Algebraic filter approach for fast approximation of nonlinear tomographic reconstruction methods

L. Plantagie,[a] K. J. Batenburg,[a, b, c]

[a] Centrum Wiskunde and Informatica, Scientific Computing, Science Park 123, Amsterdam, The Netherlands, 1098 XG

[b] iMinds-Vision Lab, University of Antwerp, Universiteitsplein 1, Building N, Wilrijk, Belgium, 2610

[c] Mathematical Institute, Leiden University, Niels Bohrweg 1, Leiden, The Netherlands, 2333 CA

Abstract. In this paper we present a computational approach for fast approximation of nonlinear tomographic reconstruction methods by filtered backprojection methods. Algebraic reconstruction algorithms are the methods of choice in a wide range of tomographic applications, yet they require significant computation time, restricting their usefulness. We build upon recent work on the approximation of *linear* algebraic reconstruction methods and extend the approach to the approximation of nonlinear reconstruction methods, which are common in practice. We demonstrate that if a blueprint image is available that is sufficiently similar to the scanned object, our approach can compute reconstructions that approximate iterative nonlinear methods, yet have the same speed as filtered backprojection.

Address all correspondence to: L. Plantagie, Centrum Wiskunde and Informatica, Scientific Computing, Science Park 123, Amsterdam, The Netherlands, 1098 XG; E-mail: Linda.Plantagie@cwi.nl

1 Introduction

Computed Tomography (CT) deals with the reconstruction of an object from a series of projections of this object, taken along a range of angles.[1–4] Depending on the application, projection images are typically acquired by a scanning device, using a photon or particle beam that is transmitted through the object (e.g. X-rays, electrons, neutrons). Besides extensive applications in medical imaging, tomography is a common technique in many other fields in academia (materials science, micro-biology) as well as industry (quality inspection, process monitoring).[5–9]

1

In this paper we focus on the reconstruction phase of the tomography pipeline, where an image of the original object is computed from the projections by a reconstruction algorithm. The characteristics of the reconstructed image depend not only on the set of input projection data, but also on the reconstruction algorithm employed. A range of reconstruction methods have been proposed in the literature, each having strong and weak points with respect to reconstruction quality, reconstruction time, and robustness.[2, 10–13] The reconstruction methods that are used often in practice can be divided into two categories: analytical methods and algebraic methods. Here, we use the term "algebraic methods" to refer to the category of algorithms that converge to a least-squares solution (e.g. the Simultaneous Iterative Reconstruction Technique (SIRT)), as well as the category of statistical methods such as the Expectation Maximization (EM) algorithm.

In the ideal situation, where the reconstruction problem can be represented by a continuous set of data, we can find an exact solution of the reconstruction problem using an analytical method. This method uses an inversion formula to obtain the analytical solution. By discretizing this inversion formula, a reconstruction algorithm is obtained that approximates the analytical result. Typically, such methods are based on the assumption that the projections can be sampled continuously and a full range of angles are available. In practice, interpolation techniques must be used to account for missing projection data.[1]

The Filtered Backprojection (FBP) algorithm – and its many variants – is the most

prominent example of an analytical reconstruction method. Due to its computational efficiency, ease of implementation, and high accuracy if sufficient data is available, FBP is extensively used in practice.[14] Reconstructions are obtained by convolution of the projection data with a filter, followed by a so-called backprojection step. The filter can be customized and affects the quality of the reconstruction. FBP is capable of computing accurate reconstructions if a large number of low-noise projections are available, sampled along the full angular range. For reconstruction tasks involving just a small number of projections, or limited-angle datasets, FBP reconstructions are typically plagued by serious artefacts which hamper image interpretation .

Algebraic methods solve a system of linear equations, which represents the discretized tomographic reconstruction problem. This equation system directly models the finite set of projections available in the actual scan, resulting in a very large and sparse matrix describing the equations. The computation time required to calculate a least squares solution is so high, that in practice iterative methods are used for its solution.[1] The algebraic reconstruction methods are known to produce more accurate reconstructions than FBP when few projections are available or with noisy data. A key drawback of these methods is their high computational cost.

A subclass of the algebraic reconstruction methods consists of the linear algebraic reconstruction methods (LARMs). An algebraic reconstruction method is linear if the

3

algorithm acts on the projection data as a linear operator. Examples of LARMs are ART, SART and SIRT.[3]

In Ref. 15, a method was introduced to create filters for FBP that are based on the operation of linear algebraic reconstruction methods. The reconstructions of FBP with these filters approximate reconstructions of the corresponding algebraic reconstruction methods. This method is known as Algebraic Filter FBP (AF-FBP). Using these filters, one can approximate the accuracy of algebraic reconstruction methods while at the same time attaining the computational efficiency of FBP. The construction of the filters is computationally intensive, but needs to be performed only once for a given set of geometrical parameters (i.e. number of projections and their corresponding angles). The filters do not depend on the object that is being reconstructed. An important limitation of the method for constructing these filters is that the underlying algebraic reconstruction method (ARM) is required to be linear.

Other methods have been proposed in literature to create filters for FBP.[16–19] There are two different approaches; creating filters based on theoretical derivations which are suited for every geometry,[16,17] and creating filters based on the geometry.[18,19] The AF-FBP method belongs to this second group. Like the filters in Ref. 19, an iterative reconstruction method is used to obtain the filters. An important difference between these two methods is that in Ref. 19 an object consisting of three thin rods is used in the calculations of the

filters, while AF-FBP for LARMs is object-independent.

Many common iterative tomographic reconstruction methods do not meet the linearity condition. Therefore, they cannot be approximated directly by the AF-FBP. Examples of these methods are *conjugate gradient least squares* (CGLS), and the statistical reconstruction method *expectation maximization* (EM).[20–22] The CGLS algorithm is mildly nonlinear, in the sense that the algorithm becomes a LARM as the number of iterations tends to infinity. In tomography, only a relatively small number of iterations are typically performed, not only for limiting the computation time but also due to the regularizing effect embedded in the method itself. The Expectation Maximization algorithm (EM) maximizes the log likelihood function of the reconstructed image for a given set of projection data, assuming that the observed data and the Radon transform of the object are related by a Poisson distribution. Such statistical reconstruction methods are known to yield superior reconstructions to LARMs if the projection data contains a high level of (Poisson distributed) noise.[23]

Even though each of the reconstruction methods mentioned above yields a different reconstruction, the forward operator in the underlying inverse problem solved by these methods is still a linear operator (i.e. a discretized Radon transform). We therefore intuitively expect that many nonlinear algebraic reconstruction methods (NLARMs) can *locally* be approximated by a linear method. By the term "locally" we refer here to the case where

the reconstruction of a similar image, called a *blueprint*, is already available and only the difference of the projections compared to this known image needs to be reconstructed. Hence we refer not to spatial locality in the image, but to locality in the space of images, where each image is considered as a point in a high-dimensional vector space.

In this paper we adapt the AF-FBP approach to approximate NLARMs that have this locally linear behavior, such that AF-FBP can still be applied in cases where a blueprint object is available. We present the results of a series of experiments performed to assess the nonlinearity present in the CGLS and EM methods. We subsequently examine the reconstruction quality of our adapted AF-FBP method and compare them with reconstructions obtained by applying the NLARM directly.

The ability to compute fast approximations of NLARM reconstructions is useful in cases where the object to be reconstructed is expected to contain (small) variations with respect to the blueprint, which is known in advance. Such cases can be found, for example, in the field of nondestructive testing and inspection. Our approach is based on the computation of an *algebraic filter* that depends on the blueprint image. This processing step is computationally demanding, as it involves a series of runs of the iterative algebraic method. Once the filter has been computed, new images can be reconstructed with the same speed as Filtered Backprojection. Our method will mainly be useful in scenarios where a large number of similar objects are scanned in sequence, and where algebraic

methods are required due to limitations in the scanning geometry or scanning time. In such cases, the approach enables the approximation of advanced algebraic methods, while still maintaining a low computation time, sufficient for real-time reconstruction.

The concept of replacing the full reconstruction problem by the task of reconstructing the difference from a blueprint object has been used in various nonlinear inverse problems in imaging (e.g. seismic inversion[24] and electrical impedance tomography) but is applied here for the first time to fast approximation of nonlinear tomographic reconstruction methods, thereby considerably extending the applicability of AF-FBP.

This paper is structured as follows. In Sect. 2 the algorithms are briefly discussed and notation is introduced. The AF-FBP method for linear algebraic reconstruction methods is briefly introduced in Sect. 3.1. In Sect. 3.2, we extend the AF-FBP approach to NLARMs. In Sect. 4 we describe the setup and results of a series of computational experiments, comparing the proposed AF-FBP method to other methods. We conclude with discussion and conclusions in Sect. 5.

