Algebraic filters for filtered backprojection

Plantagie, L.

Citation

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

Cover Page

Universiteit Leiden

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

Author: Plantagie, L.
Title: Algebraic filters for filtered backprojection
Issue Date: 2017-04-13

3

SPATIAL VARIATIONS IN RECONSTRUCTION METHODS FOR CT

Abstract – In both Filtered Backprojection and algebraic reconstruction algorithms for tomography, the reconstruction of an object can depend on the position of the object within the discretized region, even if the object is aligned perfectly with pixel boundaries. In this chapter, we investigate this effect and report on a simulation study concerning *spatial dependencies* in these reconstruction methods. We demonstrate that for algebraic methods, these dependencies are influenced not only by the discretization within the reconstruction region, but also by the *shape* of the reconstruction region itself.

3.1 INTRODUCTION

Most reconstruction algorithms for CT can be assigned to either the class of analytical reconstruction methods, which are based on analytical inversion formulas of the Radon transform, or to the class of algebraic reconstruction methods, which start with a discretized inverse problem and then apply a numerical solver [1–3].

One of the fundamental differences between these two classes relates to the spatial locality of the reconstruction properties. Analytical inversion formulas are usually spatially invariant, in the sense that the value of a particular point in the reconstruction only depends on the measured values relative to the position of that point. If this dependency is known for a single point, it can be applied to all image points (e.g., pixel centers) to obtain a full reconstruction. Also, there is no predefined window outside which the reconstruction must be zero. The well-known Filtered Backprojection (FBP) algorithm is obtained by discretizing an analytical inversion formula of the Radon transform, and can therefore be expected to have approximately similar properties.

For *algebraic methods* on the other hand, there is no intrinsic reason why the reconstruction should be spatially invariant, and the reconstruction is constrained a priori to a *reconstruction region*, which is discretized and represented by a collection of basis functions. Outside this region, the reconstruction is automatically set to zero, as the exterior region is not covered by the support of the basis functions.

For both FBP and algebraic methods, there may be differences in the way projection values are sampled to determine the value of an image pixel, depending on the position of that pixel, due to discretization and interpolation effects within the projection model. As a consequence, reconstructing an object centered at one position within the reconstruction region may yield a different result from reconstructing this same object centered at another position. We refer to these variations as *discretization-effects*.

For algebraic methods, the *shape* and position of the reconstruction region with respect to the object can also influence its reconstruction. For example, if a line intersects the reconstruction region as a short segment in a corner, noise that is present in the projection for that line can have a strong impact on the values of the pixels on the small segment. For a line segment that has a longer intersection with the reconstruction region, the noise can be distributed among many pixels on

that segment. We refer to these local reconstruction variations, which depend on the shape of the reconstruction grid, as *shape-effects*.

In this chapter, we report on a case study that was carried out to investigate both discretization-effects and shape-effects for the FBP method and the Simultaneous Iterative Reconstruction Technique (SIRT), respectively. By moving an object across the reconstruction region and observing how its reconstruction changes with position, we keep track of both effects and obtain error maps that can be interpreted visually and analyzed quantitatively.

This chapter is structured as follows: In Section 3.2, we briefly review the discretization approach followed for FBP and SIRT, respectively. Section 3.3 describes the simulation experiments performed. The results of these experiments are presented in Section 3.4, mainly by providing a sequence of images that represent two different error measures, as a function of the position within the reconstruction region. In Section 3.5, the observations are discussed and future work in this direction is briefly outlined.

