Three-dimensional anisotropic adaptive filtering of projection data for noise reduction in cone beam CT

Andreas Maier^{a)}

Department of Radiology, Stanford University, Stanford, California 94305

Lars Wigström

Department of Radiology, Stanford University, Stanford, California 94305 and Center for Medical Image Science and Visualization, Linköping University, Linköping, Sweden

Hannes G. Hofmann and Joachim Hornegger

Pattern Recognition Laboratory, Department of Computer Science, Friedrich-Alexander University of Erlangen-Nuremberg, 91054, Erlangen, Germany

Lei Zhu

Nuclear and Radiological Engineering and Medical Physics Programs, The George W. Woodruff School of Mechanical Engineering, Georgia Institute of Technology, Atlanta, Georgia 30332

Norbert Strobel

Siemens AG Healthcare, Forchheim 91301, Germany

Rebecca Fahrig

Department of Radiology, Stanford University, Stanford, California 94305

(Received 7 March 2011; revised 9 August 2011; accepted for publication 12 August 2011)
published 14 October 2011)

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, Healthcare Sector, 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 to 30 HU). Virtually no degradation of high contrast spatial resolution was observed in the modulation transfer function (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 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.
© 2011 American Association of Physicists in Medicine. [DOI: 10.1118/1.3633901]

Key words: C-arm computed tomography (CT), low dose, noise, artifacts, image enhancement, hardware acceleration, GPU

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.^{1–3} 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.⁴⁻⁷ 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 multiplanar 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 contrast to other studies that investigate structure tensors,^{9,10} our filter implementation employs orientation filters that allow a higher directional sampling for the estimation of the orientation tensor, i.e., three, six or even more directions can be employed to extract the structure tensor. 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 were 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.

I.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.¹¹ The noise in the projection domain is mainly dependent on the number of detected photons in each detector element. A commonly used model for this process employs the Poisson distribution.¹²

$$P(n) = \frac{\kappa^n}{n!} \cdot e^{-\kappa}, \quad (1)$$

where $P(n)$ is the probability of n occurrences of an event and κ is the expected value of the distribution. More complex noise models¹³ describe polychromatic noise as a composition of multiple Poisson distributions. The overall distribution which also contains electronic noise, however, can be modeled well by the use of a single Poisson distribution.¹⁴

Given a random process which generates random numbers according to the Poisson distribution, the number of detected photons per detector element can be estimated by application of Lambert-Beer's Law

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

$$N = N_0 \cdot e^{-\int_0^s \mu(\eta) d\eta}, \quad (2)$$

where s the length of an x-ray beam through the object with attenuation $\mu(\eta)$ at ray position η .

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

$$SNR = \frac{\text{Signal}}{\text{Noise}} = \frac{\kappa}{\sqrt{\kappa}} = \sqrt{\kappa} = \sqrt{N}, \quad (3)$$

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. Fig. 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.¹⁵

I.B. Anisotropic and adaptive filtering in the literature

In the literature different uses of the term “anisotropic” are found. In the following the definition by Weikert¹⁶ will be used.

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

Fig. 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.

smoothing between detector channels when the signal level is low.¹⁷ Nonlinear filtering in the raw data domain was also investigated by Kachelrieß *et al.*¹⁸ 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.

Another interesting approach to adaptive filtering utilizing raw data was presented by Borsdorf *et al.*¹⁹ 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 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.^{20,21} A similar approach using a Poisson-based cost function is presented in Ref. 22. Although these approaches 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 Ref. 23, 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 Ref. 24. 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.²⁵ This work was later extended to 3D by Gerig *et al.*²⁶ More recent enhancement filters are based on diffusion derived from the affine heat equation.²⁷ 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.²⁸ 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.*^{29,30} A set of orientation-sensitive filters were used to estimate the local orientation of structures.³¹ 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.³² This kind of adaptive anisotropic filtering has previously been applied to reconstructed MRI Ref. 33 and CT data³⁴ 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.

II.A. Adaptive anisotropic filtering using the structure tensor

The adaptive anisotropic filtering used in this study consists of the main steps depicted in Fig. 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. (A5) in the Appendix], the relaxed tensor is used to control the amount and directionality of the high-pass filtering (Eqs. 5 and 6). In addition to the anisotropic high-pass component, an isotropic low-pass component of the original data is also included Eq. (A6). 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.

