Wavelet based Noise Reduction in CT-Images using Correlation Analysis

Anja Borsdorf [1], Rainer Raupach [2], Thomas Flohr [2] and Joachim Hornegger [1]

***Abstract*— The projection data measured in computed tomography (CT) and, consequently, the slices reconstructed from these data are noisy. We present a new wavelet based structure-preserving method for noise reduction in CT-images that can be used in combination with different reconstruction methods. The approach is based on the assumption that data can be decomposed into information and temporally uncorrelated noise. In CT two spatially identical images can be generated by reconstructions from disjoint subsets of projections: using the latest generation dual source CT-scanners one image can be reconstructed from the projections acquired at the first, the other image from the projections acquired at the second detector. For standard CT-scanners the two images can be generated by splitting up the set of projections into even and odd numbered projections. The resulting images show the same information but differ with respect to image noise. The analysis of correlations between the wavelet representations of the input images allows separating information from noise down to a certain signal-to-noise level. Wavelet coefficients with small correlation are suppressed, while those with high correlations are assumed to represent structures and are preserved. The final noise-suppressed image is reconstructed from the averaged and weighted wavelet coefficients of the input images. The proposed method is robust, of low complexity and adapts itself to the noise in the images. The quantitative and qualitative evaluation based on phantom as well as real clinical data showed, that high noise reduction rates of around 40% can be achieved without noticable loss of image resolution.**

***Index Terms*— noise reduction, wavelets, computed tomography, correlation analysis**

I. INTRODUCTION

COMPUTED TOMOGRAPHY (CT) is one of the most important modalities in medical imaging. Unfortunately, the radiation exposure associated with CT is generally regarded to be its main disadvantage. With respect to patients' care, the least possible radiation dose is demanded. However, dose has a direct impact on image quality due to quantum statistics. Reducing the exposure by a factor of 2, for instance, increases the noise approximately by a factor of $\sqrt{2}$. The ratio between relevant tissue contrasts and the amplitude of noise must be sufficiently large for a reliable diagnosis. Thus, the radiation dose cannot be reduced arbitrarily. State-of-the-art automatic exposure controls, which adapt the tube current according to the attenuation of the patient's body, achieve a remarkable dose reduction [1]–[3]. Further reduction, however, increases the noise level in the reconstructed images and leads to lower image quality. Many different approaches for noise suppression in CT have been investigated, for example iterative numerical reconstruction techniques optimizing statistical objective functions [4]. Other methods model the noise properties in the projections and seek for a smoothed estimation of the noisy data followed by filtered backprojection (FBP) [5]–[7]. Furthermore, several linear or nonlinear filtering methods for noise reduction in the sinogram [8]–[10] or reconstructed images [11], [12] have been proposed. In the majority of the sinogram based methods, the filters are adapted in order to reduce the most noise in regions of highest attenuation. Thus, the main goal of these methods is the reduction of directed noise and streak artifacts. As a result, especially in the case of nearly constant noise variance over all of the projections, these filters either do not remove any noise, or the noise reduction is accompanied by noticeable loss of image resolution. The goal of the new method, described in this paper, is the structure-preserving reduction of pixel noise in reconstructed CT-images and can be applied in combination with different reconstruction methods. The proposed post-processing allows either improved signal-to-noise ratio (SNR) without increased dose, or reduced dose without loss of image quality.

A very important requirement for any noise reduction in medical images is that all clinically relevant image content must be preserved. Especially edges and small structures should not be affected. Several edge-preserving approaches for noise reduction in images are known. The goal of all of these methods is to lower the noise power without smoothing across edges. Some popular examples are nonlinear diffusion filtering [13] and bilateral filtering [14], which directly work in the spatial domain. Other approaches, in particular wavelet-domain denoising techniques, are based on the scale-space representation of the input data. Most of these algorithms are based on the observation that information and white noise can be separated using an orthogonal basis in the wavelet domain, as described e.g. in [15]. Structures (such as edges) are represented in a small number of dominant coefficients, while white noise, which is invariant to orthogonal transformations and remains white noise in the wavelet domain, is spread across a range of small coefficients. This observation

Manuscript received October 17, 2006; revised March 19, 2007, revised September 6, 2007. This work was supported by Siemens Healthcare and the IMPRS on Optics and Imaging. The Associate Editor responsible for coordinating the review of this paper and recommending its publication was X. Pan.

[1] A. Borsdorf and J. Hornegger are with with the FriedrichAlexander-University ErlangenNuremberg (FAU), Chair of Pattern Recognition, Martensstr. 3, 91058 Erlangen, Germany (see http://www5.informatik.uni-erlangen.de).

[2] R. Raupach and T. Flohr are with Siemens Healthcare, Siemensstr. 1, 91301 Forchheim, Germany.

dates back to the work of Donoho and Johnstone [16]. Using this knowledge, thresholding methods have been introduced, which erase insignificant coefficients but preserve those with larger values. The difficulty is to find a suitable threshold. Choosing a very high threshold may lead to visible loss of image structures. On the other hand, a very low threshold may result in insufficient noise suppression. Various techniques have been developed for improving the detection and preservation of edges and relevant image content, for example by comparing the detail coefficients at adjacent scales [17], [18]. The additive noise in CT-images, however, cannot be assumed to be white. Furthermore, the noise distribution is usually unknown. Making matters even more complicated, noise is not stationary, violating, for example, the assumptions in [19] for estimating the statistical distributions of coefficients representing structures or noise. Motivated by the complicated noise conditions in CT, we developed a methodology which adapts itself to the noise in the images.

Recently, Tischenko et al. [20] proposed a structure-saving noise reduction method using the correlations between two images for threshold determination in the wavelet domain. Their approach was motivated by the observation that, in contrast to the actual signal, noise is almost uncorrelated over time. Two projection radiography images, which are acquired directly one after the other, show the same information but noise between the images is uncorrelated assuming, of course, that the patient does not move. Both images are decomposed by an á-trous wavelet transformation. The two highpass filtered detail images at each decomposition level are interpreted as approximations of the gradient field of the previous approximation image. The cosine of the angle between the approximated gradient vectors of the two images is used as correlation measurement. Coefficients with low correlation are weighted down and others with high correlation are kept unchanged. The result of the inverse wavelet transformation is a noise suppressed image, which still includes all correlated structures.

