Three-dimensional Anisotropic Adaptive Filtering of Projection Data for Noise Reduction in Cone Beam CT

Andreas Maier and Rebecca Fahrig
Department of Radiology, Stanford University, Stanford, CA, USA

Lars Wigström
Department of Radiology, Stanford University, Stanford, CA, USA
Center for Medical Image Science and Visualization, Linkoeping University, Linkoeping, Sweden

Hannes G. Hofmann and Joachim Hornegger
Pattern Recognition Lab, Department of Computer Science, Friedrich-Alexander University of Erlangen-Nuremberg, Germany

Lei Zhu
George W. Woodruff School of Mechanical Engineering, Georgia Institute of Technology, Atlanta, GA, USA

Norbert Strobel
Siemens AG Healthcare, Forchheim, Germany

Purpose: The combination of quickly rotating C-arm gantry with digital flat panel has enabled the acquisition of three-dimensional data (3D) in the interventional suite. However, image quality is still somewhat limited since the hardware has not been optimized for CT imaging. Adaptive anisotropic filtering has the ability to improve image quality by reducing the noise level and therewith the radiation dose without introducing noticeable blurring. By applying the filtering prior to 3D reconstruction, noise induced streak artifacts are reduced as compared to processing in the image domain.

Methods: 3D anisotropic adaptive filtering was used to process an ensemble of 2D X-ray views acquired along a circular trajectory around an object. After arranging the input data into a 3D space (2D projections + angle), the orientation of structures was estimated using a set of differently oriented filters. The resulting tensor representation of local orientation was utilized to control the anisotropic filtering. Low-pass filtering is applied only along structures to maintain high spatial frequency components perpendicular to these. The evaluation of the proposed algorithm includes numerical simulations, phantom experiments, and in-vivo data which were acquired using an AXIOM Artis dTA C-arm system (Siemens AG Medical Solutions, Forchheim, Germany). Spatial resolution and noise levels were compared with and without adaptive filtering. A human observer study was carried out to evaluate low-contrast detectability.

Results: The adaptive anisotropic filtering algorithm was found to significantly improve low contrast detectability by reducing the noise level by half (reduction of the standard deviation in certain areas from 74 HU to 30 HU). Virtually no degradation of high contrast spatial resolution was observed in the MTF analysis. Although the algorithm is computationally intensive, hardware acceleration using Nvidia's CUDA Interface provided an 8.9-fold speed-up of the processing (from 1336 s to 150 s).

Conclusions: Adaptive anisotropic filtering has the potential to substantially improve image quality and/or reduce the radiation dose required for obtaining 3D image data using cone beam CT.

I. INTRODUCTION

The combination of quickly rotating C-arm gantry with digital flat panel has enabled the acquisition of three-dimensional data (3D), or C-arm CT, in the interventional suite [1][2][3]. However, image quality is still somewhat limited since the hardware has not been optimized for CT imaging. With the acquisition of an increasing number of projections to improve the contrast resolution of soft-tissue objects, it is essential to minimize the radiation dose delivered during a procedure. The same is true for many conventional CT examinations, with cardiac CT being one example where, despite many recent developments of dose reduction strategies, the radiation dose is still relatively high [4][5][6][7]. It is also not uncommon to acquire more than one C-arm CT scan during an intervention so that the progression of therapy can be monitored, further motivating dose reduction strategies. A number of different reconstruction filters are usually available that offer a selection of image characteristics[8]. Inevitably, these conventional shift invariant filters can only provide a fixed compromise between spatial resolution and noise reduction. Reconstruction using larger pixel sizes combined with a thick multi-planar reformatted (MPR) display will also reduce noise, but at the expense of fine details which are lost due to the inherent spatial averaging. Adaptive anisotropic filtering, on the other hand, has the ability to reduce the noise level and improve low-contrast detectability without noticeable blurring of small higher contrast spatial objects. However, most of these approaches neglect the effects of the projection noise in reconstruction domain.

In this study, we have used 3D adaptive anisotropic filtering to reduce the noise level in X-ray projection data before performing tomographic image reconstruction. In our case, the projection data was acquired using a C-arm device equipped with a flat-panel detector (FD), but the same principles would also be applicable to conventional multidetector CT (MDCT) data.

A. NOISE IN PROJECTION AND RECONSTRUCTION DOMAIN

Noise in the projection domain may lead to complex structure in the reconstruction domain and makes noise-adaptive filtering in the reconstruction domain challenging [9]. The noise in the projection domain is mainly dependent on the number of detected photons in each detector element. This behavior is known to be distributed in a Poisson manner [10].

$$P(n) = \frac{\kappa^n}{n!} * e^{-\kappa} \tag{1}$$

where P(n) is the probability of n occurrences of an event and κ is the expected value of the distribution.
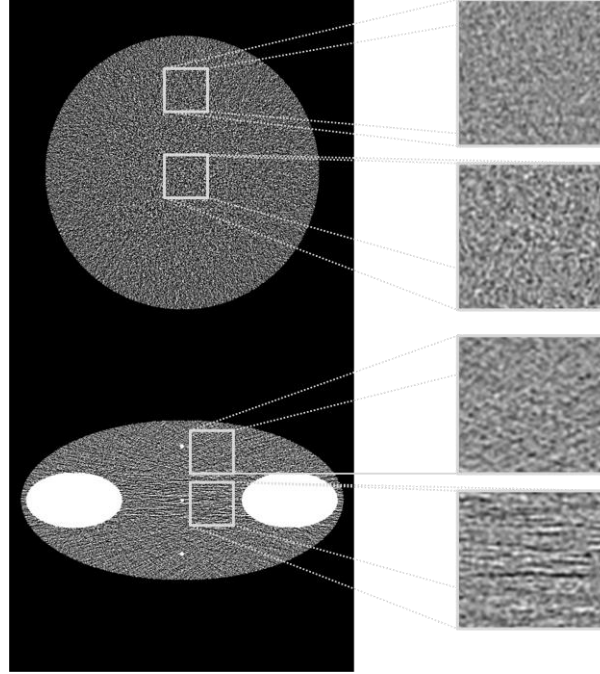
Given a random process $R(\kappa)$ which generates random numbers according to the Poisson distribution, the number of detected photons per detector element can hence be estimated by application of Lambert-Beer's Law

$$N_{\text{Noise}} = R(N) \tag{2}$$

$$N = N_0 * e^{-\int_0^S \mu(\eta)d\eta}$$

where s is the running index along an x-ray beam through the object with attenuation μ(η) at ray position η.

Figure 1: Distribution of noise in the reconstruction domain is dependent on the geometry of the imaged object. While the upper image shows almost Gaussian noise as it is perfectly round, the lower image shows streaks in some areas of the reconstruction due to the inhomogeneous mass density distribution.

This finding implies that the signal to noise ratio SNR is dependent on the number of observed photons:

$$SNR = \frac{\text{Signal}}{\text{Noise}} = \frac{\kappa}{\sqrt{\kappa}} = \sqrt{\kappa} = \sqrt{N} \tag{3}$$

Hence, the observed noise depends on the total amount of mass density through which the photon beam passed. This process makes the estimation of the noise distribution in reconstruction domain difficult as it depends on the geometry and the material properties of the object. In some areas of the reconstruction artificial noise patterns may emerge depending on the imaged object's geometry (cf. Figure 1).

Filtering in the projection domain has the advantage that noise is removed early in the pipeline inhibiting the propagation of noise into the reconstruction process, i.e. the assumption that the noise is distributed in a Poisson or Gaussian manner is still valid. This is especially important as some steps in the reconstruction process are known to worsen the noise properties of the reconstructed image [11].

75

B. ANISOTROPIC AND ADAPTIVE FILTERING IN THE LITERATURE

In the literature different uses of the terms "anisotropic" and "adaptive" are found. Weickert, for example, defines anisotropic diffusion as a process which alters the direction, i.e. the gradient, of a diffusion process [12]. This will be referred to as "anisotropic diffusion" in the following. Perona and Malik consider a process anisotropic if it is variant in either location or time [13]. In

80 contrary to Weickert's definition, they would classify any location or time variant process as anisotropic. Only a location- and time-invariant process is referred to as isotropic. As these processes are adaptive to some kind of input data, this will hereafter be called "adaptive". However, if the Greek origin of the word "isotropic" from iso (equal) and tropos (direction) is considered, one may find yet another definition of the word "anisotropic". With this definition everything which is equal in all directions is isotropic, while anything which prefers a certain direction is anisotropic. Under this definition even a gradient computation in a

85 certain direction is anisotropic. Hence, we will refer to processes which are identical in all spatial directions as "isotropic" and anything which is not uniform in all directions is called "anisotropic". Under this definition, both the filtering presented by Perona and Malik [13] as well as the anisotropic diffusion filtering by Weickert [12], are considered as anisotropic and adaptive.