3.2 METHOD

The Filtered Backprojection (FBP) algorithm is obtained by discretizing the following inversion formula of the Radon transform (see Section 3.3.2 of [1] for details):

$$f(x,y) = \int_{\theta=0}^{\pi} \int_{\tau=-\infty}^{\infty} p(\theta, \tau - x\cos\theta - y\sin\theta)g(\tau)\,d\tau\,d\theta, \quad (3.1)$$

where $f : \mathbb{R}^2 \to \mathbb{R}$ denotes the unknown image, $p(\theta, \tau)$ denotes the measured line projection at angle θ and detector coordinate τ, and g denotes a *filter*, which determines how the detector values are weighted before backprojection to form the value at position (x, y). If we assume that p corresponds to the Radon transform of a certain original object, it is easy to see that translating this object over $(\Delta x, \Delta y)$ leads to a corresponding translation in the reconstruction over $(\Delta x, \Delta y)$. As a consequence of the discretization step in FBP, interpolation steps are required to compute an approximation of Eq. (3.1), leading to violations in this translational property, which we refer to as *discretization-effects*.

In *algebraic* reconstruction methods, the image is represented as a finite weighted sum of *basis functions* (see, e.g., Chapter 7 of [1] or

Section 6.3 of [2]). For this chapter, we limited ourselves to the reconstruction of 2-dimensional (2D) slices from 1D parallel beam projections using a standard pixel basis, yet the general methodology can be applied to 3D volume reconstruction using various types of basis functions, and various acquisition geometries.

When setting up an algebraic method, it is assumed that a certain *reconstruction region* is known, which completely contains the scanned object. Typically, this region is chosen to be either square or rectangular, while sometimes it is modelled as a disk. This region is then discretized along with the projection operator, leading to the following relation between the unknown image x and the measured projection data p:

$$Wx = p, \tag{3.2}$$

where $W = (w_{ij}) \in \mathbb{R}^{m \times n}$ denotes the *projection matrix*, $x = (x_j) \in \mathbb{R}^n$ is a vector representation of the pixel values in the unknown image, and $p = (p_i) \in \mathbb{R}^m$ represents the full set of measured detector values in all projections.

The exact projection matrix W depends on the selection of the reconstruction region, the choice and distribution of basis functions to represent the image within this region, and the model used for the projection operator.

The system in Eq. (3.2) is typically solved using iterative numerical solvers, as it is both very large and sparse. In this article, we consider one such iterative method called SIRT [4, 5], which converges to a weighted least-squares solution of the equation system.

Note that not all individual linear equations in Eq. (3.2) have the same algebraic structure. Each equation corresponds to a projected line. Depending on the intersection properties of that line with the discretized reconstruction region, the number of unknown pixel values that occur in the equation can vary, as well as their coefficients. As a consequence, the shape of the reconstruction region can influence the reconstruction of an object, depending on its location within that region, referred to as *shape-effects* of the reconstruction region.

3.3 EXPERIMENTS

To investigate discretization-effects and shape-effects for both FBP and SIRT, we performed a simulation study on the reconstruction of a small object that is placed at varying positions within the reconstruction region. All experiments were carried out using two different choices for the reconstruction region: (a) a square region of size $63{\times}63$ square pixels of unit size; (b) a pixelated circular region that is circumscribed around the square region of (a).

We compare the results for Filtered Backprojection (FBP) using the ramp-filter, which exhibits only discretization-effects, with the Simultaneous Iterative Reconstruction Technique (SIRT), which is expected to show both discretization-effects and shape-effects. For SIRT, 200 iterations are performed with a relaxation factor of 1. This iteration number ensures that convergence has been reached.

Projection data were simulated for a parallel detector geometry, using a detector consisting of 91 bins of unit size, thereby ensuring that the full reconstruction region is covered by the detector. The simulation was performed using a ray-driven projector based on the Joseph kernel to determine the contribution of an image pixel to each ray [6], implemented as a parallel operation on the GPU [7]. The projection angles of the parallel beam projections are regularly distributed between 0 and 180 degrees. The number of projection angles is kept fixed at 64. For the SIRT reconstruction, a forward projector based on the Joseph kernel was used.

As test objects, the square and cross images in Fig. 3.1a were used. The reconstructions of these objects, when placed in the center of a square reconstruction region, are shown in Fig. 3.1b and 3.1c.

