf{p}$ by $\mathbf{h}$, specified by the $N_d N_\theta \times N_d N_\theta$ matrix $\mathbf{C}_h$.

One of the properties of a convolution of two vectors is that it is *commutative*. Therefore, we can exchange the positions of $\mathbf{h}$ and $\mathbf{p}$ in Eq. (2.10):

$$\mathbf{M}_h \mathbf{p} = \mathbf{W}^T \mathbf{C}_p \mathbf{h} \quad (2.11)$$

Up to this point, we have only rewritten the equation of the FBP method. Now, we will improve the method by changing the filter $\mathbf{h}$ from one of the standard filters to one specific to the problem we are solving. To calculate the specific data-dependent filter $\mathbf{h}^*$, we minimize the squared difference of the projections of the reconstruction with the measured projection data, similar to algebraic methods:

$$\mathbf{h}^* = \underset{\mathbf{h}}{\operatorname{argmin}} \|\mathbf{p} - \mathbf{W} \operatorname{FBP}_h(\mathbf{p})\|_2^2 \quad (2.12)$$

Using Eqs. (2.9) and (2.11), we can write this as:

$$\mathbf{h}^* = \underset{\mathbf{h}}{\operatorname{argmin}} \|\mathbf{p} - \mathbf{W} \mathbf{W}^T \mathbf{C}_p \mathbf{h}\|_2^2 \quad (2.13)$$

As with the algebraic methods, we can find $\mathbf{h}^*$ by solving the following linear system for $\mathbf{h}$ in the least squares sense:

$$\mathbf{A}_p \mathbf{h} = \mathbf{p} \quad (2.14)$$

where $\mathbf{A}_p = \mathbf{W} \mathbf{W}^T \mathbf{C}_p$.

After computing the least squares solution $\mathbf{h}^*$ to the linear system of Eq. (2.14), the MR-FBP reconstruction is obtained by computing the FBP reconstruction with $\mathbf{h}^*$ as filter. The complete MR-FBP algorithm is given by:

Algorithm 2.1 MR-FBP reconstruction method

1. Calculate $\mathbf{A}_p = \mathbf{W} \mathbf{W}^T \mathbf{C}_p$.
 2. Find least squares solution $\mathbf{h}^*$ of $\mathbf{A}_p \mathbf{h} = \mathbf{p}$.
 3. Return $\operatorname{FBP}_{\mathbf{h}^*}(\mathbf{p})$ as MR-FBP reconstruction.
-

The linear system we need to solve in step 2) looks similar to the system $\mathbf{W} \mathbf{x} = \mathbf{p}$, which is solved in the least squares sense by algebraic methods (Eq. (2.8)). The difference is that the system of Eq. (2.14) has fewer unknowns: $\mathbf{A}_p$ has N_d columns if we impose that $\mathbf{h}$ is angle-independent, while $\mathbf{W}$ has N^2 columns. As we will show in Section 2.4, we are able to reduce the number of columns of $\mathbf{A}_p$ to $\mathcal{O}(\log N_d)$ by exponential binning, without reducing the reconstruction quality significantly.

Because of the large size of the linear system that needs to be solved in Eq. (2.8), algebraic methods usually use an iterative method to find least squares solutions. These iterative methods can sometimes converge slowly, and they introduce a new parameter to the method: the number of iterations to perform. Since the system of MR-FBP is smaller, we can use a direct method to find the least squares solution, making it both efficient and parameter-free.

2.4 Implementation

Although the number of unknowns of the MR-FBP method is smaller than that of algebraic methods, the actual implementation of the method is important to actually obtain a method that is computationally more efficient. In this section, we give details on how we implemented the MR-FBP method in this chapter to obtain the experimental results of Section 2.7. We will begin by discussing how the matrix A_p of Eq. (2.14) can be calculated efficiently. Furthermore, we will show that the size of the linear system can be reduced by exponential binning. Finally, we discuss the computational complexity of the MR-FBP method compared to existing methods.

Calculation of A_p

The first step of the MR-FBP method is to calculate the matrix $A_p = WW^T C_p$. Usually, the projection matrix W is not used directly by algebraic methods, since it can be very large. Instead, multiplication of W with an image x is calculated implicitly by calculating the line integrals of x on-the-fly [PBS11]. Similarly, multiplication of W^T with a sinogram p is calculated by backprojecting p on-the-fly. Here, we use a similar approach to calculate A_p , column by column.