Different approaches have been proposed for adaptive filtering of CT projection data. The known statistical properties of the measured signal were utilized by Hsieh to apply smoothing between detector channels when the signal level is low [14]. Nonlinear

90 filtering in the raw data domain was also investigated by Kachelrieß et al.[15]. They extended the filtering to utilize data from multiple detector rows and also successive views. Where the attenuation exceeded a specified threshold, the value was replaced with a result obtained by low-pass filtering within the three-dimensional neighborhood. By using all available dimensions, the filters could have a smaller extent while still gathering sufficient photon statistics, thereby improving spatial resolution compared to one-dimensional filtering.

95 Another interesting approach to adaptive filtering utilizing raw data was presented by Borsdorf et al. [16]. Here, the set of projections was split into two disjoint sets which were each used to create a reconstruction. Wavelet analysis of both independent reconstructions was then used to identify noise and structure. Coefficients which describe structure should be correlated while noise should be uncorrelated. With this information both reconstructions were finally composed into a single filtered image. This method led to a significant reduction of noise in the resulting image. However, with the limited number of projection images in

100 C-arm CT, this method is not well suited to our application, since the number of projections per reconstruction would be further

reduced by a factor of two.

Modeling of the noise in CT projection data as a Gaussian distribution with a nonlinear signal-dependent variance has been investigated and penalized weighted least-square smoothing (PWLS) has also been used to suppress noise [17][18]. A similar approach using a Poisson-based cost function is presented in [19]. Although these approaches will utilize prior statistical knowledge to apply an appropriate amount of smoothing, the actual content in the acquired projection data is not taken into account when controlling the shape of the filters.

Recently, anisotropic filtering in the raw data domain was also investigated by other groups. In [20] the PWLS algorithm was extended with an anisotropic controlling term which takes into account the difference in gray value of neighboring pixels to control the anisotropy of the filtering. A similar approach is also introduced in [21]. The paper proposed bilateral filtering in the raw projection data domain. Again the anisotropy is controlled by the gray value difference of pixels in a local neighborhood. Both approaches show significant noise reduction in the reconstructed image domain.

In the reconstructed image domain, different approaches have been proposed to locally adapt filters to the orientation of structures in the image data. Perona and Malik developed a multiscale smoothing and edge detection scheme, mathematically formulated as an anisotropic process that effectively performed smoothing within regions while minimizing the blurring across edges [13]. This work was later extended to 3D by Gerig et al [22]. More recent enhancement filters are based on diffusion derived from the affine heat equation [23]. This approach yields implementations that are more numerically stable. A simulation study in 2D using nonlinear diffusion filtering of CT projection data was performed by Demirkaya [24]. It showed a substantial improvement in SNR without significant loss of image resolution.

Image enhancement using anisotropic filters that are adaptively controlled by local image structures was introduced by Knutsson et al [25][26]. A set of orientation sensitive filters were used to estimate the local orientation of structures [27]. By utilizing the resulting orientation tensor field to control the anisotropy and degree of high pass filtering, smoothing could be achieved along structures while still maintaining the high spatial frequencies across edges [28]. This kind of adaptive anisotropic filtering has previously been applied to reconstructed MRI [29] and CT data [30] for noise reduction. We extend and apply this approach in the projection domain and evaluate its efficacy.

II. METHODS

The acquired set of 2D projections is first stacked to a 3D volume. With the third dimension being the projection angle, the data set can be viewed as a 3D sinogram involving 2D X-ray projections that is processed by adaptive anisotropic filtering prior to tomographic image reconstruction.

A. ADAPTIVE ANISOTROPIC FILTERING USING THE STRUCTURE TENSOR

130   The adaptive anisotropic filtering used in this study consists of the main steps depicted in Figure 2. The input of the algorithm is a 3D volume. First, the local orientation of structures in the input data is estimated using a set of six differently oriented filters (cf. Eq. 4). This results in an orientation estimate tensor for each voxel. After relaxation of the orientation estimate has been performed using a low-pass filter (cf. Eq. A.5 in the Appendix), the relaxed tensor is used to control the amount and directionality of the high-pass filtering (Eq. 5&6). In addition to the anisotropic high-pass component, an isotropic low-pass component of the original

135   data is also included (Eq. A.6). A mapping function (cf. Eq. 7) is used to combine the high-pass and low-pass components, based on the amount of structure detected at every location.
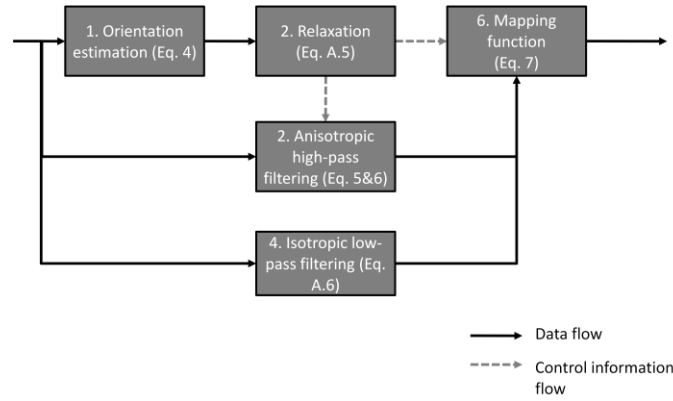
140   Figure 2 Block diagram showing the basic principle of the adaptive anisotropic filtering. The input is the stack of acquired projection data and the output can be reconstructed using any desired CT reconstruction algorithm.

By algorithm design, the sum of the low-pass and high-pass contributions will preserve all spatial frequencies orthogonal to estimated structures. In directions parallel to structures and in regions with no apparent structure, only the low frequency

145   information is maintained, resulting in a suppression of noise. These steps are described in more detail below. After adaptive filtering, the projection data are used for 3D image reconstruction using any tomographic image reconstruction algorithm.

1. Orientation estimation

150   The orientation estimation is based on a local estimate of the directionality in the image. This kind of orientation analysis tensor is often referred to as structure tensor in the literature [12]. In principle all structure tensors $\boldsymbol{T}^A$ are of the form $\boldsymbol{T}^A = A\hat{\boldsymbol{x}}\hat{\boldsymbol{x}}^\mathrm{T}$ where A is an arbitrary constant and $\hat{\boldsymbol{x}}$ is the vector pointing in the direction of the maximal signal variation [31].

We base the orientation estimation on quadrature filters. These filters are zero over one half of their frequency representation and their spatial representation is complex valued [31]. In order to obtain a description of the local orientation in every neighborhood in the input data, a set of six differently oriented 3D quadrature filters is applied [32][33]. We chose six over three directional filters as the angular sampling is improved with an increasing number of filters. The directions and angular dependence of the filters are chosen so that a rotationally invariant response is obtained, i.e. all possible orientations in 3D space are detected equally well. The radial and angular characteristics of these filters are described in detail in Appendix A. At every location $x \in \mathbb{R}^3$ of the volume $f(x) \in \mathbb{R}$, the responses from each of these quadrature filters, $q_k(x) \in \mathbb{R}$, are combined into a tensor representation, $T(x) \in \mathbb{R}^{3\times3}$, according to

$$T(x) = \sum_{k=1}^{6} q_k(x) \cdot M_k, \tag{4}$$

where

$$q_k(x) = \mathcal{F}^{-1}\{F_{\text{orient}}(u) \cdot \mathcal{F}\{f(x)\}\},$$

$$M_k = l\hat{n}_k\hat{n}_k^\top - mI,$$

$q_k(x)$ is the filter response of quadrature filter k, M_k is the orientation tensor in direction k, $\mathcal{F}$ and $\mathcal{F}^{-1}$ denote the Fourier transform and its inverse, $u \in \mathbb{R}^3$ the spatial frequency in Fourier domain, $\hat{n}_k \in \mathbb{R}^3$ is the direction of filter k, $I \in \mathbb{R}^3$ is the identity tensor, l is $\frac{4}{3}$ for 2D and $\frac{5}{4}$ for 3D and m is $\frac{1}{3}$ for 2D and $\frac{1}{4}$ for 3D [31]. $F_{\text{orient}}(u) \in \mathbb{R}$ is defined in the Appendix (cf. Eq. A.1).

