

Optimal Pre-Filtering for Linear Interpolation in Computed Tomography

Colas Schretter, Christoph Neukirchen, Georg Rose and Matthias Bertram

Abstract—Image interpolation is ubiquitous for image reconstruction in computed tomography (CT). For instance, the backprojection step of reconstruction algorithms is traditionally implemented with the simple linear interpolation model. This model is approximate but offers a good trade-off between speed and accuracy. Furthermore the implementation is natural and available on hardware graphics processing units (GPU). Approximation theory says that the image blurring induced by the triangular interpolation kernel can be compensated by enhancing the image with an all-pole recursive filter before resampling. This paper shows that the experimentally optimal pole differs from the one derived by theoretical approaches and that optimal pre-filtering leads to significant image quality improvement in term of signal to noise ratio (SNR). In fact, optimal pre-filtered linear interpolation outperforms the higher order cubic B-spline interpolation for image reconstruction in CT.

Index Terms—Image sampling, image reconstruction, interpolation, approximation.

I. INTRODUCTION

LINEAR image interpolations are widely used in the field of computed tomography (CT). Bilinear interpolation is traditionally used during backprojections when fetching the value of filtered line integrals for filtered-backprojection (FBP) tomographic reconstruction algorithms [1]. Trilinear interpolation is often used in conjunction with a numerical integrator for computing forward projections through digital volumetric images.

Linear interpolation from point samples relies on a compact triangle signal reconstruction kernel that is only a very crude approximation of the theoretically exact sinc kernel. Nevertheless, the computational performances, the ease of implementation, the implicit handling of image borders and the fair accuracy of the interpolated values made the uncontested popularity of linear interpolation schemes.

Furthermore, linear interpolations are implemented in hardware in graphical processing units (GPU) that are used to accelerate the backprojection and forward projection operations. Those two algorithms are ubiquitous in CT and are also the main bottleneck in both analytical and iterative tomographic image reconstruction algorithms. The impact of several image interpolation methods on the accuracy of forward projections has been evaluated by Xu and Muller [2].

This work received financial support from the Marie Curie program of the European Commission under contract MEST-CT-2005-020424.

C. Schretter, C. Neukirchen and M. Bertram are with Philips Research Europe, Weisshausstr. 2, 52066 Aachen, Germany (e-mail: colas.schretter@philips.com).

G. Rose is with the Otto-von-Guericke University, Universitätsplatz 2, 39016 Magdeburg, Germany.

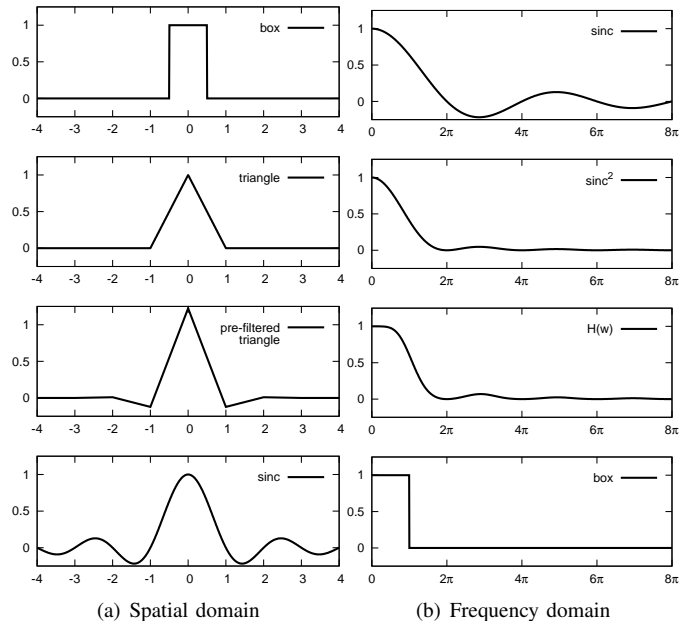


Figure 1. Spatial (a) and frequency (b) responses of various interpolation kernels. The frequency response is expressed as a fraction of the sampling frequency. The goal of pre-filtered triangle interpolation is to compensate smoothing in the pass-band while ensuring simultaneously an effective suppression of aliasing. Note that the support of the sinc kernel is infinite.

Traditional image interpolation approaches [3] do not consider the possibility to filter the image before resampling. However it has been shown a long time ago that optimal accuracy can be obtained for signal reconstruction by using a pair of optimized pre-filter and reconstruction post-filter [4]. A generalized sampling theory that does not assume band-limited signals justifies this approach [5]. Very accurate image interpolation have been reported for medical images [6].

Recently, approximation instead of interpolation has been considered in the image processing community. In particular, quasi-interpolation with infinite impulse response (IIR) filter [7] and least-square approximation with finite impulse response (FIR) filter [8] have been proposed independently. The aim of the present work is to demonstrate the potential of modern interpolation and approximation schemes for more accurate tomographic image reconstruction in CT.