Denoting a column j of A_p by $A_p(:, j)$, we have:

$$A_p(:, j) = A_p e_j \quad (2.15)$$

where e_j is a unit vector with all elements zero except for element j , which is equal to one. Using the definition of A_p , we see that:

$$A_p(:, j) = WW^T C_p e_j = WW^T C_e p = W FBP_{e_j}(p) \quad (2.16)$$

In other words, we can calculate a column j of A_p by creating a filtered backprojection reconstruction with filter e_j , and forward projecting the result. By doing this for every column, we can calculate matrix A_p .

Exponential binning

At this point, the MR-FBP linear system of Eq. (2.14) has $N_d N_\theta$ equations and N_d unknowns, one for each detector element. Although the system is smaller than the one used in algebraic reconstruction methods, which have $N_d N_\theta$ equations and N^2 unknowns, we can further reduce the number of unknowns by *exponential binning*. Exponential binning was also used successfully to reduce system sizes in [BK06] and Chapter 5 of this thesis.

In exponential binning, we assume that the filter h is a piecewise constant function of N_b pieces. Each constant region of the function is called a *bin*, and the boundary points of a bin β_i are defined by positions s_i and s_{i+1} : $\beta_i = (s_i, s_{i+1})$. The width of a bin is equal to the difference between its boundary points $d_i = s_{i+1} - s_i$. Since the filter value of a single bin is constant, we can represent a binned filter by a vector with




Figure 2.4: A filter with exponential binning, where we impose that the filter is symmetrical, and consists of several bins with a constant filter value. The size of the bins increases exponentially away from the center of the filter.

one element for each bin. The idea of exponential binning in the MR-FBP method is that we can reduce the number of unknowns of the linear system from N_d to N_b , by using fewer bins than detectors ($N_b < N_d$). The question remains how to choose the boundary points of the bins.

Looking at Fig. 2.2, we see that the Ram-Lak filter has most details around $n = 0$, and drops to zero relatively quickly for $|n| \rightarrow \infty$. This suggests that we should use small bins around $n = 0$, and can use larger bins further away from the center. In this chapter, we use bins with widths that increase exponentially away from $n = 0$. Specifically, we take $d_i = 1$ for $|i| < N_l$ and $d_i = 2^{|i| - N_l}$ for $|i| \geq N_l$, with β_0 being the central bin. The number of bins with width 1 is specified by N_l , where larger values lead to more detail around the center of the filter, but more unknowns as well. For the rest of this chapter, we used $N_l = 2$, unless specified otherwise. We can reduce the number of bins even more by making it symmetric, defining new bins $B_0 = \beta_0$ and $B_i = (\beta_i \cup \beta_{-i})$ for $i \neq 0$. A filter with exponential binning and $N_l = 2$ is shown in Fig. 2.4.

Since the bin width increases exponentially, we end up with $\mathcal{O}(\log N_d)$ bins. Therefore, by using exponential binning, we have reduced the number of unknowns of the MR-FBP method from N_d to $\mathcal{O}(\log N_d)$. The restrictions we impose on the filter by assuming it is piecewise constant and symmetrical can reduce the quality of the MR-FBP reconstructions. We will show in Section 2.7, however, that the quality does not decrease significantly by using exponential binning, while the time it takes to calculate the reconstructions greatly decreases.

The matrix $\mathbf{A}_p$ with an exponentially binned filter can again be calculated column by column. In order to do this, we change the filter $\mathbf{e}_j$ of Eq. (2.16) to a vector $\mathbf{q}_{B_i}$, in which each filter element included in bin B_i is set to one, and all other elements are set to zero.

Computational complexity

For many tomographic reconstruction methods, the most costly subroutines computationally are forward projecting and backprojecting, for which straightforward implementations take $\mathcal{O}(N_d N_\theta N)$ and $\mathcal{O}(N_\theta N^2)$ time, respectively, although faster implementations exist which use hierarchical decomposition [BB00]. We can compare the computational costs of different reconstruction methods by comparing the number of projection operations each method has to perform. Filtered backprojection consists of a single projection operation: the final backprojection of the filtered sinogram. Algebraic methods usually perform a few projection operations per iteration. The SIRT method, for example, performs two projection operations per iteration, and typically has to perform $\mathcal{O}(N_d)$ iterations to converge to an acceptable solution.