2. Relaxation

Even in data containing sharp edges, the orientation of these structures are assumed to vary smoothly within the 3D sinogram space. To get a stable local orientation estimate of the tensor field obtained using Eq. 4, a subsequent low-pass filtering step or relaxation process is performed (based on the frequency representation in Eq. A.5). The result is denoted as $T_{\text{LP}}(x)$. The L_2-norm of this tensor, $|T_{\text{LP}}(x)|$, corresponds to the total amount of structure detected, while their eigenvalues and eigenvectors represent the degree and directionality of anisotropy.

3. Anisotropic high-pass filtering

Based on a set of six differently oriented high-pass (HP) filters, $F_{\text{HP}_k}(\boldsymbol{x})$ (based on the frequency representation in Eq. A.7), a new spatially dependent HP filter $F_{\text{HP}}(\boldsymbol{x})$ can be created with the desired characteristics.

185

$$F_{\text{HP}}(\boldsymbol{x}) = \sum_{k=1}^{6} \beta_k(\boldsymbol{x}) F_{\text{HP}_k}(\boldsymbol{x}) \tag{5}$$

The directionality and anisotropy are locally controlled by the relaxed orientation tensor estimate normalized by its largest eigenvalue λ_1, $\widehat{\boldsymbol{T}}_{\text{LP}}(\boldsymbol{x}) = \frac{\boldsymbol{T}_{\text{LP}}(\boldsymbol{x})}{\lambda_1}$, according to

190

$$\beta_k(\boldsymbol{x}) = \sum_{i=1}^{3} \sum_{j=1}^{3} \ \widehat{T}_{\text{LP}}^{i,j}(\boldsymbol{x}) \cdot M_k^{i,j}, \tag{6}$$

where $\widehat{T}_{\text{LP}}^{i,j}(\boldsymbol{x})$ and $M_k^{i,j}$ are the i-th and j-th element of the tensors $\widehat{\boldsymbol{T}}_{\text{LP}}(\boldsymbol{x})$ and $\boldsymbol{M}_k$ respectively.

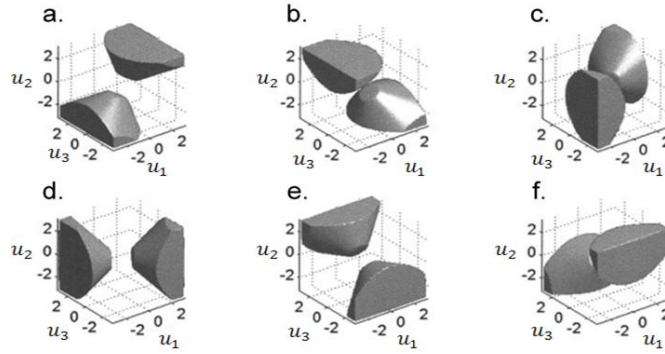The individual HP filters are shown in Figure 3.

195

Figure 3 The six differently oriented high pass filters (a-f) depicted as iso-surfaces in the Fourier domain. The responses from these filters are combined based on the low pass filtered and normalized orientation tensors in order to preserve only the high frequency components that are perpendicular to structures.

200

4. Isotropic low-pass filtering

An isotropic low-pass (LP) filter, $F_{LP}(\boldsymbol{u})$, is applied in order to obtain a uniformly smoothed data set that will have a reduced noise level compared to the original data (cf. Appendix). The low-pass filter is designed to maintain the local mean value in every
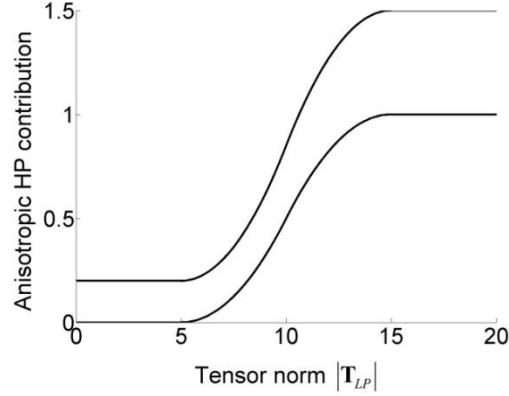
neighborhood.

205

Figure 4 Two examples of mapping functions that can be used to control the anisotropic high pass contribution based on the norm of the low-pass filtered orientation tensor. The upper curve represents a mapping that will amplify high spatial frequencies. When the lower curve is used, high spatial frequencies are merely preserved where structures are detected. When the orientation filter response $|\boldsymbol{T}_{\mathrm{LP}}(\boldsymbol{x})|$ is low, 20% or 0% of the high frequencies is maintained in these two examples respectively.

5. Mapping function

A mapping function, $\alpha(|\boldsymbol{T}_{\mathrm{LP}}(\boldsymbol{x})|)$ (cf. Eq. A.8), is used to control the amount of anisotropic high pass filtering that is added to the low-pass filtered data, based on the norm of the relaxed orientation data.

$$F_{\text{adaptive}}(\boldsymbol{x}) = F_{LP}(\boldsymbol{x}) + \ \alpha(|\boldsymbol{T}_{\mathrm{LP}}(\boldsymbol{x})|) \cdot F_{\mathrm{HP}}(\boldsymbol{x}) \tag{7}$$

Two examples of simple mapping functions are shown in Fig 3. Where the local orientation estimate indicates that a structure is present ($|\boldsymbol{T}_{\mathrm{LP}}(\boldsymbol{x})|$ is high), the mapping function will ensure that the high frequency content is preserved. If desired, the mapping function can also be designed so that the high frequencies are not only maintained, but amplified across edges (upper curve). This will result in an image that appears sharper than the original but maintains noise suppression in homogeneous regions.

6. Implementation

Implementation was done in Java using JTransforms for the fast Fourier transform on CPU and JCUDA for implementation on a graphics card. Computations were performed on a Linux server with four 2.0 GHz Intel Xeon cores. The server was also equipped with an Nvidia Tesla card with 240 CUDA-cores and 4 GB of graphics memory (cost of about $1200).

In order to avoid edge effects due to discontinuities caused by circular convolution, the data along all edges (at least 16 voxels) were mirrored to expand the 3D volume. These voxels were removed subsequent to filtering. In the current implementation, the

230 filtering is performed in the Fourier domain. A new filter is not explicitly computed for every location in the 3D data set. Instead, the input data are filtered using the LP and all six HP filters. The filter responses are subsequently combined according to the obtained control tensors.

Computation in this straight-forward manner on a standard PC may take up to several hours and may depending on the implementation require memory of up to 24 times the size of the actual projection data to be filtered depending on the accuracy of

235 the floating point operations. This is not feasible for an implementation on a graphics card. Even the latest cards only feature a memory of up to 4 GB. This is still too limited for a high resolution C-arm CT scan. Hence, we divided the problem into smaller blocks of equal size along the z-axis of the projection stack. Blocks were composed of 8, 16, or 32 projections before and after the center projection of the respective block.

240 B. ADAPTIVE ANISOTROPIC FILTERING USING THE BILATERAL FILTER

2-D Bilateral filtering was implemented as described in [34]. We chose to implement the filter without any additional modifications to decrease runtime as we were interested in maximal image quality. The adjustable parameters of the implementation were photometric and geometric standard deviation and the filter kernel size.

Table 1: Overview on the different acquisition protocols.

protocol	sweeps	frames	detector matrix	reconstruction volume	kVp	dose
#1 (simulation)	1 x 360°	800	512 x 512 1 x 1mm	512 x 512 0.5 x 0.5 mm	-	-
#2 (MTF, Catphan®)	1 x 197°	494	1240 x 960 0.31 x 0.31 mm	512 x 512 x 512 0.1 x 0.1 x 0.1 mm	109 kVp	1.20 μGy/pulse
#3 (high-dose, Catphan®)	1 x 198°	395	616 x 480 0.62 x 0.62 mm	512 x 512 x 512 0.41 x 0.41 x 0.41 mm	109 kVp	1.20 μGy/pulse
#4 (low-dose, Catphan®)	1 x 198°	395	616 x 480 0.62 x 0.62 mm	512 x 512 x 512 0.41 x 0.41 x 0.41 mm	109 kVp	0.36 μGy/pulse
#5 (CUDA, Catphan®)	1 x 191°	191	616 x 480 0.62 x 0.62 mm	512 x 512 x 512 0.41 x 0.41 x 0.41 mm	125 kVp	0.81 μGy/pulse
#6 (pig, cardiac)	4 x 198°	4 x 248	616 x 480 0.62 x 0.62 mm	512 x 512 x 512 0.41 x 0.41 x 0.41 mm	125 kVp	0.17 μGy/pulse
#7 (human, cardiac)	4 x 191°	4 x 191	616 x 480 0.62 x 0.62 mm	512 x 512 x 512 0.41 x 0.41 x 0.41 mm	90 kVp	0.54 μGy/pulse

245

C. EVALUATION

The evaluation of our proposed method covers the range from numerical simulation, to phantom experiments, and finally in-vivo applications.