This concept of image denoising serves as a basis for the suppression of pixel noise in computed tomography images, proposed in this paper. The contribution of our work is as follows: We first solved the problem of how to acquire spatially identical input images in case of CT, where noise between the two images is uncorrelated. Two images, including the same information, can be generated by separate reconstructions from disjoint subsets of projections. With the latest generation dual-source CT-scanners (DSCT), the two images can be obtained directly by separate reconstructions from the projections measured at the two detectors. Using standard CT-scanners, e.g., one image can be reconstructed from the even and the other from the odd numbered projections, respectively. Furthermore, we propose a new similarity measurement based on correlation coefficients. Pixel regions from the approximation images of the previous decomposition level, which directly influence the value of a respective detail coefficient through the computation of the wavelet transformation, build the basis for our local similarity measurement. Moreover, we investigated the use of different wavelet transformations with different properties for the noise reduction based on two input images. The nonreducing á-trous algorithm (ATR), the dyadic wavelet transformation (DWT) and the stationary wavelet transformation (SWT) are compared in combination with both similarity measurements, our correlation coefficient and the gradient approximation method. In contrast to the ATR, additional diagonal detail coefficients are needed for the DWT and SWT in order to ensure perfect reconstruction. This leads to problems if the approximated gradients are used, because some correlated diagonal structures cannot be detected by comparing the angle between the approximated gradient vectors. Visible artifacts due to wrongly down-weighted detail coefficients are the result. To circumvent this problem, we propose an alternative gradient approximation method, which is computationally very efficient and is based directly on the detail coefficients. Finally, the different approaches are evaluated with respect to reduction of pixel noise and preservation of structures. We performed experiments based on phantoms and on clinically-acquired data. We show how the modulation transfer function (MTF), a standard quality measurement in CT, can be used for directly evaluating the influence of the denoising algorithm on the edge quality for different edge-contrasts. Additionally, we performed a human observer study, comparing the low-contrast-detectability in noisy and denoised images. Lastly, we also compare our approach to a projection-based noise reduction method that is used in clinical practice.

The paper is organized as follows: In Section II, the different steps used in the noise reduction method are described in detail. Section III presents the experimental evaluation based on simulated, as well as real clinical data. Finally, Section IV concludes our work.

II. Wavelet based Noise Reduction

A. Method Overview

Figure 1 illustrates the different steps of the noise reduction method. Instead of reconstructing just one image from the complete set of projections P, two images A and B, which only differ with respect to image noise, are generated. This can be achieved by separate reconstructions from disjoint subsets of projections. Image A is reconstructed from the set of projections $\mathrm{P1}$ (e.g. from the set of projections acquired at the first detector of a DSCT) and B is reconstructed from $\mathrm{P2}$ (e.g. the set of projections acquired at the second detector of a DSCT). The two images include the same information, but noise between the two images is assumed to be uncorrelated.

Both images are then decomposed into multiple frequency bands by a 2D discrete dyadic wavelet transformation. This allows a local frequency analysis. The detail coefficients of the wavelet representations include higher frequency structure information of the images together with noise in the respective frequency bands. For the reduction of high frequency noise as it is present in CT-images, only decomposition levels covering the frequency bands of the noise spectrum are of interest. It is, thus, not necessary to compute the wavelet decomposition down to the coarsest scale. The number of decomposition levels that cover the noise spectrum depends on the reconstruction field-of-view (FOV). The smaller the FOV the smaller the pixel size and consequently the higher

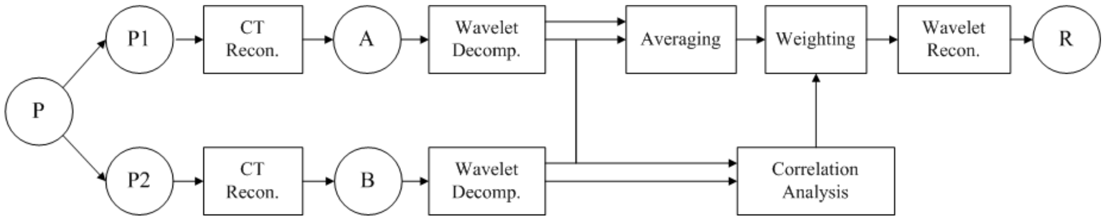

Fig. 1. Block diagram of the noise reduction method

the frequencies at the first decomposition level. Due to the logarithmic scale of the wavelet transformation, halving the FOV, e.g., means that one more decomposition level is needed. During our experiments, we found out, that in most cases few decomposition levels, e.g. 3 or 4, are sufficient because they cover approximately 90 percent of the frequencies of an image, if dyadic wavelet decompositions are used.

For each decomposition level a similarity image is computed based on correlation analysis between the wavelet coefficients of A and B. The goal is to distinguish between high frequency detail coefficients, which represent structure information and those which represent noise. High frequency structure that is present in both images should remain unchanged, while coefficients representing noise should be suppressed. A frequency dependent local similarity measurement can be obtained by comparing the wavelet coefficients of the input images. Two different approaches will be described. The similarity measurement can be based either on pixel regions taken from the lowpass filtered approximation images, or on the high frequency detail coefficients of the wavelet representation of the images.

Level dependent weighting images are then computed by applying a predefined weighting function to the computed similarity values. Ideally, the resulting masks include the value 1 in regions where structure has been detected and values smaller than 1 elsewhere. Next, the wavelet coefficients of the input images (detail- and approximation-coefficients) are averaged, what equals the computation of the wavelet coefficients of the average of the two input images because of the linearity of the wavelet transformation. The averaged detail-coefficients of the input images are then weighted according to the computed weighting image. Averaging in the wavelet domain allows the computation of just one inverse wavelet transformation in order to get a noise suppressed output image R. This output image corresponds to the reconstruction from the complete set of projections but with improved signal-to-noise ratio (SNR).

In the following subsections we will describe each step of the proposed methodology in greater detail.

B. Generation of input images

Motivated by the complicated noise conditions in CT-images (non-white, unknown distribution, non-stationary), we developed a method that is based on two spatially identical images, where noise between the images is uncorrelated. This property is used for distinguishing between structures and noise using correlation analysis in the wavelet domain. It is, however, very important to notice that the noise suppression is not performed on just one of the input images, but on the combination of both. Generally, we want to obtain a result image that corresponds to the reconstruction from the complete set of projections, but with increased SNR.