This paper is structured as follows. Traditional image interpolation and modern image approximation methods are presented in section II. Section III shows the benefit of a simple image approximation scheme implemented by a pre-filter before linear interpolation for tomographic image reconstruction. Finally, conclusions are drawn in section IV.

II. INTERPOLATION AND APPROXIMATION

This section introduces the classical interpolation and approximation problems as the reconstruction of a continuous time-varying signal from uniformly distributed point samples. For interpolation, the reconstructed continuous signal is required to match exactly the sampled values at the sampling point locations. For approximation, this constraint is not required and this unveils possible improvements over interpolation. In this section, the term “reconstruction” refers to the reconstruction of a continuous signal from point samples.

The spatial impulse and the frequency response of classical interpolation kernels are compared in Fig. 1 along with the pre-filtered linear interpolation model investigated in this work. One can remark that the frequency response of the box kernel is superior to the triangle kernel in the pass band $\omega \in [-\pi, \pi]$. This translates into sharper interpolated image. However, the large ripples outside the pass band translate in very strong aliasing artifacts.

The recovery of high frequency components in the pass band is always underestimated with the triangle kernel and this causes blurring artifacts. Unfortunately, interpolation with the theoretically optimal sinc kernel (last row in Fig. 1) is impractical since the support of the spatial impulse function is infinite. Pre-filtering the signal before linear interpolation (third row in Fig. 1) allows a trade-off between blurring and aliasing artifacts but does not ensure the interpolation property.

A. Interpolation

Let a sequence of $N \geq 2$ point samples $s(k)$, $k \in [1, N]$, being sampled from a continuous function $f(t)$, $t \in \mathbb{R}$. If f is band-limited to frequencies $\omega \in [-\pi, \pi]$ and sampled at the Nyquist rate 2π , then it is well known that an exact reconstruction of the original signal is possible between the first and last samples by using sinc interpolation [9] such that

$$f(t) = \sum_{k=1}^N s(k) \operatorname{sinc}(t - k), \quad (1)$$

Since the sinc kernel has infinite support, exact interpolation requires to convolve each sample with the kernel. While feasible if the support of the image is finite, the very large computational cost of convolutions is often impractical. Furthermore, to prevent loss of information, interpolated values have to be computed and stored also for the infinity of samples lying outside the image boundaries.

Instead, the reconstruction of a continuous function \tilde{f} close to the original signal f can be computed effectively by piecewise linear interpolation from s such that

$$\tilde{f}(t) = \sum_{k=1}^N s(k) \beta^1(t - k), \quad (2)$$

where the reconstruction function $\beta^1(t) = \max(0, 1 - |t|)$ is the second order B-spline function also known as triangle kernel. Interpolation with the triangle kernel ensures that the reconstructed signal is continuous. This property is often preferred over the simplest interpolation with a box kernel, also called nearest neighbor interpolation.




Figure 2. Tomographic image reconstructions from 1024 projections, sampled from a voxelized phantom. Linear interpolation during backprojections introduces some blurring in the reconstruction of the phantom image. The square frames delineate the borders of the two close-up views shown in Fig. 5.

The impact of linear interpolation in tomographic image reconstruction can be seen in Fig. 2. Unfortunately, the reconstructed image suffers from overall blurring artifacts when using linear interpolations during backprojections.

Modern developments have shown that very accurate interpolation can be implemented effectively by pre-filtering the image before resampling. Unser et al. [6] recommend to implement image interpolation by the application of a theoretically derived pre-filter followed by convolution with a third order B-spline basis function. The cubic B-spline interpolation model is very popular nowadays and has been evaluated in our experiments for comparison purpose.

B. Approximation

Traditional interpolation ensures that the reconstructed signal $\hat{f}(t)$ equals the original signal f at the sampling points, hence when $t = \lfloor t \rfloor$. When this constraint is not a requirement, approximation schemes instead of interpolation have the potential for better reconstructions. Approximation for image resampling has been initially proposed by Mitchell and Netravalli [10] and Blinn [11] to find a qualitatively good visual compromise between blurring, aliasing, and ringing artifacts.

Although the derivations are different, the independent works of Condat et al. [7] and Dalai et al. [8] have shown that a least-square approximation of the continuous function can be implemented by pre-filtering the signal prior to interpolation with the simple triangle kernel. Their derivations assume that the continuous function is the cubic B-spline interpolation from the know samples. However, it is likely that the true underlying function is not a linear combination of B-spline basis function. In this case, better approximation can be obtained as demonstrated in our experiments.




Figure 3. Point spread function (PSF) of the FIR pre-filter of Marco Dalai (a), the IIR pre-filter of Laurent Condat (b) and the experimentally optimal pre-filter (c). Gray correspond to zero, darker values are negative and brighter values are positive. The PSFs of (a) and (b) looks very similar, however, the support of the FIR implementation (a) is compact.