1. Numerical Simulation

250 Simulated data sets of the phantoms depicted in Figure 5 were generated according to protocol #1 (cf. Table 1). Poisson distributed noise (cf. Eq. 2) corresponding to a low-dose exposure of 30000 photons per ray was added to the projection data. The resulting 3D stack comprising all 2D projections was filtered, but only the central slice was reconstructed using a 2D FBP algorithm. Noise levels were compared for both phantoms and the central bead in Phantom B was used to estimate the spatial resolution by calculating the modulation transfer function (MTF). A large number of profiles at different angles through the bead were extracted

255 by linear interpolation within the 2D plane and the Fourier transform was computed separately for each profile. The magnitude of each transformed profile was calculated to get the MTFs which were subsequently averaged to improve the estimate. The finite size of the bead (2x2 pixels) was compensated for by normalizing with the ideal MTF.
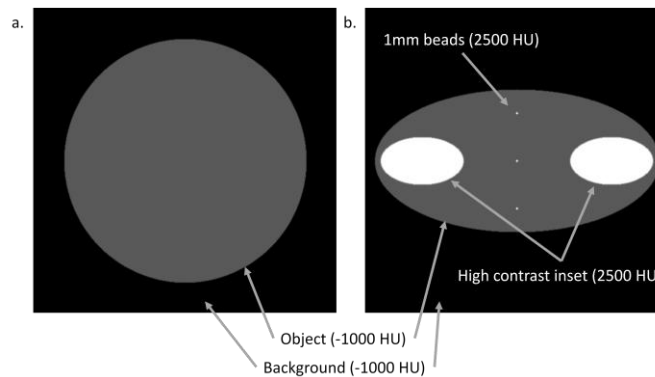

260

Figure 5 The two phantoms used for numerical simulation. Phantom A is a uniform cylinder of 1000 HU. Phantom B displays a more heterogeneous geometry with two bone-like high contrast insets and three small beads which are embedded in between.

2. Phantom Experiments

265 C-arm CT data were obtained using an AXIOM Artis dTA angiography system, software version VB30 (Siemens AG Medical Solutions, Forchheim, Germany). The original and filtered projection data were reconstructed using a Feldkamp cone beam reconstruction algorithm with additional correction algorithms for scatter, beam hardening, and truncation artifacts (DynaCT, Siemens AG Medical Solutions, Forchheim, Germany) [2].

For evaluation of spatial resolution, noise level and low contrast detectability, data were obtained using the Catphan® 600 phantom

270   similar to the results presented in [35]. In order to assess spatial resolution, one of the embedded beads was used to estimate the MTF. Acquisition was performed according to protocol #2 (cf. Table 1) with a CTDIw of 80 mGy. Data along lines passing through the bead at 3600 different angles were extracted using tri-linear interpolation to obtain the point spread function (PSF), which was subsequently Fourier transformed in order to estimate the MTF. The finite size of the bead (0.28 mm) was compensated for by normalizing with the ideal MTF. A weighted average of the MTFs from the different lines in 3D was calculated in order to assess

275   the spatial resolution along each axis separately. The angular weighting function used was $cos^2(\theta)$, where θ is the angle between each extracted line and the axis of interest.

Low contrast resolution was assessed in a detectability study. Twelve observers read four different images (high dose, high dose+filtering, low dose and low dose+filtering) obtained using the low-contrast module of the Catphan® phantom. Each observer read the images on the same monitor under the same controlled viewing conditions. Only targets at the 1.0% nominal contrast level

280   were scored, and observers were asked to determine the number of objects visible to them. For the low contrast detectability study, a high-dose scan was performed according to protocol #3 (CTDIw 60 mGy) and compared to a low-dose scan according to protocol #4 (CTDIw 18 mGy). The reconstruction kernel was a modified Shepp-Logan filter available as the system standard ("normal"). An axial slice through the low contrast section of the phantom was extracted. The display window center was adjusted such that its center value coincided with the mean background value, and the display window width was fixed at 100 HU.

285   Signal-difference-to-noise ratio (*SDNR*) measurements were performed on the same images using the largest 1.0% contrast target with an equally sized background region at the same radius in the phantom. *SDNR* was calculated as:

$$SDNR = \frac{S_a - S_b}{\sigma_b}$$
(8)

290   where S_a and S_b are the mean signal intensities in the object and background, respectively, and σ_b is the standard deviation in the background.

We also investigated the algorithm's runtime on phantom data. We processed a four-second C-arm acquisition according to protocol #5 in these experiments since the processing of the data set completely fits into the memory of the reconstruction server. In order to show that the decomposition of the problem into multiple smaller problems does not affect the filtering, we inspected the

295   outcome of the processing with respect to different block sizes. We investigated the MTF at a high frequency edge using an inset of the phantom and the standard deviation of three different homogeneous areas (at -830, 270 and 760 HU), since the algorithm's accuracy may be affected by numerical inaccuracies, a known potential limitation when using graphics hardware.

3. In-vivo Application

300   Cardiac C-arm CT data were obtained in a porcine model using the AXIOM Artis dTA angiography system (Siemens AG Medical Solutions, Forchheim, Germany). The following procedure was performed on a 41-kg female pig according to a protocol approved by Stanford University's Institutional Animal Care and Use Committee. First, the animal was sedated and intubated, placed on a ventilator and a 10-F sheath was inserted into the femoral vein. General anesthesia was maintained using isoflurane gas. Vitals were monitored throughout the procedure, and the cardiac rate was stable at ~93 bpm. A pigtail catheter was then placed proximal

305   to the femoral vein bifurcation. Data were acquired with contrast (282 mg/ml Iodine, Conray) diluted to 50% strength, an injection rate of 4.5 ml/s for a total volume of 180 ml. The contrast injection started 10 s prior to data acquisition to allow uniform opacification within the heart chambers. Four alternating forward and backward sweeps were performed over 198 degrees, with the start of each sweep synchronized to the ECG in order to ensure an adequate sampling of the projection angle – cardiac phase space [36]. Breathing was suspended throughout the image acquisition period. Acquisition was performed as listed in protocol #6. For

310   adaptive filtering, the complete set of $4\times248=992$ projections was processed as one single 3D data set (comprising $616\times480\times992$ voxels). In order to minimize edge effects due to circular convolution, the data along the edges were first mirrored extending the data to a size of $768\times512\times1024$ voxels. In order to reconstruct a specific cardiac phase after the adaptive filtering, nearest neighbor interpolation in the temporal domain was used to extract the projection data closest in time to a specific cardiac phase for all projection angles in both the original and filtered data sets.

315   Cardiac C-arm CT data were also obtained in a healthy volunteer during a single breath-hold using a similar imaging protocol. Intravenous injection of Visipaque 320 diluted to 50% was simultaneously performed into both the left (3.0 ml/s) and right (4.0ml/s) arms. Four alternating forward and backward sweeps were performed, each with a duration of four seconds, according to protocol #7. Adaptive filtering was performed using a mapping function that amplified the high spatial frequencies (see Figure 4). Subsequent to filtering of the projection data, multiple cardiac phases were retrospectively reconstructed based on the recorded

320   ECG timing information. Volume rendering was performed using the InSpace tab card on a Siemens Leonardo workstation (Siemens AG Medical Solutions, Forchheim, Germany). The study was approved by Stanford University's Institutional Review Board and informed patient consent was obtained.
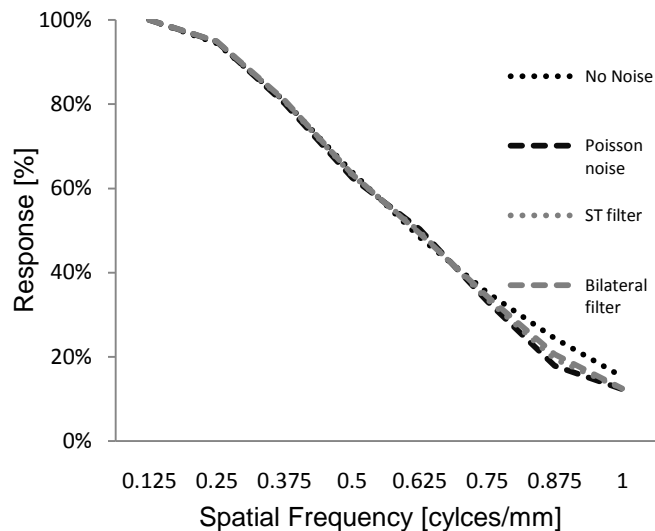
Figure 6 In-plane modulation transfer functions (MTFs) estimated in the numerical phantom. There is no significant different in the MTFs. There is no loss of
resolution in either of the methods.