The MR-FBP method has to perform one forward projection and one backprojection for every column of $\mathbf{A}_p$ during its calculation. Because there are $\mathcal{O}(\log N_d)$ unknowns, MR-FBP has to perform $\mathcal{O}(\log N_d)$ projections. The total computation time of calculating $\mathbf{A}_p$ becomes $\mathcal{O}((N_\theta N_d N + N_\theta N^2) \log N_d)$. If we assume that $N_d \approx N$, which is often the case, the total computation time becomes $\mathcal{O}(N_\theta N^2 \log N_d)$. To summarize, FBP, SIRT, and MR-FBP have to perform $\mathcal{O}(1)$, $\mathcal{O}(N_d)$, and $\mathcal{O}(\log N_d)$ projections, respectively, which shows that MR-FBP has to perform significantly fewer operations than SIRT.

Of course, the MR-FBP method also has to find the least squares solution to its linear system of Eq. (2.14). Because of its smaller size however, we can use direct methods to find this solution, instead of the iterative methods used in algebraic methods. The direct method we used in this chapter to generate the results of Section 2.7, the `gels*` lapack routine, uses singular value decomposition, and can solve an $m \times n$ system in $\mathcal{O}(mn^2)$ time. Since the MR-FBP system has $N_d N_\theta$ equations and $\log N_d$ unknowns, the least squares filter $\mathbf{h}^*$ can be found in $\mathcal{O}(N_d N_\theta [\log N_d]^2)$ time. Summing both the calculation of $\mathbf{A}_p$ and of $\mathbf{h}^*$, the total computation time of the MR-FBP method becomes $\mathcal{O}(N_\theta N^2 \log N_d + N_d N_\theta [\log N_d]^2)$.

2.5 Additional constraints

The reconstruction quality of algebraic reconstruction methods can be improved by exploiting prior knowledge about the object that was scanned. One approach of exploiting this knowledge is to add an additional constraint to the system that is minimized. Formally, such a reconstruction $\mathbf{x}^*$ can be found by solving the following equation:

$$\mathbf{x}^* = \underset{\mathbf{x}}{\operatorname{argmin}} [\|\mathbf{p} - \mathbf{W}\mathbf{x}\|_2^2 + \lambda g(\mathbf{x})] \quad (2.17)$$

where $g(\mathbf{x})$ is a function depending on the type of prior knowledge that is exploited. For example, if one knows that the object that is reconstructed has a sparse gradient, total-variation minimization can be used by setting $g(\mathbf{x}) = \|\nabla \mathbf{x}\|_1$ [SP08]. The parameter λ controls the relative strength of the additional constraint compared to the data fidelity term $\|\mathbf{p} - \mathbf{W}\mathbf{x}\|_2^2$. The optimal value of λ is often difficult to find, as it depends on the scanned object and acquired projection data.

A similar approach can be applied to the MR-FBP method, by imposing an additional constraint on the optimal filter $\mathbf{h}^*$:

$$\mathbf{h}^* = \underset{\mathbf{h}}{\operatorname{argmin}} \left[\|\mathbf{p} - \mathbf{W}\mathbf{W}^T\mathbf{C}_p\mathbf{h}\|_2^2 + \lambda g(\mathbf{h}) \right] \quad (2.18)$$

Different functions f can be used to exploit various kinds of prior knowledge. In this chapter, we will use one example, where the change in intensity of the reconstructed image in the horizontal direction and vertical direction is minimized. This can be achieved by letting $g(\mathbf{h}) = \|\nabla_x \mathbf{W}^T\mathbf{C}_p\mathbf{h}\|_2 + \|\nabla_y \mathbf{W}^T\mathbf{C}_p\mathbf{h}\|_2$, where $\nabla_x \mathbf{f}$ denotes the horizontal gradient of image $\mathbf{f}$, and $\nabla_y \mathbf{f}$ the vertical gradient. The horizontal and vertical gradient can be approximated by the linear Sobel operators $\mathbf{D}_x$ and $\mathbf{D}_y$, which convolve the image with two-dimensional kernels $\mathbf{G}_x$ and $\mathbf{G}_y$, respectively:

$$\mathbf{G}_x = \begin{bmatrix} +1 & 0 & -1 \\ +2 & 0 & -2 \\ +1 & 0 & -1 \end{bmatrix}, \mathbf{G}_y = \begin{bmatrix} +1 & +2 & +1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix} \quad (2.19)$$

If we approximate the gradients by $\mathbf{D}_x$ and $\mathbf{D}_y$, we can add the additional constraints to the linear MR-FBP system, as additional equations:

$$\begin{bmatrix} \mathbf{W}\mathbf{W}^T\mathbf{C}_p \\ \lambda\mathbf{D}_x\mathbf{W}^T\mathbf{C}_p \\ \lambda\mathbf{D}_y\mathbf{W}^T\mathbf{C}_p \end{bmatrix} \mathbf{h} = \begin{bmatrix} \mathbf{p} \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix} \quad (2.20)$$