When considering the triangle kernel as reconstruction post-filter, a general formulation of approximate reconstruction can be written as

$$\hat{f}(t) = \sum_{k=1}^N \hat{s}(k) \beta^1(t-k), \quad (3)$$

where the coefficients $\hat{s}(k)$ have to be computed from the original point samples $s(k)$ by filtering them with a symmetric pre-filter kernel φ such that

$$\hat{s}(k) = [s * \varphi](k). \quad (4)$$

Condat proposes an infinite impulse response (IIR) implementation with a simple all-pole recursive filter while Dalai proposes a finite impulse response (FIR) implementation of the pre-filter by discrete convolution. From the derivation of Condat, the negative pole for implementing the convolution with a simple IIR recursive filter is $2\sqrt{6} - 5$. The Fourier transform of the pre-filter can be extracted from the pole:

$$W_{\text{IIR}}(w) = \frac{6}{5 + \cos(2\pi w)}. \quad (5)$$

The equivalent discrete convolution implemented by FIR filtering is computed by Dalai as follows:

$$\hat{s}(k) = \frac{49}{40}s_0 - \frac{11}{90}s_1 + \frac{7}{720}s_2, \quad (6)$$

with $s_0 = s(k)$, $s_1 = s(k-1) + s(k+1)$ and $s_2 = s(k-2) + s(k+2)$. The support is arbitrarily limited to five samples. However, a larger support of seven samples have not shown any significant improvement in terms of image quality. From the coefficients of the FIR kernel, the Fourier transform of the pre-filter can be extracted:

$$W_{\text{FIR}}(w) = \frac{49}{40} - \frac{11}{45}\cos(2\pi w) + \frac{7}{360}\cos(4\pi w). \quad (7)$$

Despite very different expressions, W_{IIR} and W_{FIR} are surprisingly similar functions. Since convolutions in spatial domain are equivalent to multiplications in frequency domain, the resulting Fourier transform of the pre-filtered linear reconstruction is just

$$H(w) = \text{sinc}^2(w) W_{\text{IIR}}(w) \approx \text{sinc}^2(w) W_{\text{FIR}}(w). \quad (8)$$

This frequency response can be observed in comparison to linear interpolation in Fig. 1. The recovery of frequencies in the pass band is clearly improved at the cost of slight aliasing.

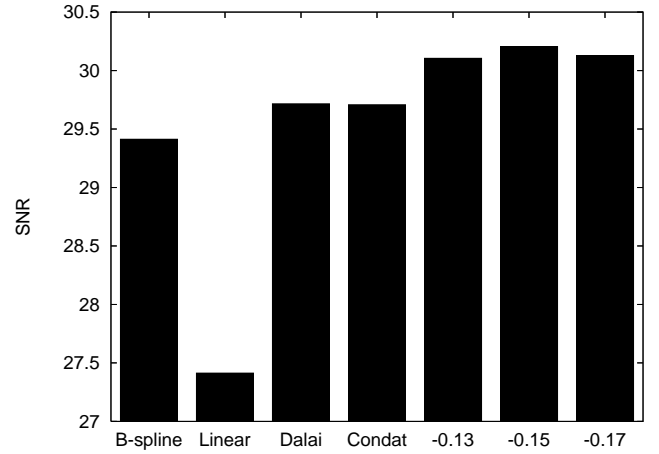


Figure 4. Quantitative analyses of image reconstruction in term of signal to noise ratio (SNR) from the original phantom image. Various values of the pole for the pre-filtering have been experimented. The experimentally optimum value of the pole (-0.15) differs from the one derived from approximation theory. The FIR filter derived by Dalai gives a very similar result than the IIR filter derived by Condat that uses a pole equals to $2\sqrt{6} - 5 = -0.1010205\dots$

Recursive filtering takes constant time per image element and requires two passes for the causal and anti-causal filtering. While discrete convolutions require more operations per pixel, in practice, in-place convolution with small kernels can be implemented to run as fast as IIR filtering. The choice between IIR and FIR is left to subjective appreciation. For two-dimensional images, the filter is applied successively in the vertical and horizontal directions.

The point spread functions of the FIR and IIR implementation of the theoretically optimal pre-filter in least-square sense are compared to the experimentally optimized pre-filter in Fig. 3. Since the support of the IIR filter is infinite, the point spread function (PSF) extends to the whole image. Despite their different intrinsic properties, experiments demonstrate that FIR and IIR implementations yield nearly identical results.