III. RESULTS

A. NUMERICAL SIMULATION

In the numerical simulation, we chose to compare noiseless reconstruction, a reconstruction with simulated Poisson noise (cf. Section II.B.1), our proposed method based on anisotropic filtering in the projection domain, and a bilateral filter in reconstruction domain.

The in-plane MTFs calculated based on the central bead in Phantom B are presented in Figure 6, showing no degradation in spatial resolution after adaptive or bilateral filtering. Both filtering methods are equally edge-preserving.

For both numerical phantoms (A & B), the noise levels were measured in the homogenous regions indicated in Figure 1. Table 2 reports results for the comparison between reconstruction without noise, reconstruction with Poisson noise, our proposed method, and a bilateral filter applied in reconstruction domain. A standard deviation (SD) of 2.7 to 4.9 HU was still observed in the slice reconstructed from noise free projection data due to minor reconstruction artifacts. With noise added to the projection data, the SD in the reconstructed image was 51.9 to 65.8 HU. Note that the most noise is observed in the center as the attenuation is highest in this region. In order to perform a fair comparison between our method and the bilateral filter applied in the reconstruction domain, both filters were adjusted in such a manner that the off-center noise level is the same. As the quantitative results show, the structure tensor-based method is more efficient at the reduction of the noise in the center of the image.
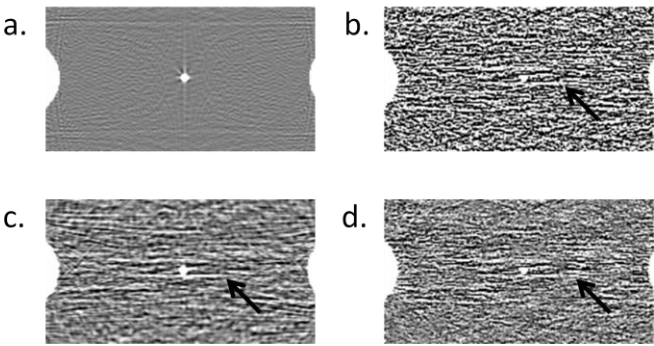
Figure 7 Detail view of Phantom B reconstructed based on noise free (a), noisy (b), structure tensor (ST) filtered projection data (c), and bilateral filtering in reconstruction domain (d). In the images reconstructed from the noisy data, horizontal streak artifacts can be seen. These structured artifacts are reduced in the ST filtered data. Bilateral filtering is not able to suppress these artifacts as well as the ST-based method. Some of the streaks are interpreted as structure and remain even in the filtered image at the position indicated by the arrow. (Window [-50, 50] HU)

Table 2: Results of the evaluation of homogenous regions. Image reconstruction in the given setting causes a noise of 2.7 to 5.0 HU. Structure tensor (ST) filter and bilateral filter were adjusted to give the same performance in the center region for better comparison.

Standard deviation [HU]	Phantom A		Phantom B	
	center	off-center	center	off-center
No noise	2.7	3.6	4.9	4.0
Poisson noise	58.3	51.9	63.6	41.1
ST projection Filtering	30.1	17.0	20.8	10.7
Bilateral reconstruction filtering	32.2	16.8	29.2	10.7

An examination of Figure 7 underscores the susceptibility of the bilateral filter to structured noise; the structure in the noise is misinterpreted as sensible structure and therefore not filtered. On the other hand, the structure tensor (ST)-based method is able to create much more uniform noise patterns, inhibiting the generation of noise-generated streaks by filtering before reconstruction.

B. PHANTOM EXPERIMENTS

The modulation transfer function estimated from the bead in the Catphan phantom showed almost no reduction in spatial resolution after adaptive filtering (see Figure 8). The spatial resolution at 10% MTF was 1.4, 1.4 and 1.5 cycles/mm in the x, y and z directions, respectively, both with and without filtering. The standard deviation in the homogeneous three-dimensional region

surrounding the bead had a standard deviation of 35 HU in the original data set, which was reduced to 15 HU after filtering. The non-filtered data set was also reconstructed using two alternative reconstruction kernels, instead of the "normal" kernel. By using a "smooth" kernel, the standard deviation was reduced to 13, but at the expense of a severe reduction in spatial resolution (10% MTF at 0.8, 0.8 and 0.9 cycles/mm in x, y and z). The "sharp" kernel preserves more of the high spatial frequencies (10% MTF at 1.7, 1.7 and 1.8 cycles/mm in x, y and z), but resulted in a standard deviation of 78 in the homogeneous area.
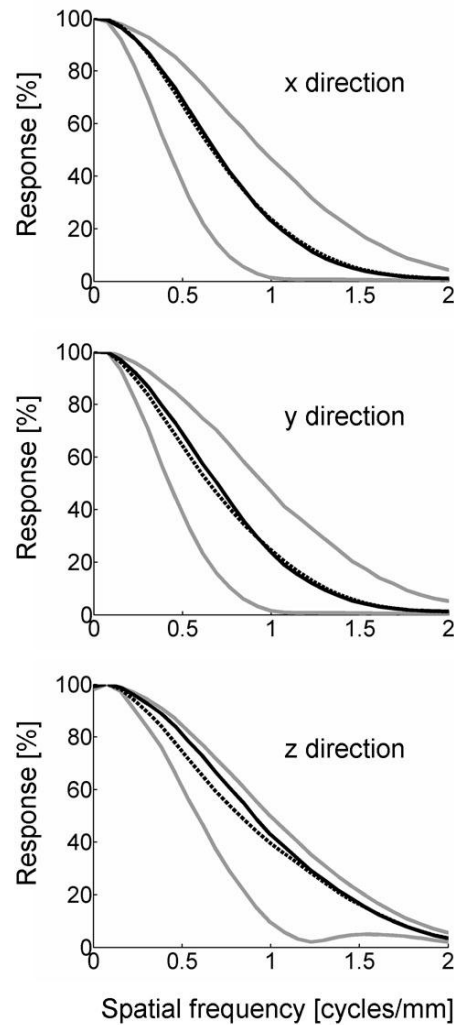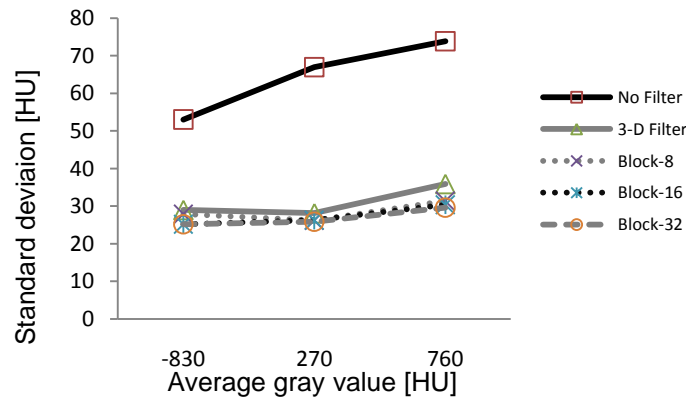