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 information is maintained, resulting in a

Fig. 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.

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.

II.A.1. Orientation estimation

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.¹⁶ In principle, all structure tensors $\mathbf{T}^A$ are of the form $\mathbf{T}^A = A\hat{\mathbf{x}}\hat{\mathbf{x}}^T$ where A is an arbitrary constant and $\hat{\mathbf{x}}$ is the vector pointing in the direction of the maximal signal variation.³⁵

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.³⁵ 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.^{36,37} We chose six over three directional filters as the angular sampling is improved with an increasing number of filters. As the filtered data in this paper are a stack of 2D projections, the same structure can be located under various angles throughout the projection stack. In theory, our approach could be extended to even more filters. However, we limited the number to six due to the computational effort. The directions and angular dependence of the filters were chosen to obtain a rotationally invariant response. As a result, all possible orientations in 3D space can be detected equally well (cf. Fig. 3). The radial and angular characteristics of these filters are described in detail in Appendix A. At every location $\mathbf{x} \in \mathbb{R}^3$ of the volume $f(\mathbf{x}) \in \mathbb{R}$, the responses from each of these quadrature filters, $q_k(\mathbf{x}) \in \mathbb{R}$, are combined into a tensor representation, $\mathbf{T}(\mathbf{x}) \in \mathbb{R}^{3 \times 3}$, according to

$$\mathbf{T}(\mathbf{x}) = \sum_{k=1}^6 q_k(\mathbf{x}) \cdot \mathbf{M}_k, \quad (4)$$

where

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

$$\mathbf{M}_k = l\hat{\mathbf{n}}_k\hat{\mathbf{n}}_k^T - m\mathbf{I},$$

$q_k(\mathbf{x})$ is the filter response of quadrature filter k , $\mathbf{M}_k$ is the orientation tensor in direction k , $\mathcal{F}$ and $\mathcal{F}^{-1}$ denote the Fourier transform and its inverse, $\mathbf{u} \in \mathbb{R}^3$ the spatial frequency in Fourier domain, $\hat{\mathbf{n}}_k \in \mathbb{R}^3$ is the direction of filter k , $\mathbf{I} \in \mathbb{R}^3$, 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.³⁵ $F_{\text{orient}}(\mathbf{u}) \in \mathbb{R}$ is the quadrature filter used for estimation of local orientation, which is defined in the Appendix [cf. Eq. (A1)].

II.A.2. Relaxation

Even in data containing sharp edges, the orientations 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 provided in Eq. (A5) in the Appendix). The result is denoted as $\mathbf{T}_{\text{LP}}(\mathbf{x})$. The L_2 -norm of this tensor, $|\mathbf{T}_{\text{LP}}(\mathbf{x})|$, corresponds to the total amount of structure detected, while their eigenvalues and eigenvectors represent the degree and directionality of anisotropy.

II.A.3. Anisotropic high-pass filtering

Based on a set of six differently oriented high-pass (HP) filters, $F_{\text{HP}_k}(\mathbf{x})$ [based on the frequency representation in Eq. (A7)], a new spatially dependent HP filter $F_{\text{HP}}(\mathbf{x})$ can be created with the desired characteristics.

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

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

$$\beta_k(\mathbf{x}) = \sum_{i=1}^3 \sum_{j=1}^3 \hat{T}_{\text{LP}}^{ij}(\mathbf{x}) \cdot M_k^{ij}, \quad (6)$$

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

The individual HP filters are shown in Fig. 3.

FIG. 3. The six differently oriented high-pass filters (a–f) depicted as isosurfaces 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.

II.A.4. Isotropic low-pass filtering

An isotropic low-pass (LP) filter, $F_{LP}(u)$, 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.

II.A.5. Mapping function

A mapping function, $\alpha(|T_{LP}(x)|)$ [cf. Eq. (A8)], 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.

$$F_{\text{adaptive}}(x) = F_{LP}(x) + \alpha(|T_{LP}(x)|) \cdot F_{HP}(x) \quad (7)$$

Two examples of simple mapping functions are shown in Fig. 4. Where the local orientation estimate indicates that a structure is present ($|T_{LP}(x)|$ 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.