A lot of research has been done in the field of CT in the recent years. Different reconstruction methods together with their influence on noise, resolution and artifacts were investigated. Detailed descriptions regarding different methods, as well as special topics like aliasing artifacts and the propagation of noise from the projections to the reconstructed slices can be found, e.g., in [21], [22]. In this section we focus on the description of different possibilities for the generation of the input images A and B.

The input images are generated by separate reconstructions from disjoint subsets of projections $\mathrm{P1} \subset \mathrm{P}$ and $\mathrm{P2} \subset \mathrm{P}$, with $\mathrm{P1} \cap \mathrm{P2} = \emptyset$, $|\mathrm{P1}| = |\mathrm{P2}|$ and $\mathrm{P} = \mathrm{P1} \cup \mathrm{P2}$, where $|\mathrm{P}|$ defines the number of samples in P. This means that

$$A = \mathcal{G}\{\mathrm{P1}\} \quad \text{and} \quad B = \mathcal{G}\{\mathrm{P2}\}, \tag{1}$$

where $\mathcal{G}$ defines the reconstruction operator, like in our case the weighted filtered backprojection (WFBP) [23]. Generally, other reconstruction techniques can be used, however, the investigation of the influence of the reconstruction technique to the denoising method is beyond the scope of this paper. Different reconstruction methods may also lead to special requirements for the valid sets of projections $P1$ and $P2$. However, the restrictions based on Shannon's sampling theorem are valid for all kinds of reconstructions (see [24]). In the following we assume that the sampling theorem is fulfilled for both single sets of projections.

Both separately reconstructed images can be written as a superposition of an ideal noise-free signal S and a zero-mean additive noise N:

$$A = S + N_A \quad \text{and} \quad B = S + N_B, \tag{2}$$

with $N_A \neq N_B$, and the subscripts describing the different images. The ideal signal, respectively the statistical expectation E, is the same for both input images $S = \mathrm{E}\{A\} = \mathrm{E}\{B\}$ and hence also for the average $M = \frac{1}{2}(A+B)$, which corresponds to the reconstruction from the complete set of projections. The noise in both images is non-stationary, and consequently the standard deviation of noise depends on the local position $\mathbf{x} = (x_1, x_2)$, but the standard deviations $\sigma_{N_A}(\mathbf{x})$ and $\sigma_{N_B}(\mathbf{x})$ at a given pixel position are approximately the same because in average the same number of contributing quanta can be assumed. Noise between the projections P1 and P2 is uncorrelated and accordingly noise between the separately reconstructed images is uncorrelated, too, leading to the following covariance:

$$\mathrm{Cov}(N_A, N_B) = \sum_{\mathbf{x}\in\Omega} N_A(\mathbf{x}) N_B(\mathbf{x}) = 0, \quad (3)$$

with $\mathbf{x}$ defining a pixel position and Ω denoting the whole image domain.

Generally, the above scheme can also be extended to work with more than two sets of projections. The reason for restricting all the following discussions on just two input images can be found in the close relation between pixel noise σ and radiation dose d [25]:

$$\sigma \propto \frac{1}{\sqrt{d}}, \quad (4)$$

which holds as long as quantum statistics are the most dominant source of noise and other effects, like electronic noise, are negligible. If the set of projections should be split up into m equally sized parts the effective dose for each separately reconstructed image decreases by a factor of m. Thus, the pixel noise increases by a factor of $\sqrt{m}$ in every single image. The detectability of edges based on correlation analysis depends on the contrast-to-noise level, as our experiments show. Therefore, it is reasonable to keep the number of separate reconstructions as small as possible if also low contrasts are of interest, leading to $m = 2$.

The simplest possibility for acquiring P1 and P2 is to use a dual-source CT-scanner (DSCT) where two X-ray tubes and two detectors work in parallel [26]. If for both tube-detector-systems the same scan and reconstruction parameters are used, two spatially identical images can be reconstructed directly. One image is reconstructed from the projections P1 acquired at the first detector and the second one from the projections P2 of the second detector. Instead of simply averaging both images, they can be used as input to the noise reduction algorithm in order to further suppress noise (see section III-E).

If no DSCT scanner is available, different approaches for generating two disjoint subsets are possible. For example, P1 and P2 can be acquired within two successive scans of the same body region using the same scanning parameters. This requires that the patient does not move between the two scans. In order to avoid scanning the same object twice we propose another possibility for generating A and B from one single scan. As we have shown in [27], for parallel projection geometry, two complete images can be reconstructed, each using only every other projection. Specifically, one image is computed from the even and the other one from the odd numbered projections:

$$\mathrm{P1} = \left\{ P_\theta \,\middle|\, \theta = 2k\frac{\pi}{|\mathrm{P}|} \right\}, \quad (5)$$

$$\mathrm{P2} = \left\{ P_\theta \,\middle|\, \theta = (2k+1)\frac{\pi}{|\mathrm{P}|} \right\}, \quad (6)$$

with $0 \leq k \leq \frac{|\mathrm{P}|}{2} - 1$, where $|\mathrm{P}|$ denotes the total number of projections and is assumed to be even. A projection acquired at rotation angle θ is denoted as P_θ. Under the constraint that noise between different projections is uncorrelated, which means that cross-talk at the detector is negligibly small, noise between A and B is again uncorrelated as stated in equation (3). The average of the two input images again corresponds to the reconstruction from the complete set of projections, what is easy to comprehend on the example of the filtered backprojection: Reconstructing images by means of backprojection is simply a numerical integration. Thus, averaging the two separately reconstructed images corresponds to the reconstruction using the complete set of projections. It has the same image resolution and the same amount of pixel noise. However, halving the number of projections might influence aliasing artifacts and resolution in A and B. With decreasing number of projections the artifact radius, within which a reconstruction free of artifacts is possible, decreases [28]. Furthermore, azimuthal resolution is reduced away from the iso-center [21]. Usually, for CT-scanners commonly available, the number of projections is set to a fixed number that ensures a reconstruction free of artifacts within a certain field of view (FOV) at a certain maximum resolution. Thus, for the application of this splitting technique, care must be taken that the number of projections for separate reconstructions is still high enough for the desired FOV in order to avoid lower correlations due to reduced resolution or artifacts in A and B. Alternatively, the scan protocol can be adapted to acquire the doubled number of projections per rotation.