2 Preliminaries

In this paper, the projection data are obtain using a source and detector that rotate around the object in the 2D plane. We use *parallel beams* that cover all image pixels for each projection angle. The approach presented here is not limited to the scanning geometry described here, but can be applied to other scanning geometries as well (e.g. fan-beam,

7

cone-beam). An additional rebinning step of the projection data is necessary for these geometries. First, we give a brief introduction to the concepts and notation of FBP, followed by a description of the key algebraic methods studied in this paper.

2.1 Filtered Backprojection

Our description of Filtered Backprojection introduces the same notation as used in Section II of Ref. 15, but is included here to keep the present article self-contained.

In the continuous tomography model that is used to describe Filtered Backprojection, the unknown image is represented by a finite and integrable function $f : \mathbb{R}^2 \to \mathbb{R}$ of bounded support. Projections $p(\theta, \cdot)$ of f are defined by the *Radon transform*

$$p(\theta, t) = (\mathcal{R}f)(\theta, t) = \int_{-\infty}^{\infty} f(t\cos\theta - s\sin\theta, t\sin\theta + s\cos\theta)\, \mathrm{d}s, \tag{1}$$

where $\theta \in [0, \pi)$ denotes the projection angle and $t \in \mathbb{R}$ refers to the signed distance of a projected line from the origin.

The unknown image f can be recovered analytically from its Radon transform based on the following inversion formula:

$$f(x, y) = \int_0^{\pi} \int_{-\infty}^{\infty} P(\theta, u)G(\theta, u)e^{2\pi i u t}\, \mathrm{d}u\mathrm{d}\theta, \tag{2}$$

where $t = x\cos\theta + y\sin\theta$, $G(\theta, u) = |u|$ and $P(\theta, \cdot)$ is the Fourier transform of $p(\theta, \cdot)$. The function $G(\theta, \cdot)$ acts as a filter on the projection data.[25]

In practical experiments, the detector that measures the projections $p(\theta, \cdot)$ is discretized as an array of detector elements, each measuring a single value. Moreover, the set of angles θ for which projections are available is also finite and discrete.

Let $R \in \mathbb{N}_{>0}$. For simplicity, we assume that the detector contains $l = 2R + 1$ detector elements of unit width and that it is symmetric around $t = 0$. Hence, the lines (θ, t) for which the Radon transform is observed by these detector elements correspond to the set $T = \{-R, -R + 1, \ldots, R - 1, R\}$ of values for the parameter t in the Radon transform. Furthermore, the finite set of projection angles is given by $\Theta = \{\theta_1, \ldots, \theta_d\}$. The FBP formula, which approximates Eq. (2), is then given by

$$f(x, y) \approx \frac{\pi}{d} \sum_{\theta \in \Theta} \sum_{\tau \in T} p(\theta, \tau) g(\theta, x\cos\theta + y\sin\theta - \tau), \tag{3}$$

with g the inverse Fourier transform of G. In practice, various filters are used in FBP, combining the ideal ramp filter with smooth windowing functions. Examples are the Ram-Lak filter (using a hard frequency cut-off) and the Hann and Cosine filters (using soft windowing).[3]

2.2 Algebraic Reconstruction Methods

In this section we introduce the class of Algebraic Reconstruction Methods (ARMs). The input for these methods is the set of measured projections and a matrix describing the projection geometry, which together define a system of linear equations. Let $\boldsymbol{p} = (p_i) \in \mathbb{R}^m$ denote a vector containing the $m = dl$ measured detector values (with d the number of angles and l the number of detectors), and let $\boldsymbol{v} = (v_i) \in \mathbb{R}^n$ denote the object to be reconstructed. We now form the system of linear equations

$$\boldsymbol{W}\boldsymbol{v} = \boldsymbol{p}, \tag{4}$$

where the matrix $\boldsymbol{W}$, called the *projection matrix*, describes the geometry of the tomography setup. Each entry p_i of the projection data corresponds to an angle $\theta \in \Theta$ and detector bin $t \in T$, and is therefore also denoted as $p_{\theta t}$. Its value corresponds to the weighted sum of the values v_j on the line parameterized by (θ, t). Various models can be used to determine the weight of the contribution of the pixels on such a line, such as the line model, strip model,[26] and Joseph's model.[27] Algebraic reconstruction methods solve the system in Eq. (4) by starting at an initial guess of the solution (which can be the zero vector) and iteratively refining the solution. If the system is inconsistent, due to noise or other artefacts in the projection data, the residual of the equation system is often minimized with respect

to a particular norm. The results of ARMs can depend substantially on the particular ARM that is selected, for several reasons. First of all, the iteration process is typically terminated before convergence has been fully reached, such that the resulting reconstruction depends on convergence properties of the algorithm. Secondly, the norm that is minimized depends on the particular ARM. Thirdly, if the equation system is underdetermined (which is quite common in tomography), the particular solution that is chosen among all solutions depends on the ARM.

Algebraic reconstruction methods can be modeled as an operator $\mathcal{R} : \mathbb{R}^m \to \mathbb{R}^n$ that maps a vector $\boldsymbol{p} \in \mathbb{R}^m$ of projection data to a reconstructed image $\boldsymbol{u} \in \mathbb{R}^n$. Some ARMs are linear, meaning that their operation can be written as $\boldsymbol{u} = \boldsymbol{R}\boldsymbol{p}$ with $\boldsymbol{R} \in \mathbb{R}^{m \times n}$ a matrix (called the *reconstruction matrix* of the ARM). By definition, a linear ARM has the property that $\mathcal{R}(\lambda \boldsymbol{p} + \mu \boldsymbol{q}) = \lambda \mathcal{R}(\boldsymbol{p}) + \mu \mathcal{R}(\boldsymbol{q})$ for all $\lambda, \mu \in \mathbb{R}$ and $\boldsymbol{p}, \boldsymbol{q} \in \mathbb{R}^m$, which makes it possible to decompose its operation as a sum of reconstructions of unit vectors. This property is essential to the filter construction presented in Ref. 15. The well-known *SIRT* algorithm is an example of a linear method that can be used to compute algebraic filters.

Many ARMs used in practice are not linear. Here, we consider two examples. The *CGLS* algorithm is nonlinear, except in the limit case of an infinite number of iterations. In this limit, CGLS converges to the same reconstruction that would result by applying the *Moore-Penrose inverse*[28] to the projection data, which is a linear operator. For any finite

number of operations, CGLS is not linear. We therefore refer to CGLS as a *mildly nonlinear ARM*. Another example of a common nonlinear ARM is the *Expectation Maximization* (EM) algorithm. Even in the case of an infinite number of iterations EM does not have the linearity property. Therefore, we refer to EM as a *nonlinear ARM*. Below, we will outline CGLS and EM in more detail.

2.2.1 CGLS

The Conjugate Gradient (CG) algorithm is commonly used for solving large systems of linear equations, due to its fast convergence. It is originally designed for large sparse systems of equations with a symmetric and positive-definite matrix.[29] We apply a variant of this method, Conjugate Gradient Least Squares (CGLS), where the CG algorithm is applied to the system of normal equations $\boldsymbol{W}^T\boldsymbol{W}\boldsymbol{v} = \boldsymbol{W}^T\boldsymbol{p}$.

Put $\boldsymbol{u}^{(0)} = \boldsymbol{0}$, $\boldsymbol{r}^0 = \boldsymbol{p}$, and $\boldsymbol{d}^0 = \boldsymbol{W}^T\boldsymbol{p}$. The iteration which computes $\boldsymbol{u}^{k+1}$ ($k = 0, 1, 2, \ldots$) is then given by Eq. (5).

$$
\begin{aligned}
\alpha_{k+1} &= \|\boldsymbol{W}^T\boldsymbol{r}^k\|_2^2 \,/\, \|\boldsymbol{W}\boldsymbol{d}^k\|_2^2, \\
\boldsymbol{u}^{k+1} &= \boldsymbol{u}^k + \alpha_{k+1}\boldsymbol{d}^k, \\
\boldsymbol{r}^{k+1} &= \boldsymbol{r}^k - \alpha_{k+1}\boldsymbol{W}\boldsymbol{d}^k, \\
\beta_{k+1} &= \|\boldsymbol{W}^T\boldsymbol{r}^{k+1}\|_2^2 \,/\, \|\boldsymbol{W}^T\boldsymbol{r}^k\|_2^2, \\
\boldsymbol{d}^{k+1} &= \boldsymbol{W}^T\boldsymbol{r}^{k+1} + \beta_{k+1}\boldsymbol{d}^k,
\end{aligned}
\tag{5}
$$

For any system of linear equations, the CGLS algorithm converges to the least squares solution of minimal norm. One way to express this property, is that CGLS converges to $\boldsymbol{W}^\dagger \boldsymbol{p}$, where $\boldsymbol{W}^\dagger$ denotes the Moore Penrose inverse. As $\boldsymbol{W}^\dagger$ is a matrix, the limit behavior for CGLS corresponding to a large number of iterations is linear. In practice, one often performs just a few iterations, thereby implicitly imposing a form of regularization on the reconstructed image. As we will demonstrate in Sect. 4.1, CGLS is not linear in this case, but a linear model can be used as an approximation.