Figure 8 Modulation transfer functions estimated along each of the in-plane axes x and y and the axial direction z based on images of the bead embedded in the Catphan phantom. The solid black line is the MTF using the "normal" reconstruction kernel. The lower and higher gray lines correspond to the "smooth" and "sharp" kernels, while the dotted black lines are the MTFs using adaptive filtering and the "normal" kernel.
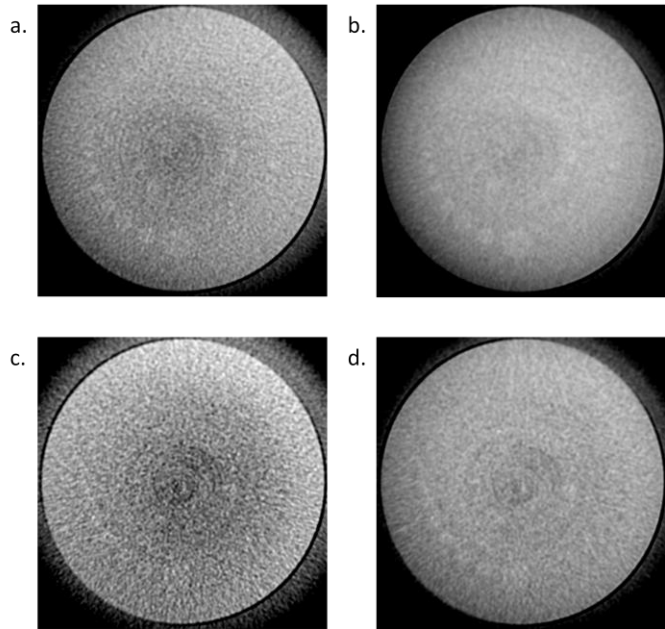
375   Figure 9 The low contrast module in the Catphan phantom imaged using a high (a) and low (c) dose. Images reconstructed using the same projection data subsequent to adaptive filtering are shown in (b) and (d).
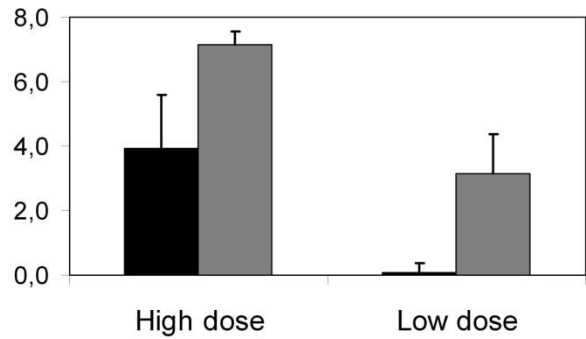
380   Figure 10   Number of objects detected by the observers in the images shown in the previous figure. The black and gray bars correspond to grading of the images reconstructed based on the original and filtered data, respectively. The error bars correspond to one standard deviation.

385 Figure 11   Noise levels (denoted as standard deviation in HU) are clearly reduced in the uniform areas of the image (measured at -830, 270, and 760 HU).
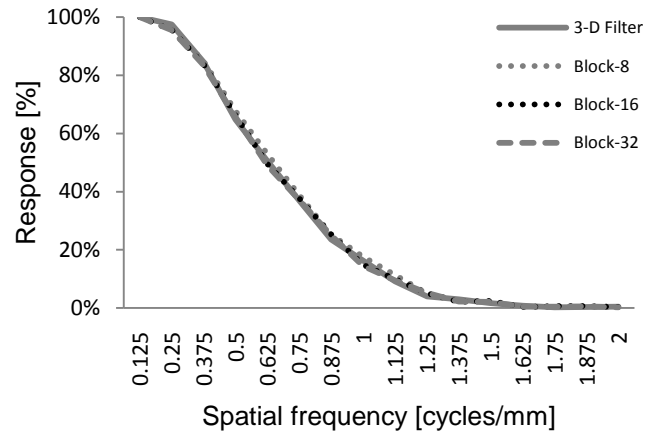

Figure 12   The modulation transfer function is in the same range for all problem subdivisions.

390   The low contrast detectability was significantly improved by applying adaptive filtering. Figure 9 shows the four images used for grading (high and low dose data, with and without adaptive filtering). The bar graph in Figure 10 illustrates the number of objects detected by the 12 observers. In the high dose data, an improvement from 3.9 ± 1.7 to 7.2 ± 0.4 detectable objects was achieved ($p<0.00001$) and from 0.1 ± 0.3 to 3.2 ± 1.2 in the low dose data ($p<0.00001$). The corresponding SDNR measurements in the same phantom images were 0.54 (high dose), 1.56 (high dose+filtering), 0.25 (low dose) and 0.76 (low dose+filtering); an improvement

395 by a factor of three after adaptive filtering.

Division of the filtering problem into different block sizes did not produce any differences with respect to noise and resolution. The noise in uniform areas is reduced drastically when compared to the reconstruction without any noise filtering (Figure 10). The MTFs remain in the same range as the MTF of the unfiltered reconstruction (Figure 11).

Since the filtering is independent of the size of processing blocks the speed-up by hardware acceleration on a graphics card was

400 investigated. Runtime evaluations were performed with a context size of 32 projections. The parallel four-core implementation on the server's CPU processed the data set in 1336 s. Computation of the Fourier transforms required in the adaptive noise filter on the graphics card provided a reasonable speed-up of 2.2 to a runtime of 601 s. With a complete implementation of the adaptive noise filter in CUDA, we further reduced the computation time to only 150 s. This corresponds to a 8.9-fold speed-up, and renders the approach clinically viable.

405

C. IN-VIVO APPLICATION

Cardiac C-arm CT images obtained in the porcine model are shown in Figure 13. Three-dimensional image volumes were reconstructed using projection data corresponding to a diastolic phase (80% of the cardiac cycle). By applying adaptive filtering to the projection data prior to retrospective cardiac gating and image reconstruction, the noise level was reduced. The gray value mean ± standard deviation in regions defined in the images in Figure 13a and Figure 13d were 229±58 and 230±32 in the left ventricular cavity (without and with filtering), and 64±62 and 62±36 in the ventricular septum. This corresponds to an improvement in the SDNR between blood and myocardium from 2.9 to 5.3.
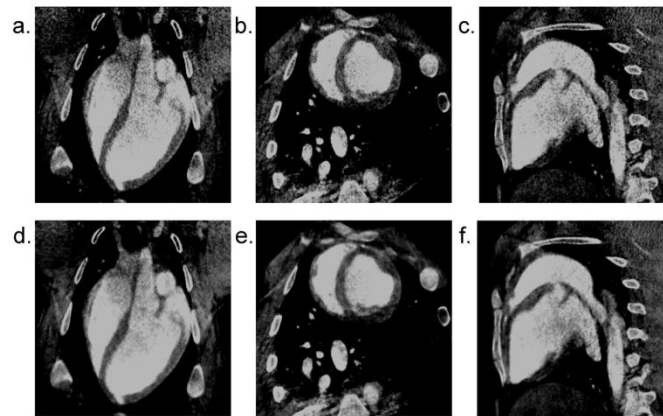

Figure 13   In-vivo data from the porcine study. The top row (a-c) shows three orthogonal views of the heart from a low dose cardiac C-arm CT acquisition. The images in the bottom row are reconstructed based on filtered projection data.
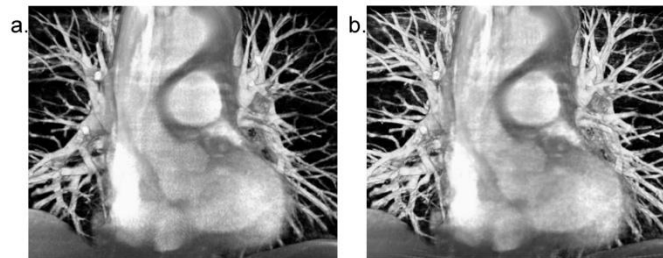

Figure 14 Human cardiac C-arm CT data without and with adaptive filtering. In the anisotropic filtering shown on the right, the high frequency contribution was amplified in order to improve the visualization of the smaller vessels

Increasing the high frequency content will sharpen the contours of small structures (see Figure 14). It does, however, require the adaptive control so that only regions with an apparent structure are enhanced. If the high frequency content was amplified equally in all regions, noise would be increased. The anisotropic nature of the filters also ensures that the high pass filtering only occurs across structures, while smoothing along structures reduces the noise level.

IV. DISCUSSION

Anisotropic adaptive filtering achieves a substantial reduction of noise while maintaining high spatial resolution. An advantage of filtering early in the reconstruction pipeline, i.e. in the raw data domain, is that noise induced streak artifacts are reduced. This was observed clearly in the numerical simulation. Since C-Arm data did not contain many high contrast insets, the noise induced streak

430   artifacts were not very prominent in the C-arm CT data, and so even at low dose, anisotropic filtering had little to improve upon. Yet a dramatic reduction of the noise level could still be seen in both in the phantom data and in-vivo examples.