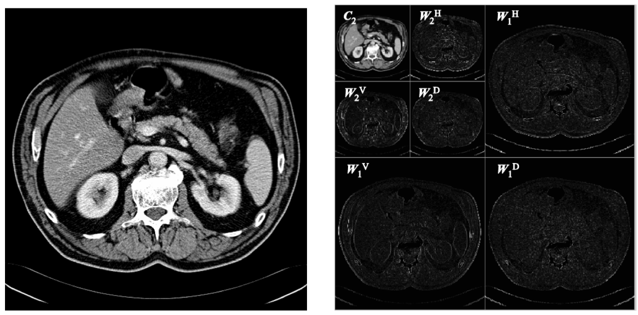

(a) (b)

Fig. 2. Example of a discrete dyadic wavelet decomposition (DWT) - (a) original image, (b) wavelet coefficients up to the second decomposition level.

C. Wavelet Transformation

This section introduces the notation and reviews the basic concepts of the three wavelet transformations used in this paper. For detailed information on wavelet theory we refer to [29]–[31].

1) DWT: The one-dimensional, discrete, dyadic, decimating (nonredundant) wavelet transformation (DWT) of a signal is a linear operation that maps the discrete input signal of length k onto the set of k wavelet coefficients. The multiresolution decomposition proceeds as an iterated filter bank. The signal is filtered with a highpass filter $\tilde{g}$ and a corresponding lowpass filter $\tilde{h}$ followed by a dyadic downsampling step respectively. This decomposition can be repeated for the lowpass filtered approximation coefficients until the maximum decomposition level $l_{\max} \leq \log_2 k$ (assumed k is a power of two) is reached. For perfect reconstruction of the signal, the dual filters g and h are applied to the coefficients at decomposition level l after upsampling. The two resulting parts are summed up leading to the approximation coefficients at level $l-1$.

When dealing with images, a two-dimensional wavelet transformation is required. The one-dimensional transformation can be applied to the rows and columns in succession, which is referred to as separable transformation. After this decomposition, four two-dimensional blocks of coefficients are available: the lowpass filtered approximation image C, and three detail images W^{H}, W^{V} and W^{D} which include high frequency structures in the horizontal (H), vertical (V) and diagonal (D) directions, respectively together with noise in the corresponding frequency bands. Like the 1D case, the 2D multiresolution wavelet decomposition can be computed iteratively from the approximation coefficients of the previous decomposition level. An example of a 2D-DWT performed on a CT-image is shown in Figure 2.

2) SWT: The computational efficiency and the constant storage complexity are advantages of DWT. Nevertheless, the nondecimating wavelet transformation, also known as stationary wavelet transformation (SWT), has certain advantages over DWT concerning noise reduction [32], [33]. Mainly, SWT works in the same way as DWT with the difference that no downsampling step is performed. In contrast to DWT, the frequency resolution is now gained by upsampling the wavelet filters $\tilde{g}$ and $\tilde{h}$ in each iteration. The number of coefficients at each decomposition level is constant, leading to an overall increased storage complexity. The reconstruction from this redundant representation is not unique. If coefficients are modified, as it is done in cases of noise reduction, an additional smoothing can be achieved by combining all possible reconstruction schemes. A further advantage is that, unlike DWT, SWT is shift-invariant.

3) ATR: A third alternative wavelet transformation we considered the two-dimensional á-trous (ATR) algorithm as described in [34]. The main difference in comparison to DWT and SWT is that only two instead of three detail images are computed at each decomposition level. The approximation coefficients C_l at decomposition level l are again computed by filtering the approximation coefficients of the previous decomposition level $l-1$ with the lowpass filter in both directions. The detail coefficients are filtered with the one-dimensional highpass only in one direction respectively, resulting in two detail images W^{H} and W^{V}. In contrast to DWT and SWT, no lowpass filtering orthogonal to the highpass filtering direction is performed. Diagonal detail coefficients are not needed for perfect reconstruction because no downsampling step is performed. For the reconstruction, however, an additional lowpass filtering orthogonal to the highpass filtering direction is necessary for the detail coefficients, in order to compensate for the missing diagonal detail coefficients [34].

D. Correlation Analysis

Detail coefficients gained from the multiresolution wavelet decomposition of the input images include structure information together with noise. The goal of the correlation analysis is to estimate the probability of a detail coefficient corresponding to structural information. This estimate is based on the measurement of the local frequency-dependent similarity of the input images.

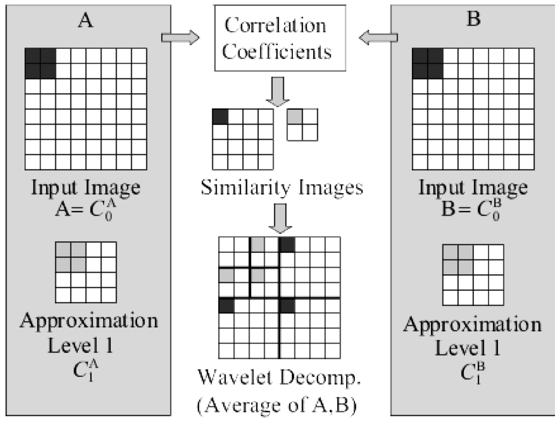

Fig. 3. Schematic description of similarity computation based on correlation coefficients between approximation coefficients of the wavelet decompositions (here DWT) obtained from the input images A and B.

Two different methods for similarity computation will be discussed. First, a correlation coefficient based measurement, comparing pixel regions from the approximation images, will be introduced. Secondly, a similarity measurement, directly based on the detail coefficients, is presented. The core idea behind both methods is similar: For all detail images of the wavelet decomposition, including horizontal, vertical (and diagonal) details, a corresponding similarity image S_l between the corresponding wavelet decompositions of the two input images A and B is computed for each level l up to the maximum decomposition level. The higher the local similarity, the higher the probability that the coefficients at the corresponding positions include structural information that should be preserved. According to the defined weighting function, the detail coefficients are weighted with respect to their corresponding values in the similarity image. Detail coefficients representing high frequency structure information are preserved, while noisy coefficients are suppressed.

1) Correlation Coefficient: One popular method for measuring the similarity of noisy data is the computation of the empirical correlation coefficient, also known as *Pearson's correlation*. It is independent from both origin and scale and its value lies in the interval $[-1; 1]$, where 1 means perfect correlation, 0 no correlation and -1 perfect anticorrelation [35]. This correlation coefficient can be used in computing the local similarity between two images, by taking blocks of pixels in a defined neighborhood around each pixel in the two images and computing their empirical correlation coefficient.