2.2.2 EM

Expectation Maximization (EM) aims at finding the reconstruction that is most likely to result in the measured projection data, where the measurements have been perturbed by Poisson noise. The detected photon counts in an X-ray scanner follow such a distribution. The reconstruction that maximizes the likelihood then satisfies Eq. (6);[2, 30]

$$\hat{\boldsymbol{u}} = \hat{\boldsymbol{u}} \, \boldsymbol{W}^T \frac{e^{-\boldsymbol{W}\hat{\boldsymbol{u}}}}{e^{-\boldsymbol{p}}}, \tag{6}$$

where $\boldsymbol{W}$ has column sum 1 and the arithmetic operations are performed element wise.[28, 31] Here, we focus on a straightforward Expectation Maximization method that iteratively solves Eq. (6). For a nonzero start solution $\boldsymbol{u}^0$, the $(k+1)$th iteration of the multiplicative

13

algorithm EM is given by Eq. (7);

$$\boldsymbol{u}^{k+1} = \boldsymbol{u}^k \, \boldsymbol{W}^T \frac{e^{-\boldsymbol{W}\boldsymbol{u}^k}}{e^{-\boldsymbol{p}}}. \tag{7}$$

Note that in practice, regularization is often employed to make the method more stable.[28]

3 Algebraic Filters for FBP

We now briefly discuss the AF-FBP approach for linear algebraic reconstruction methods, followed by our adaptation to make this method applicable to nonlinear algebraic reconstruction methods.

3.1 The Linear Case

In Ref. 15, a method is presented to created filters for FBP based on a linear algebraic reconstruction method (LARM) of choice. The reconstructions of FBP with these filters approximate the reconstruction quality of the LARM. As explained in Sect. 2.2, a linear reconstruction method can be represented by the reconstruction matrix $\boldsymbol{R}$. The value u_c of a single pixel $c \in \{1, \ldots, n\}$ in the reconstruction $\boldsymbol{u} = \boldsymbol{R}\boldsymbol{p}$ is given by

$$u_c = \sum_{\theta \in \Theta} \sum_{t \in T} r_{\theta t}^{(c)} p_{\theta t}, \tag{8}$$

where $r_{\theta t}^{(c)}$ is the entry of $\boldsymbol{R}$ in row c and column θt.

14

Define the center of pixel c as $(x_c, y_c) \in \mathbb{R} \times \mathbb{R}$, and let $t_c^{(\theta)} = x_c \cos\theta + y_c \sin\theta$ for $\theta \in \Theta$. Define the function $h^{(c)}$ by

$$h^{(c)}(\theta, \tau) = r^{(c)}_{\theta(\tau + t_c^{(\theta)})}, \tag{9}$$

where $(\theta, \tau) \in (\Theta \times (T - t_c^{(\theta)}))$. Then we can write u_c as

$$u_c = \sum_{\theta \in \Theta} \sum_{t \in T} h^{(c)}(\theta, t - t_c^{(\theta)}) p_{\theta t}. \tag{10}$$

Note that the formulas are valid for any set of projection data, either consistent (i.e. in the range of the Radon transform) or inconsistent. Comparing Eqs. (3) and (10) shows that $h^{(c)}$ acts as a filter on the projection data for pixel u_c. This filter is called an *algebraic filter*, and is determined by calculating the impulse response of pixel c for all the detector positions θt. Similarly, we can create filters for the other pixels in the image domain. It is reasonable to expect that these filters resemble $h^{(c)}$, at least for pixels in a neighborhood of u_c. As shown in Ref. 15, applying the algebraic filter of u_c to all the pixels in the image result in a reconstruction quality that is comparable to that of the LARM used to create the filter. This method is known as AF-FBP with the algebraic filter $h^{(c)}$. Note that in practice, one uses a Fourier convolution operation to evaluate the result of applying the filter to the projection data.

15

3.2 The Nonlinear Case

Nonlinear algebraic reconstruction methods (NLARMs) are used in many applications of computed tomography. As with linear reconstruction methods, a disadvantage of these methods is their expensive computational cost. If the AF-FBP approach as described in Sect. 3.1 could be applied, then this would lead to a method with relatively low computational cost that approximates the NLARM.

The method described in Sect. 3.1 requires linearity of the reconstruction operator $\mathcal{R}$ (see Sect. 2.2), which is clearly not satisfied for general NLARMs. We now introduce a variant of AF-FBP that *can* be applied to NLARMs, provided that the NLARM behaves *locally* as a linear transformation. We say that a reconstruction method is locally linear if for a set of projections $\boldsymbol{p} \in \mathbb{R}^m$ and a small perturbation $\Delta \boldsymbol{p} \in \mathbb{R}^m$, there exists a linear transformation $\boldsymbol{L_p} : \mathbb{R}^m \to \mathbb{R}^n$ such that

$$\mathcal{R}(\boldsymbol{p} + \Delta \boldsymbol{p}) \approx \mathcal{R}(\boldsymbol{p}) + \boldsymbol{L_p}(\Delta \boldsymbol{p}). \tag{11}$$

As the matrix $\boldsymbol{L_p}$ can be seen as a linear reconstruction method that reconstructs the perturbation of the projection data with respect to the projection data $\boldsymbol{p}$, we can approximate this method by FBP with an appropriately chosen filter $h^{(c)}$.

A difference with the filter method for the linear case is the dependence of the matrix

16

L_p on the projection data p. For every new set of projections p the matrix L_p changes and the filter $h^{(c)}$ has to be recalculated. In certain practical applications, in particular in industrial tomography, a blueprint of the scanned object is already available, while one aims to reconstruct the deviations from this blueprint. In this case, the measured projection data p_m is the superposition of the forward projection p_b of this blueprint and the forward projection of the variation, denoted by Δp. The reconstruction $v = \mathcal{R}(p_m)$ is then given by $v = \mathcal{R}(p_b + \Delta p)$. Since $\mathcal{R}$ is a locally linear transformation, there exists a matrix $L_{p_b} \in \mathbb{R}^{m \times n}$ such that $v \approx \mathcal{R}(p_b) + L_{p_b}\Delta p$. Note that L_{p_b} is independent of the perturbation Δp.

In the remainder of this paper we will refer to the known image as the *blueprint image*. An example of a blueprint image is the cross-section of a turbine blade. Variations on this blueprint are cracks, holes or local variations in density, while the general shape and size of the blade is equal to that of the blueprint.

For these locally linear transformations we can create a filter with a similar method as described in Sect. 3.1 for linear algorithms. Instead of calculating an impulse response as in the linear case, we now calculate a local derivative around p_b, because L_{p_b} depends on p_b and acts on the perturbation Δp. Hence for a given pixel c, we calculate the filter values at $\theta \in \Theta, t \in T$ as stated in Algorithm 1.

This filter $h^{(c)}$ can be used in FBP in the same way as the (angle dependent) standard

17

Algorithm 1: Calculate filters for locally linear algebraic reconstruction algorithms

Data: $\mathcal{R} : \mathbb{R}^n \to \mathbb{R}^m$ an NLARM,

$\boldsymbol{p}_b \in \mathbb{R}^m$ the projection data of a blueprint image,

Θ the set of all projection angles,

T the set of all detector pixels.

Result: The algebraic filter $h^{(c)}$ corresponding to $\mathcal{R}$ for $\boldsymbol{p}_b$.

begin

    Choose a pixel c in the image

    Define $(x_c, y_c) \in \mathbb{R} \times \mathbb{R}$ the center of pixel c

    Define $\boldsymbol{e}_{\theta t}$ the unit vector with value 1 at entry θt for $\theta \in \Theta, t \in T$

    for $\theta \in \Theta$ **do**

        for $t \in T$ **do**

            $l_{\theta t}^{(c)} = [\mathcal{R}(\boldsymbol{p}_b + \boldsymbol{e}_{\theta t})]_c - [\mathcal{R}(\boldsymbol{p}_b)]_c$

        end

        Define $t_c^{(\theta)} = x_c \cos \theta + y_c \sin \theta$

        Define $h^{(c)}$ for $\tau \in T - t_c^{(\theta)}$ by $h^{(c)}(\theta, \tau) = l_{\theta(\tau + t_c^{(\theta)})}^{(c)}$

    end

end

filters and the filters for the linear algebraic reconstruction methods. The implementation of AF-FBP for NLARM is stated in Algorithm 2. We will refer to this algorithm as nlAF-FBP. In cases where the distinction between the linear and nonlinear variant of AF-FBP is not relevant for the discussion, we use the general term AF-FBP.

Algorithm 2: Applying the AF-FBP filters for NLARM

Data: $BP : \mathbb{R}^n \to \mathbb{R}^m$ the unfiltered backprojection operation,

         $\boldsymbol{p}_m \in \mathbb{R}^m$ the measured projection data,

         $\boldsymbol{p}_b \in \mathbb{R}^m$ the projection data of the blueprint image,

         Θ the set of all projection angles,

         T the set of all detector pixels.