II.A.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

FIG. 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 $|T_{LP}(x)|$ is low, 20 or 0% of the high frequencies are maintained in these two examples, respectively.

were removed subsequent to filtering. In the current implementation, the 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 requires memory of up to 24 times the size of the actual projection data to be filtered depending on the accuracy of 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.

II.B. Adaptive anisotropic filtering using the bilateral filter

Two-dimensional bilateral filtering was implemented as described in Ref. 38. 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.

II.C. Evaluation

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

II.C.1. Numerical simulation

Simulated data sets of the phantoms depicted in Fig. 5 were generated according to protocol #1 (cf. Table I). Poisson distributed noise [cf. Eq. (2)] corresponding to a low dose exposure of 30,000 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 spatial resolution was calculated using the modulation transfer function (MTF) of the central and top off-center beads in Phantom B. A large number of profiles at different angles through the beads were extracted 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 beads (2×2 pixels) was compensated for by normalizing with the ideal MTF.

In addition, we investigated our method using the well known Forbild head section. All simulations were conducted in a low dose scenario using 30,000 photons and a scan setup according to protocol #1.

FIG. 5. The two phantoms used for numerical simulation. Phantom A is a uniform cylinder of 0 HU. Phantom B displays a more heterogeneous geometry with two bone-like high contrast insets between which are imbedded three small beads.

II.C.2. Phantom experiments

C-arm CT data were obtained using an AXIOM Artis dTA angiography system (Siemens AG, Healthcare Sector, Forchheim, Germany, software version VB30). 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, Healthcare Sector).²

For evaluation of spatial resolution, noise level, and low contrast detectability, data were obtained using the Catphan[®] 600 phantom similar to the results presented in Ref. 39. 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 I) with a CTDI_w of 80 mGy. Data along lines passing through the bead at 3600 different angles were extracted using trilinear 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 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 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 (CTDI_w 60 mGy) and compared to a low dose scan according to protocol #4 (CTDI_w 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. 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:

TABLE I. Overview on the different acquisition protocols.

Protocol	Sweeps	Frames	Detector matrix	Reconstruction volume	kVp	Dose
#1 (simulation)	1 × 360°	800	512 × 512 1 × 1 mm	512 × 512 0.5 × 0.5 mm	—	—
#2 (MTF, Catphan [®])	1 × 197°	494	1240 × 960 0.31 × 0.31 mm	512 × 512 × 512 0.1 × 0.1 × 0.1 mm	109 kVp	1.20 μGy/pulse
#3 (high-dose, Catphan [®])	1 × 198°	395	616 × 480 0.62 × 0.62 mm	512 × 512 × 512 0.41 × 0.41 × 0.41 mm	109 kVp	1.20 μGy/pulse
#4 (low-dose, Catphan [®])	1 × 198°	395	616 × 480 0.62 × 0.62 mm	512 × 512 × 512 0.41 × 0.41 × 0.41 mm	109 kVp	0.36 μGy/pulse
#5 (CUDA, Catphan [®])	1 × 191°	191	616 × 480 0.62 × 0.62 mm	512 × 512 × 512 0.41 × 0.41 × 0.41 mm	125 kVp	0.81 μGy/pulse
#6 (pig, cardiac)	4 × 198°	4 × 248	616 × 480 0.62 × 0.62 mm	512 × 512 × 512 0.41 × 0.41 × 0.41 mm	125 kVp	0.17 μGy/pulse
#7 (human, cardiac)	4 × 191°	4 × 191	616 × 480 0.62 × 0.62 mm	512 × 512 × 512 0.41 × 0.41 × 0.41 mm	90 kVp	0.54 μGy/pulse

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

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 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.

II.C.3. In-vivo application

Cardiac C-arm CT data were obtained in a porcine model using the AXIOM Artis dTA angiography system (Siemens AG, Healthcare Sector). 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 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° , with the start of each sweep synchronized to the ECG in order to ensure an adequate sampling of the projection angle – cardiac phase space.⁴⁰ Breathing was suspended throughout the image acquisition period. Acquisition was performed as listed in protocol #6.

For adaptive filtering, the complete set of $4 \times 248 = 992$ projections was processed as one single 3D data set (comprising $616 \times 480 \times 992$ 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 $768 \times 512 \times 1024$ 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.