The least squares solution $\mathbf{h}_{GM}^*$ of this system can be found using standard methods, by solving:

$$\mathbf{h}_{GM}^* = \underset{\mathbf{h}}{\operatorname{argmin}} \left[\|\mathbf{p} - \mathbf{W}\mathbf{W}^T\mathbf{C}_p\mathbf{h}\|_2^2 + \lambda \left(\|\mathbf{D}_x\mathbf{W}^T\mathbf{C}_p\mathbf{h}\|_2^2 + \|\mathbf{D}_y\mathbf{W}^T\mathbf{C}_p\mathbf{h}\|_2^2 \right) \right] \quad (2.21)$$

The resulting method, which we call MR-FBP_{GM}, finds a filter that minimizes a weighted sum of the residual and the horizontal and vertical gradient of the resulting reconstruction. The method can improve reconstructions of objects that have a small gradient. In the case of noise in the projection data, MR-FBP_{GM} can improve reconstructions as well, since the gradient of the object is usually much smaller than that of image noise. Therefore, by reducing the gradient of the reconstructed image, we reduce the amount of image noise as well.

Similar to the MR-FBP method, we can calculate the matrix $\mathbf{A}_p$ of the MR-FBP_{GM} method column by column. For a column j , we can calculate the FBP reconstruction with filter $\mathbf{e}_j$. We can then forward project this reconstruction to obtain the top part of column j of the linear system shown in Eq. (2.20). The remaining part of column j can be calculated by applying the Sobel operators $\mathbf{D}_x$ and $\mathbf{D}_y$ to the FBP reconstruction. Since the gradient image calculations can be performed efficiently in Fourier space, the asymptotic computational complexity of the $\mathbf{A}_p$ calculation step of MR-FBP_{GM} is identical to MR-FBP. The resulting linear system consists of $N_d N_\theta + 2N^2$ equations and $\log N_d$ unknowns. Therefore, the linear system can be solved in




Figure 2.5: The three phantom images used in this chapter. PHANTOM1 is the Shepp-Logan head phantom, PHANTOM2 represents a cross section of an engine block, and PHANTOM3 is a difficult to reconstruct object with both discrete and continuous areas.

$\mathcal{O}((N_d N_\theta + N^2)[\log N_d]^2)$ time, which is a slightly higher complexity than MR-FBP without gradient minimization. The total computational complexity of MR-FBP_{GM} is equal to $\mathcal{O}(N_\theta N^2 \log N_d + (N_d N_\theta + N^2)[\log N_d]^2)$.

2.6 Experiments

To compare the performance of the MR-FBP and MR-FBP_{GM} methods with other methods, we implemented them using Python 2.7.3, PyCUDA 2012.1, and Numpy 1.6.3 [Oli07] built with ATLAS 3.10.0 [WP05]. The GPU implementations of the forward and backprojection operations are based on the ASTRA-Toolbox [PBS11], in which backprojection is not the exact transpose of forward projection for performance reasons. We applied MR-FBP on three phantom images and two experimental μ -CT datasets, comparing the results of MR-FBP with SIRT, an algebraic method, and FBP with three different standard filters: the Ram-Lak filter, the Hann filter, and the Shepp-Logan filter. We compare MR-FBP_{GM} reconstructions with MR-FBP, FBP, and SIRT reconstructions of one of the three phantoms, with noise in the projection data.

The three phantom images are shown in Fig. 2.5. Each phantom image is represented on a 4096×4096 pixel grid, on which projections are calculated. Afterwards, the projection data is rebinned to 1024 detector elements, and all reconstructions are calculated on a 1024×1024 pixel grid. We calculate reconstructions for varying numbers of projection angles, and compare them to the original phantom image, scaled to a 1024×1024 pixel grid by averaging 4×4 squares.

For the experimental data, we used two different objects scanned by μ CT scanners, with the acquired cone-beam projection data rebinned to a parallel beam geometry. The first object is a diamond, which was scanned by a Scanco 40 μ CT scanner using 1024 detector elements and 500 equidistant projection angles. The second object is a mouse femur, scanned by a Skyscanner 1172 μ CT scanner using 1200 detector elements and 360 equidistant projection angles. Reconstructions of the experimental data are shown in Fig. 2.6, calculated using SIRT and all projection angles. To test the reconstruction quality of MR-FBP compared to other methods, we apply them on projection data from small subsets of all projection angles, with a fixed separation between the angles.




Figure 2.6: Reconstructions of the μ CT datasets used in this chapter. The reconstructions were calculated using SIRT and all available projections.