Result: Reconstruction $\boldsymbol{v} \in \mathbb{R}^n$.

begin

    // Calculate the perturbation $\Delta\boldsymbol{p}$

    $\Delta\boldsymbol{p} = \boldsymbol{p}_m - \boldsymbol{p}_b$

    for $\theta \in \Theta$ **do**

        // Apply the angle dependent algebraic filter $h^{(c)}$
           on $\Delta\boldsymbol{p}$

        $\tilde{\boldsymbol{p}} = \sum_{\tau \in T} h^{(c)}(\theta, \tau - t_c^{(\theta)}) \Delta p_{\theta\tau}$

    end

    $\boldsymbol{v} = \mathcal{R}(\boldsymbol{p}_b) + BP(\tilde{\boldsymbol{p}})$

end

---

4 Experiments

To examine the performance of the nlAF-FBP method, a series of computational experiments has been carried out based on simulated projection data. Validating the approach is not straightforward, as the reconstruction accuracy depends not only on the algebraic method and its parameters, but also on the blueprint image and the scanned object. We have chosen to focus on five key-aspects: (i) to verify that the selected nonlinear algebraic methods indeed exhibit locally linear behavior; (ii) to demonstrate the image quality of the nlAF-FBP approximation for a set of realistic differences between blueprint and scanned object; (iii) to investigate robustness with respect to noise; (iv) to determine robustness with respect to image registration errors between the scanned object and the blueprint; and
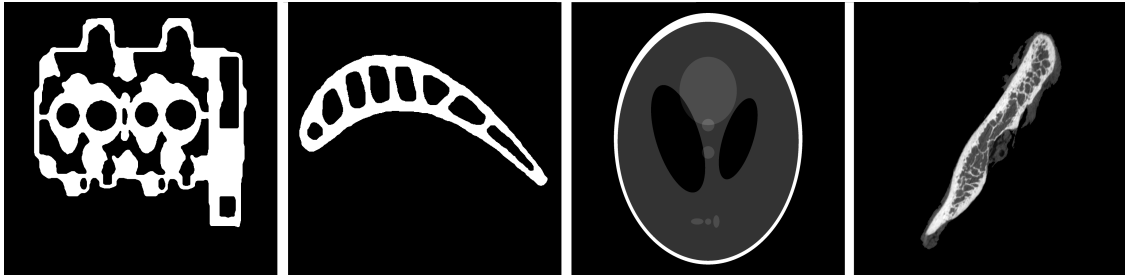
19

Fig 1 The blueprint images; (a) cylinder head, (b) turbine blade, (c) Shepp-Logan, (d) mandible.

(v) to investigate the robustness with respect to beam hardening artefacts.

As blueprint images, we consider both binary images and blueprints having continuous grey levels. They are shown in Fig. 1 and correspond to (1a) a cross-section of a cylinder head, (1b) a cross-section of a turbine blade, (1c) the well-known Shepp-Logan phantom, and (1d) a cross-section of a mandible. All the blueprint images are defined on a grid of 2044×2044 pixels. The cylinder head and turbine blade phantoms are representative for inspection tasks in industrial tomography, which we consider the primary application target of our nlAF-FBP approach. The Shepp-Logan and mandible phantom have been added to demonstrate how the algorithm performs for objects with multiple and continuous grey levels, respectively.

Real-world objects can typically not be represented on a pixel grid. To approximate the continuous nature of real objects, the reconstructions are performed on a coarser grid, where four phantom pixels correspond to one pixel in the reconstruction grid. The projection data are computed based on a detector of 511 bins with a width of four image pixels

20

per bin. We use an equiangular parallel beam geometry with a relatively small number of 64 projections, varying between 0 and 180 degrees, as this is a scenario where algebraic methods typically are preferred over Filtered Backprojection. We use the strip model to determine the contribution of each image pixel to a projection ray.[26]

The following reconstruction methods are used in the computational experiments:

Algebraic reconstruction methods

CGLS Conjugate Gradient Least Squares; see Sect. 2.2.1. The start solution in Eq. (5) is $u^0 = 0$. The total number of iterations is $K = 10$.

EM Expectation Maximization; see Sect. 2.2.2. The start solution in Eq. (7) is $u^0 = 1$. The total number of iterations is given by $K = 50$.

Filtered backprojection methods

FBP-RL, FBP-Hann, FBP-Cos Filtered Backprojection with the standard Ram-Lak, Hann and Cosine filters; see Sect. 2.1.

nlCGLS-FBP Filtered Backprojection with an algebraic filter based on CGLS with 10 iterations. The filter coefficients for the central pixel in Alg. 1 are used to create the filter.

nlEM-FBP Filtered Backprojection with an angle-dependent algebraic filter based on EM with 50 iterations. For every blueprint image a non-zero pixel close to the central pixel in Alg. 1 is selected to calculate the filter coefficients.

For nlCGLS-FBP and nlEM-FBP, the algebraic filters are applied to the perturbation Δp as described in Alg. 2 to obtain the AF-FBP reconstructions. For FBP with standard filters (FBP-RL, FBP-Hann, FBP-Cos), the filters are applied to the measured projection data.

To quantify the quality of the reconstructed images, we consider two error measures: the deviation from the phantom itself (i.e. the unknown ground truth) and the deviation from a reconstruction obtained by the algebraic method that one tries to approximate. To compare the reconstructions to the phantom, the reconstruction is upsampled by a factor 4, replacing each reconstruction pixel by a block of 4×4 pixels with the same value. We denote this enlarged image of 2044×2044 pixels by $\hat{\boldsymbol{u}} = (\hat{u}_{ij}) \in \mathbb{R}^{n^2}$ with $n = 2044$. We define the *mean reconstruction error E_r* by

$$E_r = \frac{\sum\limits_{i,j} |\hat{u}_{ij} - v_{ij}|}{\sum\limits_{i,j} v_{ij}}, \tag{12}$$

where $\boldsymbol{v} = (v_{ij}) \in \mathbb{R}^{n^2}$ denotes the phantom. The *mean ARM reconstruction error E_r^{ARM}* is computed similarly, yet the reconstruction is compared to the ARM reconstruction (either CGLS or EM) instead of the phantom. Our approach is designed to approximate the

4.1 Local Linearity

The validity of the proposed approach is based on the assumption that nonlinear tomographic reconstruction methods exhibit approximately linear behavior in the vicinity of the blueprint image. We performed a set of experiments to validate this assumption. In these experiments, the value of a single pixel in the reconstruction is monitored while a perturbation of increasing norm is introduced in the projections. If the reconstruction method is locally linear, the value of this pixel should also depend linearly on the magnitude of the perturbation.

We examined this property for CGLS and EM by considering two sets of projection data p, q, for the blueprint and the scanned object respectively. The perturbation $\Delta p \in \mathbb{R}^n$ in Eq. (11) is defined as $\lambda(q-p)$ with $\lambda \in [0, 1]$. We compute both $\mathcal{R}(p) + \lambda(\mathcal{R}(q) - \mathcal{R}(p))$ and $\mathcal{R}(p + \lambda(q - p))$ for every value of $\lambda \in [0, 1]$ and compare the results for a particular pixel (located in the interior of the phantom) in the reconstruction.

Figure 2 shows the grey level of this pixel as a function of the parameter λ, for CLGS and EM in two scenarios: (i) a disk-shaped gap of grey level 0 and radius 50 pixels is introduced in the cylinder head phantom, creating an artifical gap; the observed pixel is

23

included in this disk. (ii) The images p and q are completely different; p is the Shepp-Logan phantom, whereas q is the cylinder head phantom.

The results are visualized using a blue (dotted) curve and a red (solid) line. The red line corresponds to the pixel value of $\mathcal{R}(p) + \lambda(\mathcal{R}(q) - \mathcal{R}(p))$, and the blue curve corresponds to the pixel value of $\mathcal{R}(p + \lambda(q - p))$, for $\lambda \in [0, 1]$ variable. Large deviations of the blue curve from the red line imply that the corresponding reconstruction method does not behave locally as a linear function.

(a)	(b)	(c)	(d)
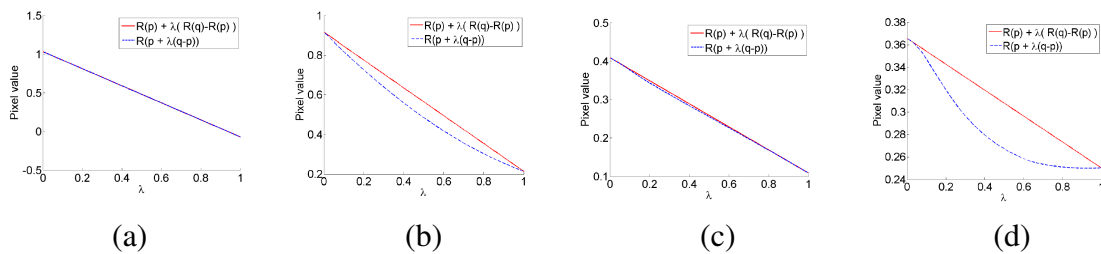
Fig 2 Plots of the local linearity for a particular image pixel; if the algorithm is locally linear, the blue and red lines should coincide. (a) CGLS, deviation is a black disk of radius 50 pixels; (b) EM, deviation is a black disk of radius 50; (c) CGLS, true object is completely different from blueprint; (d) EM, true object is completely different from blueprint.