3.3.1 *Experiments without noise*

In the first experiment, based on noiseless projection data, the test objects were moved across the reconstruction area. For each position of the object, its forward projection was computed and the object was reconstructed. The reconstruction within a small window around the object (a surrounding square, containing a boundary layer of one pixel thickness) was then shifted, placing the reconstructed object in the center of the reconstruction region. A comparison was made with the re-

(a) Original   (b) SIRT   (c) FBP

Figure 3.1: Test objects and their reconstructions when placed in the center of
the reconstruction region.

construction for which the object was placed in the center, using the
following two error measures: (a) the root mean square error (RMSE)
for all pixels in the window; (b) the absolute value of the difference in
average intensity within the object (AIE).

3.3.2 *Experiments that include noise*

In the second experiment, the test objects were again moved across
the reconstruction area. Poisson distributed noise was applied to the
projection data based on a flatfield photon count of 10^6 per detector
pixel. As the exact noise realization depends on the simulated photon
counts, which in turn are affected by discretization issues, we chose to
compare the reconstructions to the actual test object, instead of com-
paring to its reconstruction in the center. The reconstruction within a
window around the object (a surrounding square, containing a boun-
dary layer of five pixels thickness) was compared with the original ob-
ject, using the following two error measures: (a) the root mean square
error (RMSE) for all pixels in the window; (b) the absolute value of the
difference in average intensity within the object (AIE).

Figure 3.2: Intensity profile of an object at the center (blue line) and centered at pixel $(32, 8)$ (dashed red line); (a) square object, (b) plus object.

3.4 RESULTS

In this section, we examine the discretization and shape effects of SIRT and FBP for two test objects and two different shapes of the reconstruction grid.

3.4.1 *Noiseless projection data*

In the first series of experiments, the reconstructions of the noiseless projection data of a shifted object are compared with the reconstructions of the same object placed in the center of the reconstruction region. Reconstructions of the phantoms contain a variety of intensity levels within the reconstructed object. These intensity levels can be visualized by an intensity profile along a horizontal line through the center of the reconstructed object. In Fig. 3.2, the intensity profiles are shown for SIRT reconstructions of the test objects placed in the center of the reconstruction region and placed near the left boundary of the reconstruction region.

The reconstructions of the test objects clearly depend on their position in the reconstruction grid. The root mean square error for all pixels in the window is used to examine these spatial variations. In Fig. 3.3 the RMSE are shown for SIRT and FBP reconstructions of both test objects in a square reconstruction grid. The results are similar for the circular grid. Since the discretization-effects of SIRT and FBP appear to be very similar, we subtract the RMSE of FBP from that of SIRT

(a) SIRT   (b) SIRT   (c) FBP   (d) FBP

Figure 3.3: RMSE of reconstruction of the test objects using a square recon-
          struction grid without noise; (a) square object, (b) plus object, (c)
   square object, (d) plus object.

        (a)   (b)   (c)   (d)

Figure 3.4: Absolute value of the difference of RMSE between SIRT and FBP
          of the reconstruction of the test object; (a) square phantom, square
          grid, (b) square phantom, circular grid, (c) plus phantom, square
          grid, (d) plus phantom, circular grid.

to obtain an approximation of the shape-effects for SIRT. The absolute
value of this difference is shown in Fig. 3.4. Note that some scaling
was required to enhance the visibility.

For square reconstruction grids, reconstructions of an object near the
edge can differ substantially from a reconstruction of the same object
placed in the center of the reconstruction grid. Fig. 3.4 shows that, at
least in some cases, these shape-effects can be reduced by choosing a
different, for example circular, reconstruction grid. These results are
also supported by the second error measure (AIE), as is shown in Fig.
3.5 and 3.6.

3.4.2 *Noisy projection data*

In the second series of experiments, Poisson noise is applied to the
projection data of the shifted object. An example of a SIRT and FBP
reconstruction of the shifted square test object is shown in Fig. 3.7.