This concept can be extended by comparing images of wavelet coefficients. In order to estimate the probability for each detail coefficient of the wavelet decomposition to include structural information, we propose the computation of a similarity image at each decomposition level, as illustrated in Figure 3. The similarity image is of the same size as the detail images at that decomposition level, meaning that for each detail coefficient a corresponding similarity value is calculated.

An important factor is the selection of the pixel regions used for the local correlation analysis. A very close connection between the detail coefficients and the similarity values can be obtained if the approximation coefficients of the previous decomposition level $l-1$ are used for correlation analysis at level l, where the original image is the approximation image at level $l=0$. For the similarity value $S_l(\mathbf{x}_l)$ the correlation coefficient is computed between the approximation coefficients $C_{A,l-1}$ and $C_{B,l-1}$ within a local neighborhood $\Omega_{\mathbf{x}}$ around the corresponding position $\mathbf{x}_{l-1}$ of the current position $\mathbf{x}_l$ according to:

$$S_l(\mathbf{x}_l) = \frac{\mathrm{Cov}(C_{A,l-1}, C_{B,l-1})}{\sqrt{\mathrm{Var}(C_{A,l-1})\mathrm{Var}(C_{B,l-1})}}, \tag{7}$$

with covariance

$$\mathrm{Cov}(a,b) = \frac{1}{n} \sum_{\mathbf{x}\in\Omega_{\mathbf{x}}} \left(a(\mathbf{x}) - \bar{a}\right)\left(b(\mathbf{x}) - \bar{b}\right), \tag{8}$$

and variance

$$\mathrm{Var}(a) = \frac{1}{n} \sum_{\mathbf{x}\in\Omega_{\mathbf{x}}} \left(a(\mathbf{x}) - \bar{a}\right)^2, \tag{9}$$

where n defines the number of pixels in the neighborhood $\Omega_{\mathbf{x}}$ and $\bar{a} = \frac{1}{n}\sum_{\mathbf{x}\in\Omega_{\mathbf{x}}} a(\mathbf{x})$ defines the average value within $\Omega_{\mathbf{x}}$.

With this definition it is possible to directly use those approximation coefficients for the correlation analysis, which mainly influenced the detail coefficient at position $\mathbf{x}_l = (x_{1,l}, x_{2,l})$ through the computation of the wavelet transformation. The multiresolution wavelet decomposition is computed iteratively. Thus, the detail coefficients at level l are the result of the convolution of the approximation image at level $l-1$ with the respective analysis lowpass and highpass filters. During the computation of the inverse wavelet transformation, the approximation image at level $l-1$ is reconstructed by summing up the approximation and detail coefficients at level l filtered with the synthesis filters. The wavelets we used, all lead to spatially limited filters. Consequently, a detail coefficient at a certain position is influenced by a fixed number of pixels from the approximation image and has influence to a defined region of pixels in the approximation image due to the reconstruction. Therefore, we define $\Omega_{\mathbf{x}}$ to be a squared neighborhood according to:

$$\Omega_{\mathbf{x}} = \left\{ \tilde{\mathbf{x}}_{l-1} \,\middle|\, |\tilde{x}_{k,l-1} - x_{k,l-1}(k)| \le \frac{s}{2}, \forall\, k \in \{1,2\} \right\}, \tag{10}$$

where the length s of the four analysis and synthesis filters ($\tilde{g}$, $\tilde{h}$, g, h) is, without loss of generality, assumed to be equal and even. Consequently, the number of pixels used for the correlation analysis is adapted to the length of the wavelet filters. This is necessary in order to ensure that those coefficients, which include high frequency information of an edge can be preserved. Care must be taken if redundant wavelet transformations without downsampling are used. Then, analogously to the upsampling of the wavelet filters, the pixel regions used for correlation analysis also need to be adapted, leading to:

$$\begin{aligned} \Omega_{\mathbf{x}} \;=\; & \left\{ \tilde{\mathbf{x}}_{l-1} \,\middle|\, \left(|\tilde{x}_{k,l-1} - x_{k,l-1}| \le \frac{2^{l-1}s}{2} \right) \right. \\ & \wedge \quad \left(\mathrm{mod}\left(|\tilde{x}_{k,l-1} - x_{k,l-1}|, 2^{l-1} \right) = 0 \right), \\ & \left. \forall \quad k \in \{1,2\} \right\}, \end{aligned} \tag{11}$$

where the overall number of pixels used for correlation analysis is kept constant across the decomposition levels.

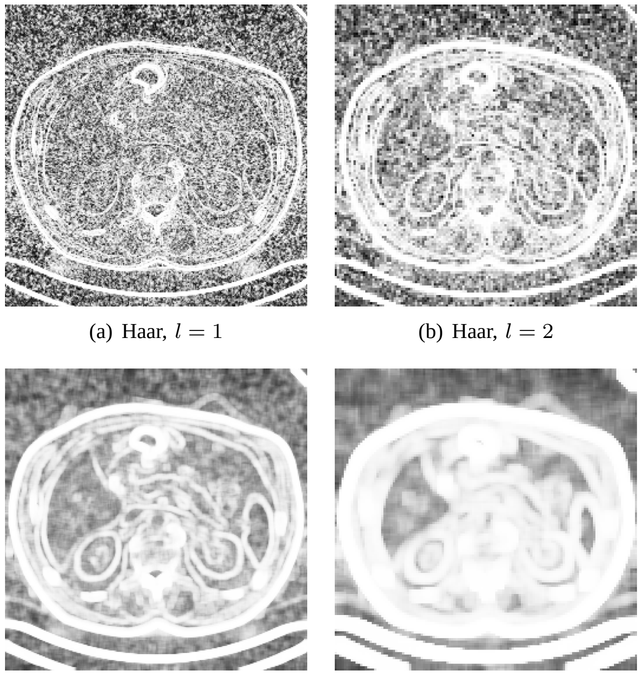

(a) Haar, $l=1$ (b) Haar, $l=2$

(c) CDF9/7, $l=1$ (d) CDF9/7, $l=2$

Fig. 4. Similarity measurement based on correlation coefficients using the Haar and CDF9/7 wavelet for the first two decomposition levels of DWT.