Although these results depend on the particular pixel for which the grey value is plotted, they illustrate general observations about the two algorithms that we found in a broad set of observed pixels: (i) the CGLS algorithms shows almost perfectly linear behavior for small perturbations and close-to-linear behavior for large deviations; (ii) the EM algorithm clearly shows nonlinear behavior already for small deviations, although it may still be sufficiently linear for our purpose; for large deviations, EM exhibits strongly nonlinear

24

behavior.

4.2 Variations with cracks

To evaluate how the nlAF-FBP approach performs for realistic deviations between the blueprint and the scanned object, we now consider the cylinder head and turbine blade phantoms, which resemble objects that are common in the field of nondestructive testing. Four sets of test images have been manually created by introducing artificial cracks to the phantom objects. For each of the two objects, a set of *broad cracks* was created, as well as a set of *narrow cracks*. Each set consists of six images. An example for each set is shown in Figure 3, where the crack is magnified in the red box.
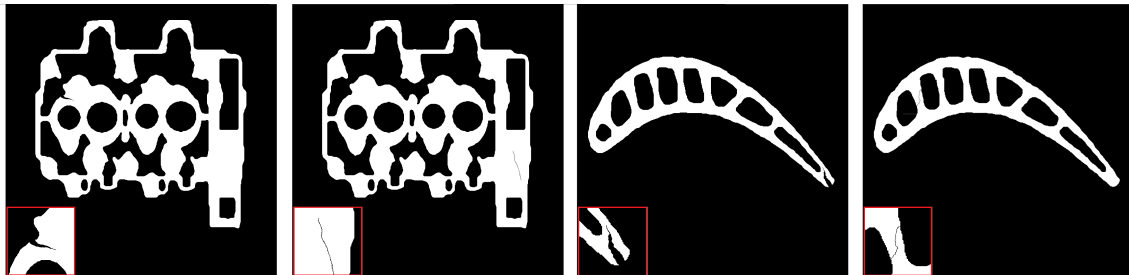
Fig 3 Example from each of the four categories of crack images: (a) cylinder head (broad); (b) cylinder head (narrow); (c) turbine blade (broad); (d) turbine blade (narrow).

In Table 1 we show the mean ARM reconstruction errors (i.e. compared to either CGLS or EM) over each of the four sets, for the proposed nlAF-FBP method vs. FBP with three standard filters (FBP-Cos, FBP-Hann, and FBP-RL). We observe that the mean

25

Class	nlAF-FBP	FBP-Cos	FBP-Hann	FBP-RL
CGLS				
Cylinder head (broad)	1.7e-3 (7.5e-4)	1.9e-2 (1.7e-3)	1.6e-2 (1.6e-3)	2.6e-2 (2.1e-3)
Cylinder head (narrow)	4.0e-4 (1.1e-4)	6.4e-3 (8.3e-4)	4.9e-3 (6.1e-4)	1.1e-2 (1.4e-3)
Turbine blade (broad)	4.7e-3 (3.0e-3)	4.7e-2 (1.9e-2)	3.9e-2 (1.5e-2)	6.9e-2 (3.3e-2)
Turbine blade (narrow)	8.4e-4 (3.1e-4)	1.3e-2 (3.5e-3)	1.0e-2 (2.7e-3)	2.3e-2 (5.9e-3)
EM				
Cylinder head (broad)	6.8e-3 (1.2e-3)	6.4e-1 (9.3e-4)	6.0e-1 (8.4e-4)	7.2e-1 (1.1e-3)
Cylinder head (narrow)	1.9e-3 (2.1e-4)	6.4e-1(1.6e-4)	6.0e-1 (1.3e-4)	7.2e-1 (2.7e-4)
Turbine blade (broad)	2.8e-2 (1.4e-2)	8.8e-1 (3.2e-3)	8.3e-1 (2.7e-3)	9.8e-1 (4.5e-3)
Turbine blade (narrow)	5.9e-3 (1.9e-3)	8.7e-1 (4.6e-4)	8.2e-1 (3.4e-4)	9.7e-1 (8.8e-4)

Table 1 Mean ARM reconstruction errors for the different categories of cracks, for the different reconstruction methods. For each set of experiments, the mean ARM reconstruction error is shown in black, and the standard deviation (over 6 crack images) in light grey.

ARM reconstruction error for nlCGLS-FBP is very small compared to FBP with standard filters. For nlEM-FBP the approximation of EM is not as good compared to CGLS, but the ARM reconstruction errors are still substantially smaller than those of FBP with standard filters for all cracks examined.

An illustration of the results for a particular testcase is shown in Figure 4. In all images, the red box contains a zoomed version of the crack. We first note that a standard Filtered Backprojection reconstruction without the use of a blueprint object (Fig. 4a) clearly shows the crack, yet also contains a considerable number of streak artefacts, making it difficult to distinguish between defects and artefacts. The CGLS (Fig. 4b) and EM (Fig. 4d) reconstructions are less prone to such artefacts. For both CGLS and EM, the nlAF-FBP reconstruction is visually very similar to the result of the algebraic method, whereas the FBP-Cos method yield quite different results. We observed similar results using other
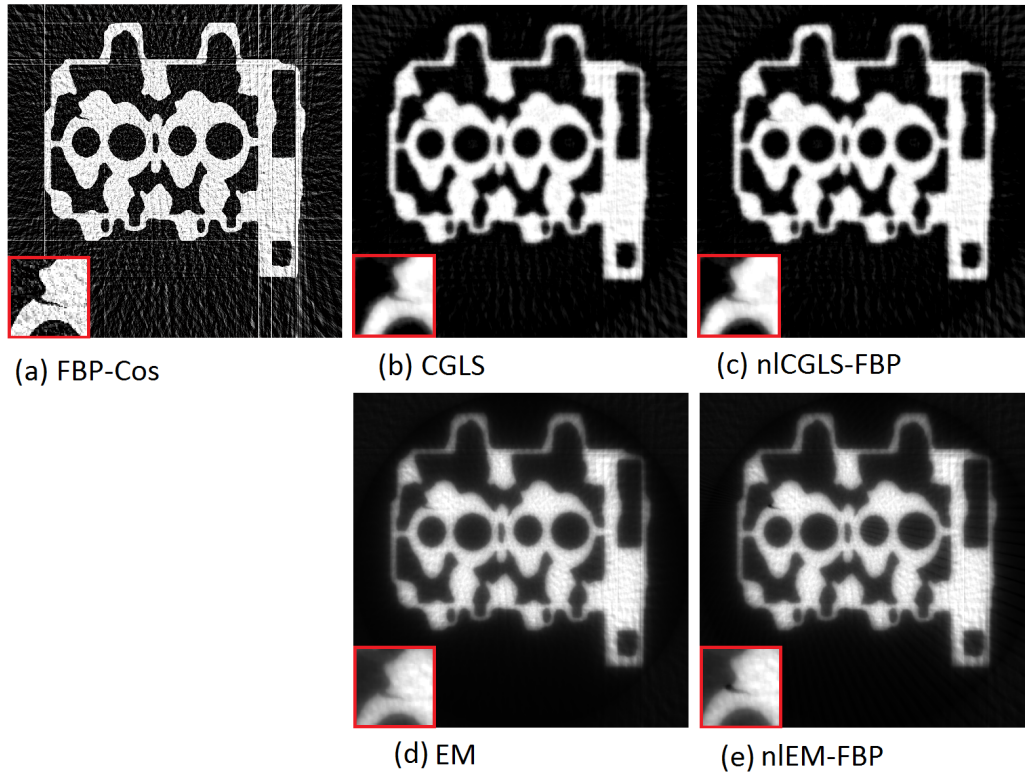
Fig 4 Illustration of reconstruction results for a broad crack in the cylinder head. The image in (a) is obtained by applying FBP with a Cosine filter, without using a blueprint. The images (c) and (e) are based on reconstruction of the difference between the object and the blueprint.

standard FBP filters.

We emphasize that the goal of our approach is to provide an accurate approximation of the algebraic method, which is not necessarily the same as providing the most accurate reconstruction. The results of Table 1 and Fig. 4 demonstrate that indeed the nlCGLS-FBP and nlEM-FBP methods provide a reconstructed image highly similar to the respective CGLS and EM reconstructions.