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 (4.0 ml/s) arms. Four alternating forward and backward sweeps were performed, each with a duration of 4 s, according to protocol #7. Adaptive filtering was performed using a mapping function that amplified the high spatial frequencies (see Fig. 4). Subsequent to filtering of the projection data, multiple cardiac phases were retrospectively reconstructed based on the recorded ECG timing information. Volume rendering was performed using the InSpace tab card on a Siemens Leonardo workstation (Siemens AG, Healthcare Sector). The study was approved by Stanford University's Institutional Review Board and informed patient consent was obtained.

III. RESULTS

III.A. Numerical simulation

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

The in-plane MTFs calculated based on the central and the top off-center beads in Phantom B are presented in

Fig. 6. In-plane modulation transfer functions (MTFs) estimated in the numerical phantom. The left plot shows the MTFs at the center bead while the right plot shows the MTFs at the top off-center bead. There is no significant difference in the MTFs. There is no loss of resolution in either of the methods.

TABLE II. Results of the evaluation of homogeneous regions. Image reconstruction in the given setting causes a noise of 2.7 to 4.9 HU. The parameters of the structure tensor (ST) filter and bilateral filter were adjusted to give the same performance in the off-center region of the image 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

Fig. 6. They show no degradation in spatial resolution after adaptive or bilateral filtering. Both filtering methods are equally edge-preserving.

For both numerical phantoms (A and B), the noise levels were measured in the homogeneous regions indicated in Fig. 1. Table II reports results for the comparison between reconstruction without noise, reconstruction with Poisson noise, our proposed method, and a bilateral filter applied in the reconstruction domain. A standard deviation (SD) of 2.7–4.9 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 51.9–65.8 HU. Note that the maximum 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.

An examination of Fig. 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.

Examination of the Forbild head section showed again that the noise can be reduced drastically using the ST filter. Figure 8 shows the phantom under different conditions. The top left shows the phantom reconstruction without any additional noise. Due to the high contrast objects in the phantom, SD of 19.4 HU is measured in the homogeneous area in the center of the phantom and 19.3 HU in the lower left homogeneous area. After the simulation of a low dose scan, the noise in the center increases to 85.6 HU and the off-center noise to 77.5 HU. Application of the ST-filter reduces the noise in the center to 64.8 HU and to 41.2 HU in the off-center region. Interestingly, the streak artifacts due to undersampling and view aliasing that are visible in the noise-free (“no noise”) image are obscured by noise in the noisy simulation, but reappear in the ST-filtered image. A plot through the high contrast ear pattern shows that the edges are preserved in all three versions.

III.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 Fig. 9). 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 nonfiltered 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 HU, 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 HU in the homogeneous area.

FIG. 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 the 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 in the filtered image, e.g., at the position indicated by the arrow (Window [−50, 50] HU).

FIG. 8. For the Forbield phantom, the noise in the reconstruction domain is significantly reduced (Window $[-50, 150]$ HU). As shown by the profiles through the high contrast pattern of the ear on the right side of the numerical phantom, the spatial resolution is also preserved.

FIG. 9. 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.

The low contrast detectability was significantly improved by applying adaptive filtering. Figure 10 shows the four images (high and low dose data, with and without adaptive filtering) used for grading in a randomized trial. The bar graph in Fig. 11 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), demonstrating an improvement 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 (Fig. 12). The MTFs remain in the same range as the MTF of the unfiltered reconstruction (Fig. 13).

Since the filtering is independent of the size of processing blocks the speed-up by hardware acceleration on a graphics card was 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 1337s. 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 150s. This corresponds to an 8.9-fold speed-up, and renders the approach clinically viable.

III.C. In-vivo application

Cardiac C-arm CT images obtained in the porcine model are shown in Fig. 14. Three-dimensional image volumes were reconstructed using projection data corresponding to a diastolic phase (80% of the cardiac cycle). By applying

FIG. 10. 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).

adaptive filtering to the projection data prior to retrospective cardiac gating and image reconstruction, the noise level was reduced. The gray value mean $\pm$ standard deviation in regions defined in the images in Figs. 14(a) and 14(d) 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, a factor of 1.8 (Fig. 15).

Increasing the high frequency content will sharpen the contours of small structures (see Fig. 15). It does, however, require the adaptive control so that only regions with an apparent structure are enhanced. If the high frequency con-

tent 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. Interestingly, the view aliasing that is visible in the

FIG. 11. 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.