For each experiment, we report the mean absolute error and structural similarity (SSIM) index [Wan+04] of reconstructions of the various methods. The mean absolute error is defined as:

$$e_p(\mathbf{x}, \mathbf{o}) = \frac{N_D^{-1} \sum_{i \in D} |\mathbf{x}_i - \mathbf{o}_i|}{\max \mathbf{o} - \min \mathbf{o}} \quad (2.22)$$

where $\mathbf{x} \in \mathbb{R}^{N^2}$ is the reconstructed image, $\mathbf{o} \in \mathbb{R}^{N^2}$ the correct image, and the average is taken over all N_D pixels within the central disc D of radius $N/2$. For the experimental data, the mean absolute errors and SSIM indices are calculated with respect to SIRT reconstructions from projection data of all available projections, shown in Fig. 2.6. The SSIM index measures the similarity between two images, and was designed to represent human visual perception more accurately than other metrics. A higher SSIM index corresponds with larger perceptual similarity between the compared images. For the phantom experiments, we also report the mean absolute residual, defined as:

$$e_r(\mathbf{x}, \mathbf{p}) = (N_d N_\theta)^{-1} \sum_{i=0}^{N_d N_\theta - 1} |(W\mathbf{x})_i - \mathbf{p}_i| \quad (2.23)$$

where $\mathbf{x} \in \mathbb{R}^{N^2}$ is the reconstructed image, and $\mathbf{p} \in \mathbb{R}^{N_\theta N_d}$ the measured projection data.

2.7 Results and discussion

Results for simulation phantoms

The mean absolute error, SSIM index, and mean absolute residual for PHANTOM1 are shown in Fig. 2.7 as a function of the number of projection angles N_θ . The results show that the mean absolute error and mean absolute residual of the MR-FBP reconstructions are significantly lower than those of all tested FBP methods. Even though FBP with the Shepp-Logan filter or Hann filter produces more accurate reconstructions than FBP with the Ram-Lak filter, MR-FBP is more accurate than all three. A similar result is




Figure 2.7: The mean absolute error, SSIM, and mean absolute residual of reconstructions calculated with different methods, for PHANTOM1. The methods shown are FBP with the Ram-Lak filter (FBP-RL), FBP with the Shepp-Logan filter (FBP-SL), FBP with the Hann filter (FBP-HN), SIRT with 200 iterations (SIRT-200), SIRT with 1000 iterations (SIRT-1000), and the MR-FBP method (MR-FBP).

found for the SSIM, with significantly higher indices for SIRT and MR-FBP, compared to all tested FBP methods. Compared to SIRT, MR-FBP produces reconstructions with slightly higher errors, lower SSIM indices, and higher residuals. Later results in Fig. 2.9 will show, however, that MR-FBP is significantly faster than SIRT at producing these reconstructions. Results for the other two phantom images are similar to those of PHANTOM1.

For all three phantoms, reconstructions of FBP with the Ram-Lak filter, SIRT, and MR-FBP are shown in Fig. 2.8 for 32 projection angles. Note that in all comparison images in this chapter, the pixel value that a certain graylevel represents is identical for all compared methods. A zoomed inset is included in most images, giving a better indication than the entire image of how the reconstruction will look at full resolution. Figure 2.8 shows that both MR-FBP and SIRT are able to reduce the number of streak artifacts compared to standard FBP. Visually, the sharpness of the MR-FBP and SIRT reconstructions is slightly lower than that of the FBP reconstructions. In some applications, the higher sharpness of the FBP reconstructions might be preferable despite its artifacts, especially when the user is familiar with the scanned objects and FBP artifacts. In other applications, and in common post-processing steps such as segmentation, the artifacts present in FBP reconstructions can be problematic, and MR-FBP might be preferable.

The time it takes to calculate the reconstructions of PHANTOM1 using the different methods is shown in Fig. 2.9. In Fig. 2.9a, the reconstruction time is shown as a function of the number of projections, for a fixed number of detectors $N_d = 1024$. The results show that MR-FBP is significantly faster than SIRT, but slower than FBP with a standard filter. Specifically, MR-FBP is around 20 times faster than SIRT with 200 iterations in these cases, which is expected since MR-FBP has to perform around $2 \log N_d = 20$ forward projections and backprojections, while SIRT with 200 iterations has to perform 400. Similar results are shown in Fig. 2.9b, where the reconstruction




Figure 2.8: Reconstructions of the phantom images, from data with 32 projections, for different reconstruction methods.