4.3 Robustness with respect to noise

So far, the projection data used in the experiments was noiseless. Since real (i.e. measured) datasets often contain noise, we also examine the robustness of nlAF-FBP with respect to noise. We applied Poisson noise to the projection data by first transforming the Radon transform data into photon counts (using the exponential function), subsequently generating noisy photon count by drawing from a Poisson distribution for each detector value, and then using the logarithm to convert the noisy data back to linearized projection data. Note that the resulting linearized projections can contain negative values, which are set to 0. In the results, the noise level is indicated by I_0, the photon count measured at a detector pixel without an object between source and detector (higher value means less noise). Fig. 5 shows a series of examples of reconstructed images for the mandible and turbine blade phantoms. Fig. 6 shows the mean reconstruction error and mean ARM reconstruction error for the mandible phantom as a function of noise level, averaging the results over a large number of noise realizations.

By visually comparing the reconstructions in Fig. 5, we notice that nlCGLS-FBP handles noisy projection data well compared to FBP with a standard filter. This is confirmed by the corresponding reconstruction errors, where the accuracy of nlCGLS-FBP is similar to that of the CGLS reconstructions. The accuracy of nlEM-FBP is also much better than that of FBP with standard filters, but it is not as close to EM as nlCGLS-FBP is to CGLS.

It is, however, a good approximation of EM based on the results of E_r^{ARM}. We observe similar results for the other phantoms and for other numbers of projection angles.

Hence in case of noisy projection data, nlAF-FBP yields results that approximate the NLARM well for both CGLS and EM, yet a more accurate approximation is observed for CGLS compared to EM. Compared to FBP using standard filters, the nlAF-FBP method (for both CGLS and EM) yields reconstruction with a strongly reduced noise level.

4.4 Robustness with respect to registration errors

Even if a blueprint of the scanned object is available, there may be registration errors between the scanned object and the blueprint, in addition to the deviations of the actual object structure. To examine the effect of such errors on the reconstruction accuracy of the AF-FBP method, we performed a series of experiments where the scanned object is obtained by rotating the blueprint around its center. The rotations are denoted by the number of degrees, where a positive number corresponds to a rotation clockwise and a negative number to a rotation counterclockwise.

In Fig. 7 the mean reconstruction error is shown for the Shepp-Logan phantom. For small rotations of at most a few degrees, nlCGLS-FBP is almost as accurate as CGLS. Using FBP with standard filters to obtain reconstructions results in considerably higher reconstruction errors. The nlEM-FBP is less tolerant to registration errors and becomes less accurate than EM already for rotations of 1 degree. We observed similar results for

29

the other phantoms.

4.5 Beam hardening

The previous experiments have been conducted using a monochromatic X-ray beam. For a monochromatic beam, the law of Lambert-Beer states that the measured projections (after log-correction) increase linearly with the thickness of a homogeneous object. In practice, however, a polychromatic X-ray beam is often used and the measured projections depend on the thickness of the object in a nonlinear way, resulting in beam hardening artefacts.[2] Especially for objects with metal parts, such as the cylinder head phantom and turbine blade phantom, beam hardening should be taken into account.

We performed a series of experiments to determine the behaviour of the nlAF-FBP method when the projection data are obtained using a polychromatic source. For the sake of brevity, we focus here on the CGLS algorithm. To simulate beam hardening, we apply the correlation between material thickness and attenuation as found by Ref. 32. The reconstructions are shown in Fig. 8. The corresponding mean ARM reconstruction errors for CGLS for both the cylinder head phantom and the turbine phantom are shown in Table 2.

We observe that also for a polychromatic X-ray beam the mean ARM reconstruction error for nlCGLS-FBP is smaller than that of FBP with standard filters. Hence nlCGLS-FBP reconstructions resemble the CGLS reconstruction better than reconstructions of FBP

Class	nlCGLS-FBP	FBP-Cos	FBP-Hann	FBP-RL
CGLS				
Cylinder head (broad)	1.7e-3 (4.8e-4)	2.7e-1 (3.2e-4)	2.4e-1 (2.7e-4)	3.4e-1 (4.3e-4)
Cylinder head (narrow)	9.9e-4 (5.2e-5)	2.7e-1 (5.5e-5)	2.4e-1 (4.5e-5)	3.4e-1 (9.8e-5)
Turbine blade (broad)	1.3e-1 (4.8e-4)	3.4e-1 (1.2e-3)	3.0e-1 (1.0e-3)	4.2e-1 (1.9e-3)
Turbine blade (narrow)	1.3e-1 (5.4e-5)	3.3e-1 (1.7e-4)	2.9e-1 (1.2e-4)	4.1e-1 (4.7e-4)

Table 2 Mean ARM reconstruction errors for the different categories of cracks for CGLS, using a polychromatic X-ray beam. For each set of experiments, the mean ARM reconstruction error is shown in black, and the standard deviation (over 6 crack images) in light grey.