FIG. 12. Noise levels in the Catphan[®] phantom (denoted as standard deviation in HU) are clearly reduced in the uniform areas of the image (measured at -830 , 270 , and 760 HU).

FIG. 13. The modulation transfer function of the Catphan[®] phantom is in the same range for all problem subdivisions.

noise-free reconstruction reappeared in the reconstruction after filtering in projection domain. As the view aliasing is caused by the reconstruction process itself, it cannot be reduced by filtering in the projection domain. The effect, however, was only observed in the simulated data, as view aliasing was not a predominant source of artifact in the real (low dose) C-arm data. In addition, since the C-Arm data did not contain many high contrast insets, the noise-induced streak artifacts were not very prominent, and so even at low dose, anisotropic filtering had little to improve upon. Because view aliasing produces structure-like streaks, we expect postreconstruction filters will not reduce this type of artifact. Yet a dramatic reduction of the noise level could still be seen in both the phantom data and the in-vivo examples.

In the detectability study, the reduction of noise due to anisotropic filtering improved detectability of low contrast objects such 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 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 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 also be relatively low.

If the raw projection data are 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 enhanced by adaptive filtering. Another example of this effect is seen if filtering is applied before correction for the

FIG. 14. 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 without filtering. The images in the bottom row are reconstructed based on ST-filtered projection data.

FIG. 15. Human cardiac C-arm CT data without (a) and with adaptive filtering (b). In the anisotropic filtering shown in the middle (b), the high frequency contribution was amplified in order to improve the visualization of the smaller vessels. The difference of both is depicted on the right (c).

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 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 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 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 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. Since the third dimension (angle) may be sampled differently than the first two dimensions (pixel distances in the horizontal and vertical directions), we chose to define all filters with respect to π as the maximal frequency that may occur in the respective dimension, i.e., the frequency axis is rescaled from $-\pi$ to π in each dimension. The quadrature filters $F_{\text{orient}}(\mathbf{u})$ 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 $D(\mathbf{u})$ as

$$F_{\text{orient}}(\mathbf{u}) = R(\rho) \cdot D(\mathbf{u}), \quad (\text{A1})$$

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

$$R(\rho) = e^{-\frac{4}{B^2 \ln(2)} \ln^2(\rho/\rho_i)} \quad (\text{A2})$$

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

$$D_k(\mathbf{u}) = \max(\mathbf{u} \cdot \hat{\mathbf{n}}_k, 0) \quad (\text{A3})$$

The directions in 3D space of the six filters are

$$\begin{aligned} \hat{\mathbf{n}}_1 &= c(a, 0, b)^T & \hat{\mathbf{n}}_2 &= c(-a, 0, b)^T \\ \hat{\mathbf{n}}_3 &= c(b, a, 0)^T & \hat{\mathbf{n}}_4 &= c(b, -a, 0)^T \\ \hat{\mathbf{n}}_5 &= c(0, b, a)^T & \hat{\mathbf{n}}_6 &= c(0, b, -a)^T \end{aligned} \quad (\text{A4})$$

where

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

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

$$F_{\text{relax}}(\mathbf{u}) = e^{-\alpha^2 |\mathbf{u}|^2 / 2} \quad (\text{A5})$$

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

$$F_{\text{LP}}(\mathbf{u}) = \begin{cases} \cos^2 \frac{\pi \rho}{2 \rho_{\text{LP}}} (\mathbf{u} \cdot \hat{\mathbf{n}}_k) & \text{if } 0 < \rho < \rho_{\text{LP}} \\ 0 & \text{otherwise} \end{cases} \quad (\text{A6})$$

and for the anisotropic high-pass filters by

$$F_{\text{HP}_k}(\mathbf{u}) = (1 - F_{\text{LP}})(\hat{\mathbf{n}}_k \cdot \mathbf{u})^2 \quad (\text{A7})$$

The value 1.5 was used as a cut-off frequency ρ_{LP} between the LP and HP filters and the directions $\hat{n}_k$ are the same as in Eq. (A4).