Figure 2.9: The reconstruction time of the PHANTOM1 image. In (a), $N = N_d = 1024$, and the reconstruction time is shown as a function of the number of projections N_θ . In (b), the number of projections is 64, and the reconstruction time is shown as a function of the number of detectors $N = N_d$.




Figure 2.10: Mean absolute error and SSIM of reconstructions of PHANTOM1 from data of 64 projections with various amounts of Poisson noise applied.

time is shown as a function of the number of detectors, for a fixed number of projections $N_\theta = 64$.

In Fig. 2.10, the mean absolute error of reconstructions of PHANTOM1 with different methods is shown, for data of 64 projections with various amounts of Poisson noise applied. The parameter I_0 indicates the amount of applied Poisson noise, with lower values corresponding to higher amounts of noise. Specifically, noise was applied by first transforming the projection data to virtual photon counts, where I_0 corresponds to the largest photon count of all detector elements. For each detector element, a new photon count is sampled from a Poisson distribution with the original photon count as expected value. The resulting noisy photon counts are transformed back to obtain noisy projection data.

Since the Ram-Lak filter amplifies high-frequency signals, the reconstructions of FBP with the Ram-Lak filter are of low quality when noise is present in the projection data. Other filters, like the Hann filter, suppress high-frequency signals, and therefore yield reconstructions of higher quality. Algebraic techniques, like SIRT, often include a form of regularization on the reconstruction image, yielding reconstructions of even higher quality when noise is present. The results of Fig. 2.10 show that, as expected, SIRT reconstructions have the lowest error and highest SSIM, while the FBP method with the Hann filter yields reconstructions with higher errors and lower SSIM, and FBP with the Ram-Lak filter produces reconstructions with the highest error and lowest SSIM. The MR-FBP method yields reconstructions with similar errors and SSIM indices to the SIRT method, for every noise level, but requires less computation time. Examples of reconstructions of data with two different noise levels are shown in Fig. 2.11, for FBP with the Ram-Lak filter, SIRT, and the MR-FBP method.

Results for experimental data

The mean absolute error and SSIM indices of the reconstructions of the experimental data, obtained by the different methods, is shown in Fig. 2.12. For the DIAMOND




Figure 2.11: Reconstruction images of PHANTOM1, for FBP, SIRT-200, and MR-FBP, with various amounts of Poisson noise applied to data from 64 projections.




Figure 2.12: The mean absolute error and SSIM of reconstructions of the experimental data, for different numbers of projections.

dataset, the mean absolute error of the MR-FBP method is slightly higher than SIRT when reconstructing from very few projections, but both are significantly lower than FBP with fixed filters. For the FEMUR dataset, the mean absolute error of the MR-FBP method is similar to that of SIRT, with the error of FBP again significantly higher. Similar results can be observed for the SSIM indices of the different methods for both experimental sets. Reconstructions of both experimental datasets, with data from 25 projections for the DIAMOND set, and 24 projections for the FEMUR set, are shown in Fig. 2.13. Similar to previous results, the reconstructions of SIRT and MR-FBP are




Figure 2.13: Reconstructed images of both experimental datasets, for data of 25 projections (DIAMOND) and 24 projections (FEMUR).

visually similar, both having less artifacts than the FBP reconstructions, but a lower sharpness as well.

Additional constraints

To test the reconstruction quality of the $\text{MR-FBP}_{\text{GM}}$ method, we used the PHANTOM2 phantom image, which has a sparse gradient, and calculated the mean absolute error of reconstructions obtained by different methods from data of 64 projections and various amounts of applied Poisson noise. The weight λ in Eq. (2.20) of the $\text{MR-FBP}_{\text{GM}}$ method was set to $27 + 1600/I_0$, which was experimentally verified to be a reasonable choice for this phantom. Results are shown in Fig. 2.14, where it is clear that by using $\text{MR-FBP}_{\text{GM}}$, which exploits prior knowledge about the gradient of the image, we are able to obtain significantly lower mean absolute errors and significantly higher SSIM compared to other methods. The computation times of a single reconstruction using the FBP, SIRT-200, MR-FBP, and $\text{MR-FBP}_{\text{GM}}$ methods were 3.02 ms, 970 ms, 68.5 ms, and 577 ms, respectively. Reconstructions of the PHANTOM2 image from data with a large amount of applied Poisson noise ($I_0 = 2^6$) are shown in Fig. 2.15.