with standard filters.
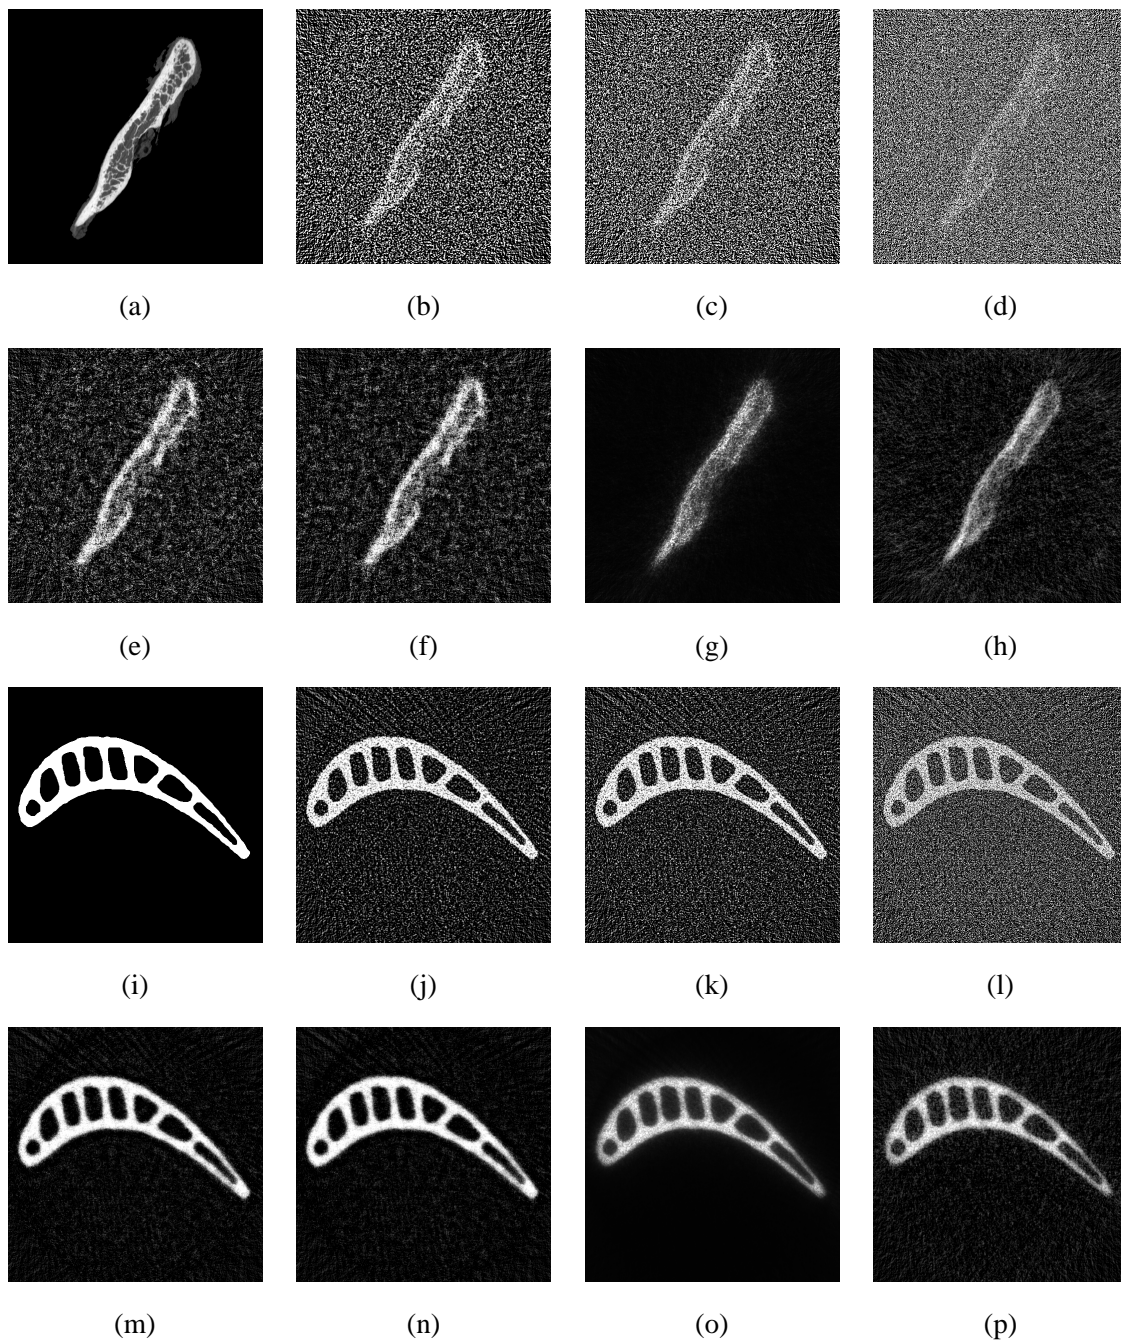
Fig 5 Reconstructions for the mandible and turbine phantom with noisy projection data; (a)-(h): $I_0 = 10^5$, (i)-(p): $I_0 = 10^6$; (a) blueprint, (b) FBP-Hann, (c) FBP-Cos, (d) FBP-RL, (e) CGLS, (f) nlCGLS-FBP, (g) EM, (h) nlEM-FBP, (i) blueprint, (j) FBP-Hann, (k) FBP-Cos, (l) FBP-RL, (m) CGLS, (n) nlCGLS-FBP, (o) EM, (p) nlEM-FBP.
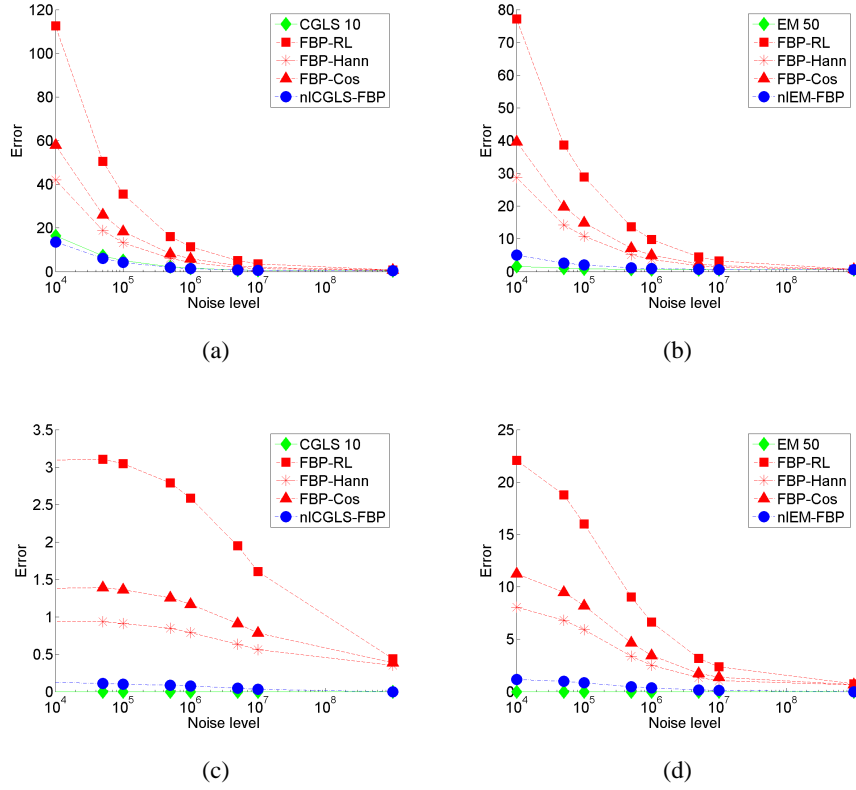
Fig 6 Mean reconstruction error with either the original phantom (E_r) or the ARM reconstruction (E_r^{ARM}) for the mandible phantom with varying Poisson noise levels I_0 applied to the projection data; (a) CGLS, E_r, (b) EM, E_r, (c) CGLS, E_r^{ARM}, (d) EM, E_r^{ARM}.
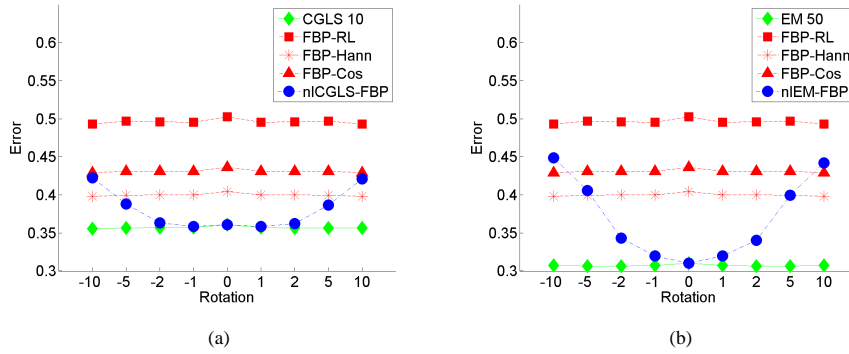

Fig 7 Mean reconstruction error with the rotated phantom (E_r) for the Shepp-Logan phantom; (a) CGLS, (b) EM.
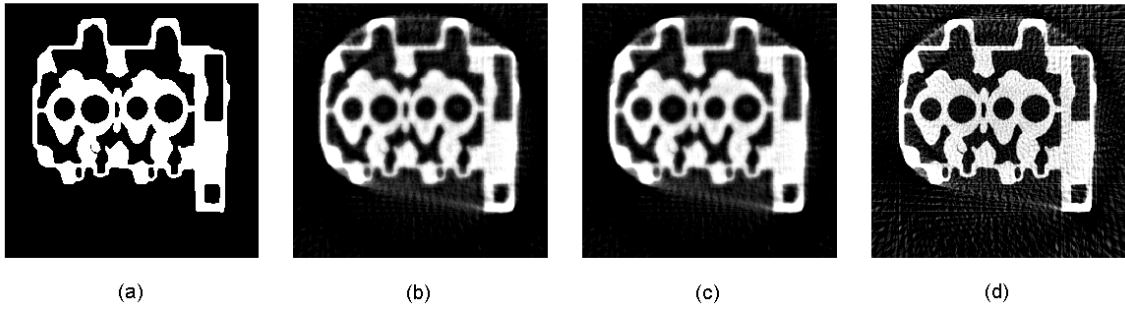
33

Fig 8 Reconstructions of the cylinder head phantom with a broad crack; (a) blueprint, (b) CGLS, (c) nlCGLS-FBP, (d) FBP-Hann. Note that the subfigures (b)-(d) have a different grey level range compared to the phantom, to provide a more clear visualization of the beam hardening artefacts.

5 Discussion and Conclusions

We have presented a novel approach for computing algebraic filters that can be used in FBP. With these filters we can approximate algebraic reconstruction methods with the computational efficiency of filtered backprojection. Contrary to the original AF-FBP method, which requires the underlying algebraic reconstruction method to be linear, our new approach is aimed at approximating *nonlinear* reconstruction methods, provided that they exhibit *local* linearity for reconstructions close to a give blueprint image. We experimentally investigated this local linearity property for two nonlinear methods: Conjugate Gradient Least Squares with a small number of iterations, and Expectation Maximization. The results show that while for CGLS the local linearity assumption is satisfied quite well, EM shows significant deviations from linear behavior. Using our approach for computing filters that act on the *difference* between the measured projection data and the projections of the blueprint, we performed experiments to assess the capabilities of our algorithm for both CGLS and EM. For CGLS, our approach results in reconstructions that are highly similar to the result of applying CGLS directly to the measured data, while reducing the computation cost to that of applying FBP (once the pre-processing has been done). The variations can be confined to small regions (such as cracks), but moderate registration errors with respect to the blueprint image are also tolerated well. These results are not restricted to monochromatic beams. Also for polychromatic beams nlCGLS-FBP approx-

imates the reconstruction accuracy of CGLS. For the EM algorithm, which has a stronger nonlinear nature, the results are mixed. Noisy projection data and cracks are in general handled well. For variations concerning the whole blueprint image, such as rotations, the approximation accuracy of nlEM-FBP degrades. Hence in several scenarios the nonlinear algebraic filter approach yields more accurate reconstructions than using a standard filter for FBP, but there are also cases where it fails. Investigating the exact conditions under which our method is favorable, and also the influence of the pixel location for which the filter is computed, will require further research.

Compared to FBP, algebraic methods are more suitable for limited-data scenarios. The ability to approximate the results of slow, nonlinear algebraic methods using very fast FBP methods opens up the possibilities of reducing the acquisition time, while keeping reconstruction quality constant. As already discussed in Ref. 15, computing the algebraic filters is computationally highly demanding, as it requires one to carry out a large number of ARM reconstructions (one per detector element). Even when using a moderately sized GPU cluster, the computation of a new filter may take a full day of computation time, based on the implementation in Ref. 33. As outlined in Ref. 15, the computational load can sometimes be considerably reduced by angle independent filters.