(a) SIRT   (b) SIRT   (c) FBP   (d) FBP

Figure 3.5: AIE of reconstruction of the test objects using a square recon-
struction grid without noise; (a) square object, (b) plus object, (c)
square object, (d) plus object.

(a)   (b)   (c)   (d)

Figure 3.6: Absolute value of the difference of AIE between SIRT and FBP
of the reconstruction of the test phantom; (a) square phantom,
square grid, (b) square phantom, circular grid, (c) plus phantom,
square grid, (d) plus phantom, circular grid.

(a) SIRT   (b) FBP

Figure 3.7: Reconstruction of the shifted square test object with Poisson noise
applied to the projection data.

Figure 3.8: Absolute value of the difference of RMSE between SIRT and FBP of the reconstruction of the test object with Poisson noise; (a) square phantom, square grid, (b) square phantom, circular grid, (c) plus phantom, square grid, (d) plus phantom, circular grid.


Figure 3.9: Absolute value of the difference of AIE between SIRT and FBP of the reconstruction of the test object with Poisson noise; (a) square phantom, square grid, (b) square phantom, circular grid, (c) plus phantom, square grid, (d) plus phantom, circular grid.

As mentioned in section 3.3 the reconstructions are compared to the original shifted image instead of the reconstruction of the object placed at the center of the reconstruction region. The spatial variations of SIRT due to shape-effects are again visualized by comparing both the RMSE and AIE measures of SIRT and FBP; see Fig. 3.8 and 3.9 for the RMSE and AIE measure, respectively.

Apparently, as suggested by Fig. 3.9, the total intensity within the object is invariant under the position of the test object. Fig. 3.8 shows that, also in the case of noisy projection data, SIRT reconstructions of an object depend on the position of the object in the reconstruction region. These spatial variations are influenced by the shape of the reconstruction region.

3.5 DISCUSSION AND CONCLUSIONS

The results of our case study demonstrate that significant discretization-effects can be observed in both FBP and SIRT reconstructions. Moreover, this effect is highly similar for both algorithms. For SIRT, the shape-effect also comes into play, yet mainly near the corners of a square reconstruction region. It appears that this effect can be mitigated by using a disk-shaped reconstruction region. The magnitude of shape-effects is increased by the influence of noise in the projection data, which can cause serious artefacts near the corners of the reconstruction region.

The actual position dependency may well depend strongly on the particular projection model used for the reconstruction. Here, we only considered the Joseph's method, which is broadly used in tomographic algorithms. In ongoing and future research, we are now focusing on the influence of different types of discretizations (e.g., blobs, wavelets) on the spatial dependencies, along with various projection models (e.g., lines, strips).

3

BIBLIOGRAPHY

[1]   A. C. Kak and M. Slaney. *Principles of Computerized Tomographic Imaging*. Philadelphia: SIAM, 2001.

[2]   G. T. Herman. *Fundamentals of Computerized Tomography: Image Reconstruction from Projections*. Berlin: Springer, 2009.

[3]   X. Pan, E. Y. Sidky, and M. Vannier. Why do commercial CT scanners still employ traditional, filtered back-projection for image reconstruction? *Inv. Problems* 2009; 25(12): 123009.

[4]   P. Gilbert. Iterative methods for the three-dimensional reconstruction of an object from projections. *J. Theor. Biol.* 1972; 36(1): 105–117.

[5]   J. Gregor and T. Benson. Computational analysis and improvement of SIRT. *IEEE Trans. Med. Imag.* 2008; 27(7): 918–924.

[6]   P. M. Joseph. An improved algorithm for reprojecting rays through pixel images. *IEEE Trans. Med. Imag.* 1982; 1(3): 192–196.

[7]   W. J. Palenstijn, K. J. Batenburg, and J. Sijbers. Performance improvements for iterative electron tomography reconstruction using graphics processing units (GPUs). *J. Struct. Biol.* 2011; 176(2): 250–253.