In the detectability study, the reduction of noise due to anisotropic filtering improved detectability of low contrast objects so much that a better performance was obtained even when dose was reduced by a factor of three. In particular, with twelve observers, no significant difference could be seen between the *unfiltered high dose (HD)* and the *filtered low dose (LD)* data. The SDNR for the

435   filtered LD data (0.76) was even slightly higher than for the unfiltered HD data (0.54), while the latter required an X-ray dose that was more than three times higher. Note that this increase in detectability was achieved without loss of high-contrast resolution as would have been the case when using an isotropic low pass filter.

The 3D filtering approach used in this study takes advantage of the measured signal within every three-dimensional neighborhood in projection space. In this way, filters with a narrow radius can be used while still obtaining good statistical measures. In the case

440   of time-resolved 3D (4D) cardiac imaging, additional benefits could be achieved if the adaptive filtering was performed using all available dimensions. For example, if the projection angle, α, and cardiac phase, t, are considered to be two separate dimensions, the acquired projection data will be distributed in 4D (x,y,α,t). The sampling pattern will, however, not be equidistant along the cardiac time axis which will make the implementation more challenging. With the relatively low temporal resolution currently achieved in both conventional and C-arm cardiac CT, the correlation between projections acquired at different cardiac phases will

445   also be relatively low.

If the raw projection data is filtered too early in the reconstruction pipeline, inconsistencies propagate into the reconstruction. For example, a difference in gain and/or offset between adjacent detector elements will result in ring artifacts in a reconstructed CT image. When performing adaptive filtering on the projection data, it is essential to compensate for any inhomogeneities in the detector sensitivity prior to filtering. Otherwise, the gain differences in the three-dimensional sinogram domain might be further

450   enhanced by adaptive filtering. Another example of this effect is seen if filtering is applied before correction for the automatic exposure control function of the C-arm system. Hence computation of the filter after I_0 normalization is recommended.

As outlined in the introduction, many authors have presented other anisotropic filtering algorithms. Comparison of the performance of this broad variety of different anisotropic filtering algorithms, however, is difficult. Each of these algorithms has a unique parameterization which has to be carefully adjusted. Often parameters are dependent on the range of values of the input data

455   and are tweaked to optimize the image quality for the presented data sets. In our method, we alleviate this problem by

normalization of the tensors, which helps to separate the parameterization from the input data. However, a few parameters must still be chosen manually in order to optimize the image quality for a given task domain. To allow a fair comparison between the different anisotropic filtering methods a specific target domain and a representative collection of test cases must be defined. We tried to do this in the numerical study, but we are aware that even this comparison is not comprehensive. Eventually, fair

460   comparisons will only be possible if common implementations of algorithms are shared with the scientific community. For this reason, we are currently developing an open-source framework which will allow comparison of this algorithm and other algorithms, with appropriate evaluation methods and test cases.

Since the adaptive filtering is a shift-variant process it is likely that further improvements could be obtained by filtering the data both before and after image reconstruction. Noise will be reduced by analyzing neighborhoods in both the raw data and the image

465   domains independently. This possibility will also be further investigated.

V. CONCLUSION

Adaptive anisotropic filtering can reduce the noise level of an image without introducing noticeable blurring. As a consequence, this technique has the potential to substantially lower the radiation dose requirements for 2D X-ray projections used for 3D

470   tomographic image reconstruction. When radiation dose was within recommended levels, adaptive anisotropic filtering was found to significantly improve low contrast detectability. Furthermore, we have demonstrated that computationally expensive adaptive noise filters can also be computed within the time constraints of the interventional suite using hardware acceleration.

APPENDIX

This appendix gives a more detailed description of all filters used in the adaptive filtering algorithm. The quadrature filters

475    $F_{\text{orient}}(\boldsymbol{u})$ used for estimation of local orientation are spherically separable, i.e. they can be described by one radial function $R(\rho)$ and one directional function $D(\boldsymbol{u})$ as

$$F_{\text{orient}}(\boldsymbol{u}) = R(\rho) \cdot D(\boldsymbol{u}), \tag{A.1}$$

480   where $\boldsymbol{u}$ is the coordinate in Fourier space, and $\rho = |\boldsymbol{u}|$. The radial function used in this work is a lognormal function:

$$R(\rho) = e^{-\frac{4}{B^2 \ln(2)} \ln^2(\rho/\rho_i)} \tag{A.2}$$

$\rho_i = 1.5$ is the center frequency and $B = 2$ the relative bandwidth of the filters. The angular dependence of quadrature filter $D_k(\boldsymbol{u})$

485   is described by

$$D_k(\boldsymbol{u}) = \begin{cases} \boldsymbol{u} \cdot \hat{\boldsymbol{n}}_k & \text{if } \boldsymbol{u} \cdot \hat{\boldsymbol{n}}_k > 0 \\ 0 & \text{else} \end{cases} \tag{A.3}$$

The directions in 3D space of the six filters are

490

$$\hat{\boldsymbol{n}}_1 = c(a, 0, b)^{\mathsf{T}} \qquad \hat{\boldsymbol{n}}_2 = c(-a, 0, b)^{\mathsf{T}}$$

$$\hat{\boldsymbol{n}}_3 = c(b, a, 0)^{\mathsf{T}} \qquad \hat{\boldsymbol{n}}_4 = c(b, -a, 0)^{\mathsf{T}} \tag{A.4}$$

$$\hat{\boldsymbol{n}}_5 = c(0, b, a)^{\mathsf{T}} \qquad \hat{\boldsymbol{n}}_6 = c(0, b, -a)^{\mathsf{T}}$$

495   where $a = 2, b = 1 + \sqrt{5}$, and $c = \frac{1}{\sqrt{10 + 2\sqrt{5}}}$.

A Gaussian low pass filter $F_{\text{relax}}(\boldsymbol{x})$, with $\alpha = 1.0$ was used to regularize each tensor component of the orientation estimate.

$$F_{\text{relax}}(\boldsymbol{u}) = e^{-\alpha^2 |\boldsymbol{u}|^2 / 2} \tag{A.5}$$

500

The spatial frequency characteristics of the isotropic low-pass filter $F_{\text{LP}}(\boldsymbol{u})$ used in the final filtering is given by

$$F_{\text{LP}}(\boldsymbol{u}) = \begin{cases} \cos^2 \frac{\pi \rho}{2 \rho_{\text{LP}}} (\boldsymbol{u} \cdot \hat{\boldsymbol{n}}_k) & \text{if } 0 < \rho < \rho_{\text{LP}} \\ 0 & \text{else} \end{cases} \tag{A.6}$$

505

and for the anisotropic high pass filters by

$$F_{\text{HP}_k}(\boldsymbol{u}) = \left(1 - F_{LP}(\boldsymbol{u})\right)(\hat{\boldsymbol{n}}_k \cdot \boldsymbol{u})^2 \tag{A.7}$$

510   The value 1.5 was used as a cut-off frequency ρ_{LP} between the LP and HP filters and the directions $\hat{\boldsymbol{n}}_k$ are the same as in A.4.

The sigmoid function $\alpha(x)$ which is applied in the mapping function is given by

$$a(x) = \begin{cases} \hat{a}(x) & \text{if } \hat{a}(x) \geq 0 \\ 0 & \text{else} \end{cases} \tag{A.8}$$

$$\hat{a}(x) = \frac{l_{\text{output}} + h_{\text{output}}}{2} + \left(signum\left(\hat{x} - \frac{l_{\text{input}} + h_{\text{input}}}{2}\right) \cdot \left(1 - \frac{2 * \left|\left(\hat{x} - \frac{l_{\text{input}} + h_{\text{input}}}{2}\right)\right|}{h_{\text{input}} - l_{\text{input}}}\right)^{a_{\text{sm}}}\right) / 2$$

515

$$\hat{x} = \begin{cases} l_{\text{input}} & \text{if } x < l_{\text{input}} \\ x & \text{if } l_{\text{input}} < x < h_{\text{input}} \\ h_{\text{input}} & \text{if } h_{\text{input}} < x \end{cases}$$

where $a_{\text{sm}} = 2.0$ defines the slope of the sigmoid, l_{input} and h_{input} are the high and low input levels, and l_{output} and h_{output} are the high and low output levels.

ACKNOWLEDGMENTS

525   REFERENCES

[1] R. Fahrig, A. J. Fox, S. Lownie, and D. W. Holdsworth, "Use of a C-arm system to generate true three-dimensional computed rotational angiograms: preliminary in vitro and in vivo results," *AJNR Am J Neuroradiol*, vol. 18, pp. 1507-14, 1997.

[2] M. Zellerhoff, B. Scholz, E. P. Rührnschopf, and T. Brunner, "Low contrast 3-D reconstruction from C-arm data," in *Proceedings of the SPIE, Volume 5745*, San Diego, United States, 2005, pp. 646-55.