Figure 4 shows an example of the similarity measurement based on the correlation coefficients for the first two decomposition levels of DWT. The results are compared for two different wavelets: the Haar and the Cohen-Daubechies-Fauraue (CDF9/7) wavelet. White pixels correspond to high correlation and black to low correlation. It can be seen that especially in regions of edges high correlations are present. Additionally, it can be seen, that the area with high correlation at an edge increases from the first to the second decomposition level. The reason for this is that at the second decomposition level lower frequencies with larger spatial extension are analyzed. Furthermore, two important differences between the different wavelets, which influence the final result can be seen. Firstly, for longer reaching wavelets the region around edges where high correlations are obtained increases. Secondly, in homogeneous regions the correlation result is smoother. The Haar wavelet is the shortest existing wavelet. The corresponding analysis and synthesis filters have a length of $s=2$. Thus only those coefficients very close to the edge include information about the edge and the pixel region $\Omega_{\mathbf{x}}$ can be chosen to be very small without destroying the edge.

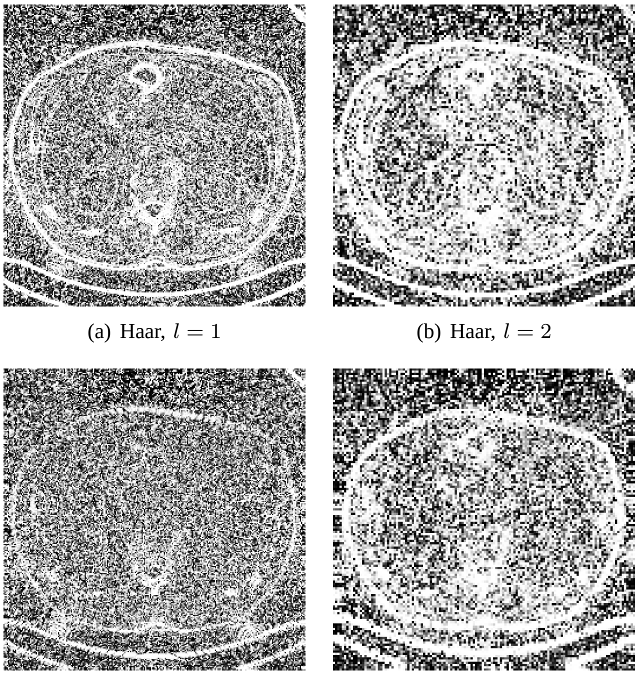

(a) Haar, $l = 1$ (b) Haar, $l = 2$

(c) CDF9/7, $l = 1$ (d) CDF9/7, $l = 2$

Fig. 5. Similarity measurement based on approximated gradients using the Haar and CDF9/7 wavelet for the first two decomposition levels of DWT.

Consequently noise can also be removed close to the edges. In contrast to that, the CDF9/7 wavelet results in filters of length $s = 10$. Thus, coefficients farther away from the edge still include information that should be preserved. This again explains the reason for adapting $\Omega_{\mathbf{x}}$ to the filterlength. Edges are preserved, but the noise reduction around high contrast edges decreases as a consequence. Because of the increased pixel region, however, a stronger smoothing can be achieved in homogeneous regions. The smaller the number of pixels used for correlation analysis, the higher the probability that noise is wrongly detected as structure. This is reflected in the higher number of white spots in combination with Haar. These observations are also confirmed by our experimental evaluation in section III.

2) Gradient Approximation: The core idea of a gradient-based similarity measurement is to exploit the fact that the horizontal and vertical detail coefficients W_l^{H} and W_l^{V} can be interpreted as approximations of the partial derivatives of the approximation image $\boldsymbol{C}_{l-1}$. In the case of the Haar wavelet, for example, the application of the highpass filter is equivalent to the computation of finite differences. Coefficients in W_l^{H} show high values at positions where high frequencies in the x_1-direction are present, while coefficients in W_l^{V} have high values where high frequencies in the x_2-direction can be found. If these two aspects are considered together, we get an approximation of the gradient field of C_{l-1}:

$$\nabla C_{l-1} = \begin{pmatrix} \frac{\partial C_{l-1}}{\partial_{x_1}} \\ \frac{\partial C_{l-1}}{\partial_{x_2}} \end{pmatrix} \approx \begin{pmatrix} W_l^{\mathrm{H}} \\ W_l^{\mathrm{V}} \end{pmatrix}. \tag{12}$$

The detail coefficients in horizontal and vertical direction of both decompositions approximate the gradient vectors with respect to Equation (12). The similarity can then be measured by computing the angle between the corresponding gradient vectors. The goal is to obtain a similarity value in the range $[-1;1]$, similar to the correlation computations of eq. 7. Therefore, we take the cosine of the angle:

$$S_l = \frac{W_{A,l}^{\mathrm{H}} W_{B,l}^{\mathrm{H}} + W_{A,l}^{\mathrm{V}} W_{B,l}^{\mathrm{V}}}{\sqrt{\left(W_{A,l}^{\mathrm{H}}\right)^2 + \left(W_{A,l}^{\mathrm{V}}\right)^2} \sqrt{\left(W_{B,l}^{\mathrm{H}}\right)^2 + \left(W_{B,l}^{\mathrm{V}}\right)^2}}, \tag{13}$$

where the index A refers to the first and B to the second input image. An example of the results of the similarity computation with the gradient approximation method is shown in Figure 5 again for the first two decompostion levels of the DWT and the Haar and CDF9/7 wavelets. Here it can already be seen that the masks look more noisy than for the correlation coefficient based approach shown in Figure 4. The difference between the Haar and CDF9/7 wavelet are very small. The edges, however seem to be better detected in combination with the Haar wavelet. These observations will also be confirmed by our quantitative evaluation III.

This kind of similarity measurement has also been used by Tischenko [20] in combination with the á-trous wavelet decomposition. As already explained above, only horizontal and vertical detail coefficients are computed in the case of the á-trous algorithm. However, the additional lowpass filtering orthogonal to the highpass filtering direction in the case of DWT and SWT is advantageous with respect to edge detection. The only problem is that the gradient approximation, as introduced so far, in the case of DWT and SWT, can sometimes lead to visible artifacts. Figure 6(a) and the difference images in Figure 6(c) show four example regions where this problem can be seen using the Haar wavelet.

Noticeably, artifacts predominantly emerge where diagonal structures appear in the image, and their shape, in general further justifies the assumption that diagonal coefficients are falsely weighted down. The different sizes of the artifacts are due to errors at different decomposition levels. Suppression of correlated diagonal structures at a coarser level influences a larger region in the reconstructed image. The reason for these types of artifacts is that diagonal patterns exist, which lead to vanishing detail coefficients in horizontal and vertical direction. If the norm of one of the approximated gradient vectors is too small or even zero, no reliable information about the existence of correlated diagonal structures can be obtained from Equation (13).