Exponential binning

To investigate the influence of exponential binning on the reconstruction quality and computation time of the MR-FBP method, we calculated the mean absolute error and reconstruction time both with and without exponential binning, for 64 projections of PHANTOM1. With exponential binning, the mean absolute error was equal to 0.0287,




Figure 2.14: Mean absolute error and SSIM of reconstructions of PHANTOM2 from data of 64 projections with various amounts of Poisson noise applied. Shown are errors of FBP, SIRT-200, MR-FBP, and MR-FBP_{GM}.




Figure 2.15: Reconstructed images from 64 projections of PHANTOM2 with a large amount of applied Poisson noise ($I_0 = 2^6$).

and it took 0.0684 seconds to reconstruct the image. Without exponential binning, the mean absolute error increased slightly to 0.0289, while the computation time increased significantly to 6.1698 seconds. This result shows that by using exponential binning, the reconstruction quality of MR-FBP does not decrease significantly, while the reconstruction time is greatly improved. In fact, it seems that exponential binning has a slight regularizing effect on the reconstruction, since the mean absolute error is actually slightly smaller with exponential binning than without.

Computed filters

In Fig. 2.16, filters that are computed by the MR-FBP method (Eq. (2.13)) are shown in Fourier space for different reconstruction problems, with the standard Ram-Lak filter shown as well. A first observation is that the computed filters are significantly different from the standard Ram-Lak filter, except for low frequencies. A second observation is that the computed filters also differ significantly for different reconstruction problems. The computed filter for $N_\theta = 128$ is closer to the Ram-Lak filter than the one for $N_\theta = 64$, which is expected, since the Ram-Lak filter is ideal when an infinite number of projections is available. Furthermore, when noise is present in the projection data,




Figure 2.16: Filters computed by the MR-FBP method for different reconstruction problems ($N_d = N = 1024$) and the standard Ram-Lak filter, shown in Fourier space.

the computed filter has a lower response for high frequencies, filtering out some of the high-frequency noise. Finally, when reconstructing PHANTOM2 instead of PHANTOM1, the computed filter is different as well. These results show that the approach of calculating a filter specific to the reconstruction problem is valid, and that there is no single filter that is ideal for every problem.

Limited angular range

In some applications of computed tomography, it is impossible to acquire projections over the full $[0, \pi]$ angular range. Since analytical methods assume that data of the full angular range is available, a limited angular range can lead to severe artifacts in the reconstructed image. Algebraic methods include the limited angular range in their model, and can usually produce more accurate reconstructions than analytical methods. When presented with data with a limited angular range, reconstructions of the MR-FBP method contain similar artifacts to those reconstructed by standard FBP as is visible in Fig. 2.17a and Fig. 2.17c. A possible cause for these artifacts is that the MR-FBP method uses an angle-independent filter: the projections of each angle are filtered by the same filter. Indeed, if we extend MR-FBP to use an angle-dependent filter, MR-FBP reconstructions are similar to those of algebraic methods, as shown in Fig. 2.17b and Fig. 2.17d. To use angle-dependent filters in MR-FBP we increase the size of the unknown filter $\mathbf{h}$ in Eq. (2.14) such that it includes a different filter for each projection angle. Of course, the size of the resulting linear system is increased as well, and the method will be significantly slower than angle-independent MR-FBP.

Truncated projection data

When the scanned object is larger than the field of view of the experimental setup, the acquired projection data is truncated at the edges of the detector. The MR-FBP method can be adapted to these cases by combining the standard approach of handling




Figure 2.17: Reconstructions from 64 projections of PHANTOM2 with a limited angular range of $\frac{3}{4}\pi$. Reconstructions are shown for standard FBP, SIRT with 200 iterations, standard MR-FBP, and MR-FBP with an angle-dependent filter (MR-FBP- θ).




Figure 2.18: Reconstruction images of PHANTOM2, for FBP and MR-FBP, from projection data of 64 projections truncated to 256 detector elements with applied Poisson noise.

truncated data in the FBP method with the standard approach in algebraic methods. First, the projection data has to be padded with a smoothly decreasing function in order to avoid the sharp drop to zero at the edge of the detector, similar to the standard approach for FBP. Secondly, during calculation of the MR-FBP matrix, we backproject the padded projection data onto a reconstruction grid as wide as the padded detector, but forward project onto the unpadded detector, similar to the standard approach for iterative methods. By using this combined approach, the MR-FBP method is able to reconstruct truncated projection data. An example of a reconstruction from truncated projection data is shown in Fig. 2.18, where we artificially truncated projection data of PHANTOM2. The results show that the reconstruction quality of the MR-FBP method when reconstructing truncated data is similar to the quality achieved when reconstructing untruncated data.