[3] G. Benndorf, C. M. Strother, B. Claus, R. Naeini, H. Morsi, R. Klucznik, and M. E. Mawad, "Angiographic CT in cerebrovascular stenting," *AJNR Am J Neuroradiol*, vol. 26, pp. 1813-8, 2005.

[4] M. K. Kalra, S. Rizzo, M. M. Maher, E. F. Halpern, T. L. Toth, J. A. Shepard, and S. L. Aquino, "Chest CT performed with

z-axis modulation: scanning protocol and radiation dose," *Radiology*, vol. 237, pp. 303-8, 2005.

[5] T. F. Jakobs, C. R. Becker, B. Ohnesorge, T. Flohr, C. Suess, U. J. Schoepf, and M. F. Reiser, "Multislice helical CT of the heart with retrospective ECG gating: reduction of radiation exposure by ECG-controlled tube current modulation," *Eur Radiol*, vol. 12, pp. 1081-6, 2002.

[6] J. Hausleiter, T. Meyer, F. Hermann, M. Hadamitzky, M. Krebs, T. C. Gerber, C. McCollough, S. Martinoff, A. Kastrati, A. Schömig, S. Achenbach, "Estimated Radiation Dose Associated With Cardiac CT Angiography," *JAMA*, vol. 301, no. 5, pp. 500-5007, 2009.

[7] T. Pflederer, L. Rudofsky, D. Ropers, S. Bachmann, M. Marwan, W. G. Daniel, and S. Achenbach, "Image Quality in a Low Radiation Exposure Protocol for Retrospectively ECG-Gated Coronary CT Angiography," *AJR*, vol. 192, pp. 1045-50, 2009.

[8] J. Hsieh, *Computed tomography: Principles, design, artifacts, and recent advances*.: SPIE Press, 2003.

[9] A. Borsdorf, S. Kappler, R. Raupach, and J. Hornegger, "Analytic Noise-Propagation in Indirect Fan-Beam FBP Reconstruction," in *30th Annual International IEEE EMBS Conference*, Vancouver, Canada, 2008, pp. 2701-4.

[10] T. Buzug, *Computed Tomography - From Photon Statistics to Modern Cone-Beam CT*. Berlin, Germany: Springer, 2008.

[11] L. Zhu, J. Wang, L. Xing, "Noise suppression in scatter correction for cone-beam CT," *Med Phys*, vol. 36, pp. 741-52, 2009.

[12] Joachim Weikert, *Anisotropic Diffusion in Image Processing*. Stuttgart, Germany: B.G. Teubner , 1998.

[13] P. Perona and J. Malik, "Scale space and edge detection using anisotropic diffusion," *IEEE Trans Pattern Anal Machine Intel*, vol. 12, pp. 629-39, 1990.

[14] J. Hsieh, "Adaptive streak artifact reduction in computed tomography resulting from excessive x-ray photon noise," *Med Phys*, vol. 25, pp. 2139-47, 1998.

[15] M. Kachelriess, O. Watzke, and W. A. Kalender, "Generalized multi-dimensional adaptive filtering for conventional and spiral single-slice, multi-slice, and cone-beam CT," *Med Phys*, vol. 28, pp. 475-90, 2001.

[16] R. Raupach, T. Flohr, J. Hornegger A. Borsdorf, "Wavelet based Noise Reduction in CT-Images Using Correlation Analysis," *IEEE Transactions on Medical Imaging*, vol. 27, no. 12, pp. 1685-1703, 2008.

[17] T. Li, X. Li, J. Wang, J. Wen, H. Lu, J. Hsieh, and Z. Liang, "Nonlinear sinogram smoothing for low-dose x-ray CT," *IEEE Transactions on Nuclear Science*, vol. 51, pp. 2505-13, 2004.

[18] J. Wang, T. Li, H. Lu, and Z. Liang, "Penalized weighted least-squares approach to sinogram noise reduction and image

reconstruction for low-dose X-ray computed tomography," *IEEE Trans Med Imaging*, vol. 25, pp. 1272-83, 2006.

[19] P. J. La Riviere, "Penalized-likelihood sinogram smoothing for low-dose CT," *Med Phys*, vol. 32, pp. 1676-83, 2005.

[20] J. Wang, T. Li, Z. Liang, L. Xing, "Dose reduction for kilovotage cone-beam computed tomography in radiation therapy," *Phys Med Biol*, vol. 53, pp. 2897-909, 2008.

[21] A. Gabinger, R. Weigel, S. Oeckl, P. Schmitt, "Enhancement of CT Image Quality via Bilateral Filtering of Projections," in *The First International Conference on Image Formation in X-Ray Computed Tomography*, Salt Lake City, United States, 2010, pp. 140-3.

[22] G. Gerig, O. Kubler, R. Kikinis, and F. Jolesz, "Nonlinear anisotropic filtering of MRI data," *IEEE Transactions on Medical Imaging*, vol. 11, pp. 221-32, 1992.

[23] G. Sapiro and A. Tannenbaum, "Affine invariant scale space," *Int J Comp Vision*, vol. 11, pp. 25-44, 1993.

[24] O. Demirkaya, "Reduction of noise and image artifacts in computed tomography by nonlinear filtration of the projection images," in *Proceedings of the SPIE, Volume 4322*, San Diego, United States, 2001, pp. 917-23.

[25] H. Knutsson, R. Wilson, and G. H. Granlund, "Anisotropic non-stationary image estimation and its applications - Part I: Restoration of noisy images," *IEEE Transactions on Communications*, vol. 31, pp. 388-97, 1983.

[26] H. Knutsson, L. Haglund, H. Bårman, and G. H. Granlund, "A framework for anisotropic adaptive filtering and analysis of image sequences and volumes," in *Internationcal Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, San Francisco, United States, 1992, pp. 469-72.

[27] H. Knutsson, "A tensor representation of 3D structures," in *5th IEEE-ASSP and EURASIP Workshop on Multidimensional Signal Processing*, Noordwijkerhout, The Netherlands, 1987, p. no pagination, available online at http://www.imt.liu.se/mi/Publications/pdfs/k87.pdf.

[28] L. Haglund, *Adaptive multidimensional filtering*. Linköping, Sweden: Dept. of Electrical Engineering, Linköpings universitet, 1992.

[29] C. F. Westin, L. Wigström, T. Loock, L. Sjöqvist, R. Kikinis, and H. Knutsson, "Three-dimensional adaptive filtering in magnetic resonance angiography," *J Magn Reson Imaging*, vol. 14, pp. 63-71, 2001.

[30] C. F. Westin, J. Richolt, V. Moharir, and R. Kikinis, "Affine adaptive filtering of CT data," *Med Image Anal*, vol. 4, pp. 161-77, 2000.

[31] C.-F. Westin, *A Tensor Framework for Multidimensional Signal Processing*, Linkoeping Studies in Science and Technology

ed. Linkoeping, Sweden: Linkoeping University, 1994.

[32] G. H. Granlund and H. Knutsson, *Signal processing for computer vision*. Dordrecht, The Netherlands: Kluwer Academic Publishers, 1995.

[33] H. Knutsson, H. Bårman, and L. Haglund, "Robust orientation estimation in 2D, 3D and 4D using tensors," in *2nd International Conference on Automation, Robotics and Computer Vision*, Singapore, Republic of Singapore, 1992, p. no pagination.

[34] Maduchi R Tomasi C, "Bilateral Filtering for Gray and Color Images," in *ICCV '98: Proceedings of the Sixth International Conference on Computer Vision*, Washington, DC, United States, 1998, pp. 839-46.

[35] R. Fahrig, R. Dixon, T. Payne, R. L. Morin, A. Ganguly, and N. Strobel, "Dose and image quality for a cone-beam C-arm CT system," *Med Phys*, vol. 33, pp. 4541-50, 2006.

[36] G. Lauritsch, J. Boese, L. Wigström, H. Kemeth, and R. Fahrig, "Towards cardiac C-arm computed tomography," *IEEE Trans Med Imaging*, vol. 25, pp. 922-34, 2006.

[37] J. Hausleiter, T. Meyer, M. Hadamitzky, E. Huber, M. Zankl, S. Martinoff, A. Kastrati, A.t Schömig, "Radiation Dose Estimates From Cardiac Multislice Computed Tomography in Daily Practice," *Circulation*, vol. 113, pp. 1305-1310, 2006.