The simplest solution for eliminating such artifacts is to weight only the detail coefficients W_l^{H} and W_l^{V} based on the similarity measurement S_l and leave the diagonal coefficients W_l^{D} unchanged. As expected, this avoids artifacts in the resulting images, but, unfortunately, noise included in the diagonal coefficients is not removed, leading to a lower signal-to-noise ratio in the denoised images. Equation (13) shows that the similarity value is computed only from W_l^{H} and W_l^{V}. The diagonal coefficients do not influence S_l. The idea of extending the approximated gradient vector (see Equation (12)) by the diagonal coefficients to a three dimensional vector does not lead to the desired improvements. In the cases of vanishing detail coefficients in the horizontal and vertical direction, no

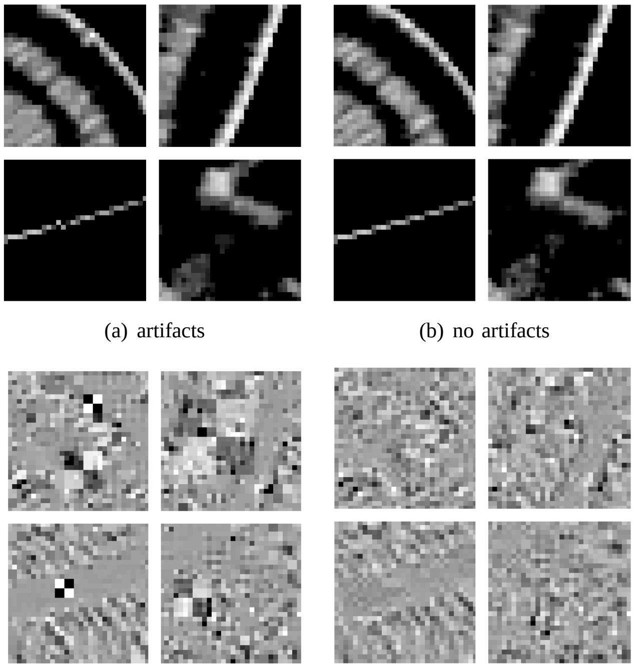

(a) artifacts (b) no artifacts

(c) difference, artifacts (d) difference, no artifacts

Fig. 6. Artifacts due to weighting down correlated diagonal coefficients with the gradient approximation method - (a) four detailed regions showing artifacts, (b) same image regions without artifacts, (c) difference between noise suppressed and original image regions showing artifacts, (d) differences free of artifact after appropriate weighting of diagonal detail coefficients.

quantitative relation between the diagonal coefficients can be obtained. Moreover, the extension of the approximated gradient vector by the diagonal coefficient is more errorprone. A diagonal coefficient can be interpreted as a second order derivative because of highpass filtering to both directions and is, therefore, very sensitive to noise. Mixing it with the horizontal and vertical detail coefficients generally leads to less reliable similarity measurements.

In order to avoid artifacts while still reducing noise in the diagonal coefficients, we propose weighting only the detail coefficients W_l^{H} and W_l^{V} depending on the similarity measurement computed from Equation (13). The diagonal detail coefficients are then treated separately. The new weighting function for the diagonal coefficients is based on the following correlation analysis between $W_{A,l}^{\mathrm{D}}$ and $W_{B,l}^{\mathrm{D}}$:

$$S_l^{\mathrm{D}} = \frac{2W_{A,l}^{\mathrm{D}} W_{B,l}^{\mathrm{D}}}{\left(W_{A,l}^{\mathrm{D}}\right)^2 + \left(W_{B,l}^{\mathrm{D}}\right)^2}. \tag{14}$$

Using this extension for a separate weighting of diagonal coefficients, denoising results are free of artifacts (see Figure 6(d)).

Note that, equations (7, 13, 14) are only defined for non-zero denominators. However, in all three cases it can be assumed that no relevant high frequency details are present if the denominator is 0 and, therefore, the similarity value is set to 0.

E. Weighting of Coefficients

The result of the correlation analysis is a set of similarity images S_l with values in the range $[-1;1]$. The closer the values are to 1, the higher the probability that structure is present. Consequently, the detail coefficient at the corresponding position should remain. The lower the similarity value, the higher the probability that the corresponding detail coefficient includes only noise and, therefore, should be suppressed. We now have to define a weighting function $f(S_l)$, that maps the values in the similarity images to weights in the range $[0;1]$. These weights are then pointwise multiplied to the averaged detail wavelet coefficients of the two input images:

$$W_{\mathrm{R},l} = \frac{1}{2}\left(W_{A,l} + W_{B,l}\right) \cdot f(S_l), \quad \forall\, l \in [1, l_{\max}], \tag{15}$$

obtaining the detail coefficients $W_{\mathrm{R}l}$ of the output image R. The approximation images of the two input images are only averaged:

$$C_{\mathrm{R},l_{\max}} = \left(C_{A,l_{\max}} + C_{B,l_{\max}}\right) \cdot 0.5. \tag{16}$$

The simplest possible method for a weighting function is to use a thresholding approach. If the similarity value S_l at a certain position is above a defined value the weight is 1 and the detail coefficient is kept unchanged, otherwise it is set to zero. Generally, the use of continuous weighting functions, where no hard decision about keeping or discarding coefficients is required, leads to better results. In principle one can use any continuous, monotonically decreasing function with range $[0;1]$, such that 1 maps to similarity values close to 1. We use the weighting function

$$f(S_l) = \left(\frac{1}{2}\left(S_l + 1\right)\right)^p \in [0,1]\,, \tag{17}$$

which has a simple geometric interpretation. In the case of the gradient approximation method, the similarity values correspond to the cosine of the angle between the gradient vectors. In the case of the correlation coefficients, the similarity value can be interpreted as the cosine of the angle between the n-dimensional vectors a and b freed by their mean, where n defines the number of pixels in $\Omega_{\mathbf{x}}$. Equation (17), therefore, leads to a simple cosine weighting, shifted and scaled to the interval $[0;1]$, where the power p controls the amount of noise suppression. With increasing p values the function goes to 0 more rapidly, but still leads to weights close to 1 for similarity values close to 1. The influence of the parameter p on noise and resolution has been evaluated in section III-A.5 and is shown in Figure 11. In all other experiments we set $p = 1$, to have a simple cosine weighting function.