III. EXPERIMENTS

For experiments, a set of 1024 parallel-beam tomographic projections of 256×198 pixels has been computed from a phantom image of $256 \times 256 \times 198$ isotropic voxels of size equal to 1.36 mm. The goal is to reconstruct the original phantom image from the projection data with the best possible accuracy. The tomographic reconstruction algorithm is FBP and the ideal Ram-Lak ramp filter is used to preserves all frequency content. A large number of projection is used to alleviate possible issues with angular aliasing that typically translates into streak artifacts.

Several image reconstruction results have been compared using various interpolation methods for sampling filtered line integrals in projection space during backprojection. The traditional linear interpolation is compared to cubic B-spline interpolation and several pre-filtered linear interpolation models. The FIR and IIR implementations of the pre-filter for least-square optimal linear approximations give similar results.

Quantitative analyses conducted in terms of signal to noise ratio (SNR) are presented in Fig. 4. The SNR is conventionally




Figure 5. Close-up views of tomographic image reconstructions using various image interpolation methods. The resolution recovery is much better when using the cubic B-spline interpolation model in comparison to first order linear interpolations. However, when a suitable pre-filtering is applied before interpolations, linear interpolation have the potential to outperform the more costly cubic B-spline interpolation.

used for assessing the quality of image interpolation procedures [6]. Given a reference image P and a reconstructed image Q defined by N image elements, the SNR is a classical metric of the relative image similarity defined by

$$\text{SNR}(P, Q) = -10 \log \left(\frac{\sum_{i=1}^N (P_i - Q_i)^2}{\sum_{i=1}^N P_i^2} \right). \quad (9)$$

A transversal and a coronal slice of reconstructed volumetric images are shown in Fig. 2. Two selected close-up views are shown in Fig. 5. Gray is set to the attenuation of water and the window width equals 1000 HU such that black corresponds to the attenuation of air. The experimentally optimal pre-filter demonstrates dramatic improvement in image fidelity.

IV. CONCLUSIONS

This paper shows results of a heuristic approach to compensate for the typical blurring that can be observed when reconstructing an image with FBP. In the backprojection step of FBP, linear interpolation is used to fetch filtered line integrals in projection space. This simple interpolation model is exact only if the interpolation points exactly match pixel centers. A simple pre-filtering is used to transform interpolations into approximations and it has been observed that a sharper tomographic reconstruction can be obtained this way.

Linear interpolations are ubiquitous when using GPU implementations for backprojection. Therefore, optimal pre-filtering can improve significantly the accuracy of current image reconstruction codes. For analytic FBP algorithms, the pre-filter is applied in projection space before the backprojection step. For more accurate high-performance computation of line integrals

through volumetric images, the pre-filter could be applied in image space before sampling points along integration lines.

ACKNOWLEDGMENT

Colas Schretter thanks Laurent Condat, Marco Dalai, Thierry Blu and Philippe Thévenaz for discussions on image interpolation and approximation. He would also acknowledge Fabian Pilatus for reviewing the first draft of the manuscript.

REFERENCES

- [1] A. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging*. IEEE Press, New York, 1988.
- [2] F. Xu and K. Mueller, "A comparative study of popular interpolation and integration methods for use in computed tomography," in *Proc. IEEE International Symposium on Biomedical Imaging*, 2006.
- [3] T. Lehmann, C. Gonner, and K. Spitzer, "Survey: Interpolation methods in medical image processing," *IEEE Trans. Medical Imaging*, vol. 18, no. 11, pp. 1049–1075, 1999.
- [4] H. Malvar and D. Staelin, "Optimal FIR pre- and postfilters for decimation and interpolation of random signals," *IEEE Trans. Communications*, vol. 36, no. 1, pp. 67–74, 1988.
- [5] M. Unser, "Sampling - 50 years after shannon," *Proceedings of the IEEE*, vol. 88, no. 4, pp. 569–587, 2000.
- [6] P. Thevenaz, T. Blu, and M. Unser, "Interpolation revisited," *IEEE Trans. Medical Imaging*, vol. 19, no. 7, pp. 739–758, 2000.
- [7] L. Condat, T. Blu, and M. Unser, "Beyond interpolation: Optimal reconstruction by quasi-interpolation," in *IEEE International Conference on Image Processing*, vol. 1, 2005.
- [8] M. Dalai, R. Leonardi, and M. P., "Efficient digital pre-filtering for least-squares linear approximation," in *9th International Workshop on Visual Content Processing and Representation*, pp. 161–169, 2005.
- [9] S. Shannon, "Communication in the presence of noise," *Proc. of the Inst. of Radio Eng.*, vol. 37, no. 1, pp. 10–21, 1949.
- [10] D. Mitchell and A. Netravali, "Reconstruction filters in computer-graphics," *ACM SIGGRAPH*, vol. 22, no. 4, pp. 221–228, 1988.
- [11] J. Blinn, "Return of the jaggy," *IEEE Computer Graphics and Applications*, vol. 9, no. 2, pp. 82–89, 1989.