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

$$\begin{aligned}\alpha(x) &= \max(\hat{\alpha}(x), 0) \\ \hat{\alpha}(x) &= 0.5(l_{\text{output}} + h_{\text{output}} + \text{sign}(2\hat{x} - l_{\text{input}} - h_{\text{input}})) \\ &\quad \cdot \left(1 + \left(\frac{|2\hat{x} - l_{\text{input}} - h_{\text{input}}|}{l_{\text{input}} - h_{\text{input}}}\right)^{a_{\text{sm}}}\right) \\ \hat{x} &= \begin{cases} l_{\text{input}} & \text{if } x < l_{\text{input}} \\ x & \text{if } l_{\text{input}} < x < h_{\text{input}} \\ h_{\text{input}} & \text{if } h_{\text{input}} < x \end{cases} \quad (\text{A8})\end{aligned}$$

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

We thank Wendy Baumgardner and Dr. Amin Al-Ahmad at Stanford University for their assistance with the animal and patient studies and Teri Moore, Jan Boese and Günter Lauritsch at Siemens AG, Healthcare Sector for all their help and support. This work was supported in part by NIH under grants R01EB003524 and R01HL087917, Siemens AG, Healthcare Sector, AX Division, the Lucas Foundation, and the Swedish Heart and Lung Foundation.

^{a)} Author to whom correspondence should be addressed. Electronic mail: akmaier@stanford.edu

¹R. 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.*, **18**, pp. 1507–1514 (1997).

²M. Zellerhoff, B. Scholz, E. P. Rührnschopf, and T. Brunner, "Low contrast 3-D reconstruction from C-arm data," *Proc. SPIE*, **5745**, pp. 646–655 (2005).

³G. 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.*, **26**, pp. 1813–1818 (2005).

⁴M. 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*, **237**, pp. 303–308 (2005).

⁵T. 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.*, **12**, pp. 1081–1086 (2002).

⁶J. 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, J. Am. Med. Assoc.*, **301**(5), pp. 500–5007 (2009).

⁷T. 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, Am. J. Roentgenol.*, **192**, pp. 1045–1050 (2009).

⁸J. Hsieh, *Computed tomography: Principles, design, artifacts, and recent advances* (SPIE Press, Bellingham, Washington, United States, 2003).

⁹J. Weickert, H. Schar, "A Scheme for Coherence-Enhancing Diffusion Filtering with Optimized Rotation Invariance," *J. Visual Commun. Image Represent.*, **13**(1), pp. 103–118 (2002).

¹⁰A. Mendrik, E. Vonken, A. Rutten, M. Viergever, B. van Ginneken, "Noise reduction in computed tomography scans using 3D anisotropic hybrid diffusion with continuous switch," *IEEE Trans. Med. Imaging*, **28**(10), pp. 1585–1594 (2009).

¹¹A. Borsdorf, S. Kappler, R. Raupach, and J. Hornegger, "Analytic Noise-Propagation in Indirect Fan-Beam FBP Reconstruction," *30th Annual International IEEE EMBS Conference* (Vancouver, Canada, 2008), pp. 2701–2704.

¹²T. Buzug, *Computed Tomography—From Photon Statistics to Modern Cone-Beam CT* (Springer, Berlin, Germany, 2008).

¹³P. Massoumzadeh, O. A. Earl B. R. Whiting, "Properties of preprocessed sinogram data in x-ray computed tomography," *Med. Phys.*, **33**(9), pp. 3290–3303 (2006).

¹⁴J. Bian, P. A. Vargas P. J. La Rivière, "Penalized-likelihood sinogram restoration for computed tomography," *IEEE Trans. Med. Imaging*, **25**, pp. 1022–1036 (2006).

¹⁵L. Zhu, J. Wang, L. Xing, "Noise suppression in scatter correction for cone-beam CT," *Med. Phys.*, **36**, pp. 741–752 (2009).

¹⁶Joachim Weikert, *Anisotropic Diffusion in Image Processing* (B.G. Teubner, Stuttgart, Germany, 1998).

¹⁷M. 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.*, **28**, pp. 475–490 (2001).

¹⁸R. Raupach, T. Flohr, J. Hornegger A. Borsdorf, "Wavelet based Noise Reduction in CT-Images Using Correlation Analysis," *IEEE Trans. Med. Imaging*, **27**(12), pp. 1685–1703 (2008).

¹⁹T. Li, X. Li, J. Wang, J. Wen, H. Lu, J. Hsieh, and Z. Liang, "Nonlinear sinogram smoothing for low-dose x-ray CT," *IEEE Trans. Nucl. Sci.*, **51**, pp. 2505–2513 (2004).

²⁰J. 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*, **25**, pp. 1272–1283 (2006).