The computational overhead for calculating a set of filters for a blueprint becomes cumbersome when the blueprint changes frequently. The presented method is therefore

especially suitable for nondestructive testing and inspection of a small range of industrial objects, for which blueprints are readily available. In situations where algebraic reconstruction methods are preferred over FBP with standard filters, such as a limited range of available projection angles, the extensive computation time can be a bottleneck in the testing process. For scenarios where a large number of similar objects must be scanned, the nlAF-FBP approach allows for very fast – or even real-time – reconstruction of batches of objects once the filters have been computed.

References

1 F. Natterer, *The Mathematics of Computerized Tomography*, SIAM, Philadelphia (2001).

2 T. M. Buzug, *Computed Tomography: From Photon Statistics to Modern Cone-Beam CT*, Springer, Berlin (2008).

3 A. C. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging*, SIAM, Philadelphia (2001).

4 G. T. Herman, *Fundamentals of Computerized Tomography: Image Reconstruction from Projections, 2nd edition*, Springer, Berlin (2009).

5 P. Targowski, B. Rouba, M. Wojtkowski, and A. Kowalczyk, "The application of optical coherence tomography to non-destructive examination of museum objects," *Studies in Conservation* **49**, 107–114 (2004).

6 M. Defrise and G. T. Gullberg, "Image reconstruction," *Phys. Med. Biol.* **51**, 139–54 (2006).

7 W. A. Kalender, "X-ray computed tomography," *Phys. Med. Biol.* **51**, 29–43 (2006).

8 H. Sipila, "Moving object computer-tomography for luggage inspection," in *Applications of Signal and Image Processing in Explosives Detection Systems*, M. C. Connelly and S. M. Cheung, Eds., *Proc. SPIE* **1824**, 39–40 (1993).

9 S. V. Aert, K. J. Batenburg, M. D. Rossell, R. Erni, and G. Van Tendeloo, "Three-dimensional atomic imaging of crystalline nanoparticles," *Nature* **470**, 374–377 (2011).

10 L. A. Feldkamp, L. C. Davis, and J. W. Kress, "Practical cone-beam algorithm," *J. Opt. Soc. Am.* **1**, 612–619 (1984).

11 K. J. Batenburg and J. Sijbers, "Dart: a practical reconstruction algorithm for discrete tomography," *IEEE Trans. Im. Proc.* **20**, 2542–2553 (2011).

12 D. C. Durairaj, M. C. Krishna, and R. Muregesan, "A neural network approach for image reconstruction in electron magnetic resonance tomography," *Comput. Biol. Med.* **37**, 1492–1501 (2007).

13 R. J. Gaudette, D. H. Brooks, C. A. DiMarzio, M. E. Kilmer, E. L. Miller, T. Gaudette, and D. A. Boas, "A comparison study of linear reconstruction techniques

for diffuse optical tomographic imaging of absorption coefficient," *Phys. Med. Biol.* **45**, 1051–1070 (2000).

14 X. Pan, E. Y. Sidky, and M. Vannier, "Why do commercial CT scanners still employ traditional, filtered back-projection for image reconstruction?," *Inv. Problems* **25**, 123009 (2009).

15 K. J. Batenburg and L. Plantagie, "Fast approximation of algebraic reconstruction methods for tomography," *IEEE Trans. Im. Proc.* **21**(8), 3648–3658 (2012).

16 J. Zhou, P. Munshi, M. Maisl, and H. Reiter, "An investigation of higher-order tomographic filters using fourier transforms of reconstructed images," *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* **88**(4), 485–489 (1994).

17 Y. Wei, G. Wang, and J. Hsieh, "An intuitive discussion on the ideal ramp filter in computed tomography (i)," *Computers and Mathematics with Applications* **49**(5–6), 731–740 (2005).

18 T. Nielsen, S. Hitziger, M. Grass, and A. Iske, "Filter calculation for x-ray tomosynthesis reconstruction," *Phys. Med. Biol* **57**(12), 3915–3930 (2012).

19 H. Kunze, W. Haerer, J. Orman, T. Mertelmeier, and K. Stierstorfer, "Filter determination for tomosynthesis aided by iterative reconstruction techniques," *Proc. Fully 3D* , 309–312 (2007).

20 A. Björck, *Numerical Methods for Least Squares Problems*, SIAM, Philadelphia (1996).

21 L. A. Shepp and Y. Vardi, "Maximum likelihood reconstruction for emission tomography," *IEEE Trans. Med. Im.* **1**, 113–122 (1982).

22 L. Zhang, G. Zhang, C. Zhiqiang, Y. Xing, J. Cheng, and Y. Xiao, "X-ray spectrum estimation from transmission measurements using the expectation maximization method," *IEEE Nucl. Sci. Symp. Conf. Rec.* **4**, 3089–3093 (2007).

23 C. A. Bouman and K. Sauer, "A unified approach to statistical tomography using coordinate descent optimization," *IEEE Trans. Im. Proc.* **5**, 480–492 (1996).

24 L. De Barros, M. Dietrich, and B. Valette, "Full waveform inversion of seismic waves reflected in a stratified porous medium," *Geophysical Journal International* **182**(3), 1543–1556 (2010).

25 A. Markoe, *Analytic Tomography*, Cambridge University Press, New York (2006).

26 J. Zhu, X. Li, L. Yangbo Ye, and G. Wang, "Analysis on the strip-based projection model for discrete tomography," *Discrete Applied Mathematics* **156**(12), 2359–2367 (2008).

27 P. Joseph, "An improved algorithm for reprojecting rays through pixel images," *IEEE Trans. Med. Imaging* **1**(3), 192–196 (1982).

28 F. Natterer, *Mathematical methods in image reconstruction*, SIAM, Philadelphia (2001).

29 P. C. Hansen, *Rank-deficient and discrete ill-posed problems: numerical aspects of linear inversion*, SIAM, Philadelphia (1998).

30 K. Lange, M. Bahn, and R. Little, "A theoretical study of some maximum likelihood algorithms for emission and transmission tomography," *IEEE Trans. Med. Imaging* **6**(2), 106–114 (1987).

31 W. Van den Broek, S. Van Aert, and D. Van Dyck, "A model based atomic resolution tomographic algorithm," *Ultramicroscopy* **109**(12), 1485–1490 (2009).

32 J. J. Lifton, A. A. Malcolm, and J. W. McBride, "The application of beam hardening correction for industrial x-ray computed tomography," *5th International Symposium on NDT in Aerospace* (2013).

33 W. J. Palenstijn, K. J. Batenburg, and J. Sijbers, "Performance improvements for iterative electron tomography reconstruction using graphics processing units (gpus)," *J. Struct. Biol.* **176**(2), 250–253 (2011).

Linda Plantagie received the M. Sc. degree in mathematical sciences from the University of Utrecht, Utrecht, The Netherlands in 2009. She is currently pursuing the Ph. D. degree at the Center of Mathematics and Computer Science, Amsterdam, The Netherlands, focusing on tomographic reconstruction algorithms.

Kees Joost Batenburg received the M. Sc. degree in mathematics and the M. Sc. degree in computer science from the University of Leiden, Leiden, The Netherlands, in 2002 en 2003 respectively, and the Ph. D. degree mathematics in 2006. Since March 2010, he is a Senior Researcher at the Center of Mathematics and Computer Science, Amsterdam, The Netherlands, where he performs research in the filed of computational imaging. He also has two part-time professorships at the University of Antwerp (Wilrijk, Belgium) and at the University of Leiden (Leiden, The Netherlands) respectively.

List of Figures

4   Illustration of reconstruction results for a broad crack in the cylinder head. The image in (a) is obtained by applying FBP with a Cosine filter, without using a blueprint. The images (c) and (e) are based on reconstruction of the difference between the object and the blueprint.

5   Reconstructions for the mandible and turbine phantom with noisy projection data; (a)-(h): $I_0 = 10^5$, (i)-(p): $I_0 = 10^6$; (a) blueprint, (b) FBP-Hann, (c) FBP-Cos, (d) FBP-RL, (e) CGLS, (f) nlCGLS-FBP, (g) EM, (h) nlEM-FBP, (i) blueprint, (j) FBP-Hann, (k) FBP-Cos, (l) FBP-RL, (m) CGLS, (n) nlCGLS-FBP, (o) EM, (p) nlEM-FBP.

6   Mean reconstruction error with either the original phantom (E_r) or the ARM reconstruction (E_r^{ARM}) for the mandible phantom with varying Poisson noise levels I_0 applied to the projection data; (a) CGLS, E_r, (b) EM, E_r, (c) CGLS, E_r^{ARM}, (d) EM, E_r^{ARM}.

7   Mean reconstruction error with the rotated phantom (E_r) for the Shepp-Logan phantom; (a) CGLS, (b) EM.

8   Reconstructions of the cylinder head phantom with a broad crack; (a) blueprint, (b) CGLS, (c) nlCGLS-FBP, (d) FBP-Hann. Note that the sub-figures (b)-(d) have a different grey level range compared to the phantom, to provide a more clear visualization of the beam hardening artefacts.