Comparison to gaussian filtered FBP

A popular improvement to the standard FBP method is to apply a gaussian filter to the image after reconstruction. Since different filter widths can be used, there is a certain trade-off between the amount of artifacts and the sharpness in the resulting image. By using a wider filter, the artifacts of the FBP reconstructions are suppressed more, but the resulting image will be more blurred as well. In Fig. 2.19, a comparison




Figure 2.19: Partial line profiles of reconstructions of PHANTOM1 projection data of 64 projections, produced by MR-FBP and standard FBP with gaussian filtering with different window widths (σ values). The partial line profile of the ground truth, the PHANTOM1 image, is shown by a dotted line.

between the MR-FBP method and the FBP method with gaussian filtering is made by examining a partial line profile of reconstructions of PHANTOM1. The results show that the MR-FBP reconstruction matches the ground truth better, even with a similar level of artifact reduction. The mean absolute error of the different methods for the entire reconstruction is 0.023, 0.036, 0.029, and 0.027 for MR-FBP and FBP with $\sigma = 1$, $\sigma = 2$, and $\sigma = 4$, respectively. Figure 2.19 also shows that the reconstruction quality of the gaussian filtered FBP reconstructions depends highly on the chosen window width. For a given dataset, it can be hard to determine the correct width to use, since it depends on the number of projections and the amount of noise present in the projection data. Furthermore, the parameters are often determined by visual inspection, which is subjective by nature. The MR-FBP method, on the other hand, determines the filter by minimizing a well-defined metric, making it both objective and reproducible, and automatically adjusts the filter to the projection data and geometry, as shown in Fig. 2.16.

2.8 Conclusion

In this chapter, we introduced a novel reconstruction method for 2D parallel-beam tomographic reconstruction problems, the minimum residual filtered backprojection method (MR-FBP). A reconstruction calculated by MR-FBP is a standard FBP reconstruction with a custom filter. For each reconstruction, a new filter is calculated by minimizing the squared difference of the projections of reconstructed image with the measured projections. Since FBP is a linear operation on the projection data, the

optimal filter can be found by solving a linear system in a least squares sense, similar to algebraic reconstruction techniques. In this chapter, we showed that the linear system can be constructed efficiently column by column. The result is an efficient method that automatically determines a data-dependent filter based on an objective quality criterion, eliminating the need for subjective manual filter selection. Furthermore, the method can be extended to exploit certain forms of prior knowledge to improve reconstruction quality.

Results for both phantom data and experimental data show that MR-FBP is able to produce more accurate reconstructions than FBP with standard filters when presented with data from few projections. The reconstructed images obtained from MR-FBP are visually similar to SIRT, an algebraic method, but can be calculated significantly faster. Similar results are found when reconstructing from projection data with various amounts of Poisson noise applied. Additional experiments show that exponential binning does not decrease reconstruction quality, while reducing computation time significantly.

We showed that MR-FBP can be extended to incorporate certain forms of regularization based on features of the image to improve reconstruction quality. One example, where a term minimizing the gradient of the reconstructed image is added to MR-FBP, was examined in detail. Results for the resulting method, called MR-FBP_{GM}, show that by exploiting prior knowledge we indeed obtain more accurate reconstructions from noisy projection data compared to methods that do not exploit prior knowledge. Furthermore, because MR-FBP_{GM} is an extension of the MR-FBP method, it is able to calculate reconstructions in less time than SIRT.

The results from this chapter show that by exponential binning we are able to reduce the computation time of the MR-FBP method without reducing its reconstruction quality. Other bases for reducing the number of unknowns in the linear system can be used however, which might enable us to reduce computation time even further, or improve reconstruction quality. For example, it might be possible to include some prior knowledge about the scanned object in the choice of basis. Whether better bases can be found, and whether they are useful in practice, is subject to further research.

The current study focused on two-dimensional parallel beam tomography. An approach similar to the one used in this chapter can, however, also be applied to other tomographic settings for which linear filter-based reconstruction methods exist. For example, in three-dimensional cone-beam settings, a minimum residual Feldkamp-David-Kress (FDK) method could be formulated, based on the standard FDK reconstruction method [FDK84]. The reconstruction quality of such methods is subject to further research.

In this chapter, we showed that the linear system of MR-FBP can be calculated column-by-column by creating several FBP reconstructions with specific filters. Therefore, MR-FBP can be implemented relatively easy using existing FBP implementations. Furthermore, if an optimized implementation of FBP is available, MR-FBP can use this implementation, resulting in an optimized version of MR-FBP as well. Therefore, MR-FBP can be used relatively easily to improve reconstruction quality in practical applications.