²¹P. J. La Rivière, "Penalized-likelihood sinogram smoothing for low-dose CT," *Med. Phys.*, **32**, pp. 1676–1683 (2005).

²²J. Wang, T. Li, Z. Liang, L. Xing, "Dose reduction for kilovoltage cone-beam computed tomography in radiation therapy," *Phys. Med. Biol.*, **53**, pp. 2897–2909 (2008).

²³A. 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* (The University of Utah, Salt Lake City, 2010), pp. 140–143.

²⁴P. Perona and J. Malik, "Scale space and edge detection using anisotropic diffusion," *IEEE Trans. Pattern Anal. Mach. Intell.*, **12**, pp. 629–639 (1990).

²⁵J. Hsieh, "Adaptive streak artifact reduction in computed tomography resulting from excessive x-ray photon noise," *Med. Phys.*, **25**, pp. 2139–2147 (1998).

²⁶G. Gerig, O. Kubler, R. Kikinis, and F. Jolesz, "Nonlinear anisotropic filtering of MRI data," *IEEE Trans. Med. Imaging*, **11**, pp. 221–232 (1992).

²⁷G. Sapiro and A. Tannenbaum, "Affine invariant scale space," *Int J Comput. Vis.*, **11**, pp. 25–44 (1993).

²⁸O. Demirkaya, "Reduction of noise and image artifacts in computed tomography by nonlinear filtration of the projection images," *Proc. SPIE*, **4322**, pp. 917–923 (2001).

²⁹H. Knutsson, R. Wilson, and G. H. Granlund, "Anisotropic non-stationary image estimation and its applications - Part I: Restoration of noisy images," *IEEE Trans. Commun.*, **31**, pp. 388–397 (1983).

³⁰H. Knutsson, L. Haglund, H. Bärman, and G. H. Granlund, "A framework for anisotropic adaptive filtering and analysis of image sequences and volumes," *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)* (San Francisco, 1992), pp. 469–472.

³¹H. Knutsson, "A tensor representation of 3D structures," *5th IEEE-ASSP and EURASIP Workshop on Multidimensional Signal Processing* (Noordwijkerhout, The Netherlands, 1987), available online at https://www.imt.liu.se/mi_old/Publications/pdfs/k87.pdf.

³²L. Haglund, *Adaptive multidimensional filtering* (Department of Electrical Engineering, Linköpings universitet, Linköping, Sweden, 1992).

³³C. 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*, **14**, pp. 63–71 (2001).

³⁴C. F. Westin, J. Richolt, V. Moharir, and R. Kikinis, "Affine adaptive filtering of CT data," *Med. Image Anal.*, **4**, pp. 161–177 (2000).

³⁵C.-F. Westin, *A Tensor Framework for Multidimensional Signal Processing*, in *Linköping Studies in Science and Technology* (Linköping University, Linköping, Sweden, 1994).

- ³⁶G. H. Granlund and H. Knutsson, *Signal Processing for Computer Vision*. (Kluwer Academic Publishers, Dordrecht, The Netherlands, 1995).
- ³⁷H. Knutsson, H. Bårman, and L. Haglund, "Robust orientation estimation in 2D, 3D and 4D using tensors," *2nd International Conference on Automation, Robotics and Computer Vision* (Singapore, Republic of Singapore, 1992).
- ³⁸Maduchi R Tomasi C, "Bilateral Filtering for Gray and Color Images," *ICCV'98: Proceedings of the Sixth International Conference on Computer Vision ((ICCV-98)*, (Bombay, India, January 4-7, 2008. Narosa Publishing House, Washington, DC, 1998), pp. 839–846.
- ³⁹R. 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.*, **33**, pp. 4541–4550 (2006).
- ⁴⁰G. Lauritsch, J. Boese, L. Wigström, H. Kemeth, and R. Fahrig, "Towards cardiac C-arm computed tomography," *IEEE Trans. Med. Imaging*, **25**, pp. 922–934 (2006).
- ⁴¹J. Hausleiter, T. Meyer, M. Hadamitzky, E. Huber, M. Zankl, S. Martinoff, A. Kastrati, A. t Schömig, "Radiation Dose Estimates From Cardiac Multi-slice Computed Tomography in Daily Practice," *Circulation*, **113**, pp. 1305–1310 (2006).