We now have described all the different steps of the noise reduction method, as shown in Figure 1. We described how to generate the input images A and B, explained different possibilities for wavelet decomposition, introduced a new similarity measure between the wavelet coefficients of the input images based on correlation analysis, presented an artifact-free extension to gradient-based approximations of correlation analysis and proposed a technique for weighting the averaged details. The final step is to reconstruct the noise suppressed result image R by an inverse wavelet transformation from the averaged and weighted wavelet coefficients.

(a) no noise, $10\,\mathrm{HU}$

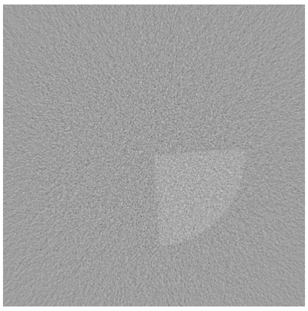

(b) noisy, $10\,\mathrm{HU}$

(c) no noise, $100\,\mathrm{HU}$

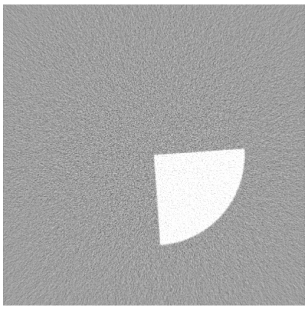

(d) noisy, $100\,\mathrm{HU}$

Fig. 7. Reconstructed simulated phantom images using S80 kernel.

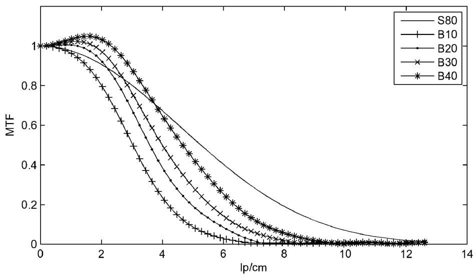

Fig. 8. MTFs of different reconstruction kernels.

III. Experimental Evaluation

For the evaluation of the described methods, experiments both on phantom data and clinically-acquired data were performed.

A. Noise and Resolution

In order to evaluate the performance of the noise reduction methods, mainly two aspects are of interest: the amount of noise reduction and, even more importantly, the preservation of anatomical structures.

1) Phantom: For our experiments we used reconstructions from a simulated cylindrical water phantom ($r = 15\,\mathrm{cm}$), with an embedded, quartered cylinder ($r = 6\,\mathrm{cm}$). The contrast of the embedded object in comparison to water varied between 10 and 100 HU. The dose of radiation ($100\,\mathrm{mAs}/1160\,\mathrm{Projections}$) is kept constant for all simulations, leading to a nearly constant pixel noise in the homogeneous area of the water cylinder. All simulations were performed with the *DRASIM* software package provided by Karl Stierstorfer [36]. The advantage of simulations is that in addition to noisy projections (with Poisson distributed noise according to quantum statistics), ideal, noise-free data can also be produced. All slices are of size 512×512 and were reconstructed within a field of view of $20\,\mathrm{cm}$ using: a) a sharp Shepp-Logan (S80) filtering kernel, leading to a pixel noise of approximately $7.6\,\mathrm{HU}$ in the homogeneous image region in the reconstruction from the complete set of projections; and b) a smoother body kernel (B40), leading to a pixel noise of approximately $5.2\,\mathrm{HU}$. The MTFs of all used kernels are shown in Figure 8. The standard deviation of noise in the separately reconstructed images is about $\sqrt{2}$ times higher. Two examples (10 and 100 HU) are shown in Figure 7. For both contrast levels, one of the noisy input images, reconstructed from every second projection, together with the ideal, noise-free image, reconstructed from the complete set of projections, are shown.

2) MTF Computation: First, we want to investigate the capability of the noise reduction algorithm to detect edges of a given contrast in the presence of noise. We are interested in how the local modulation transfer function (MTF), measured at an edge, changes due to the weighting of wavelet coefficients during noise suppression. It is possible to determine the MTF directly from the edge in an image. For this purpose, we manually selected a fixed region of 20×125 pixels around an edge (with a slope of approx. 4 degrees). The slight tilt of the edge allows a higher sampling of the edge profile, which is additionally average along the edge. The derivation of the edge profile leads to the line-spread function (LSF). The Fourier transformation of the LSF results in the MTF, which is additionally normalized so that $\mathrm{MTF}(0) = 1$. Reliable measurements of the MTF from this *edge technique* can only be achieved if the contrast of the edge is much higher than the pixel noise in the images [37]. Ideally, one should measure MTF on noise-free images. However, we are interested in measuring the quality of edge preservation based on the contrast of the edge in the presence of noise. In order to enable the measurement of a smooth MTF curve, usually, several noise realizations are needed (the number of images needed for reliable results increases, if the contrast of the edge decreases). However, the same results can be achieved even faster using the simulated data described above. We want to measure the impact of the weighting in the wavelet domain during noise suppression to the ideal signal. For that purpose, in addition to the noisy input images, which are a superposition of ideal signal and noise, an ideal image, free of noise, is simulated and reconstructed. The noise-free image is also decomposed into its wavelet coefficients. The weighting image is generated from the similarity computations from the wavelet coefficients of the noisy input images, as explained in the previous section. In order to measure the impact of the weighting to the ideal signal, the detail coefficients of the noise-free image are pointwise multiplied with the computed weights. The image gained from the inverse wavelet transformation of the weighted coefficients of the noise-free image shows the influence of the noise suppression method on structures directly. Edges, which were detected as correlated structures, are preserved. If an edge has not been detected correctly, the edge gets blurred, which influences the MTF.

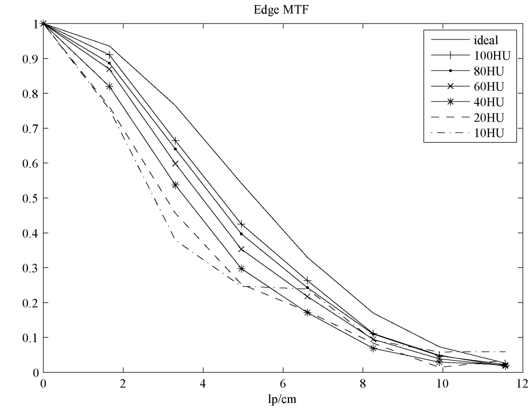

(a) Grad - ATR

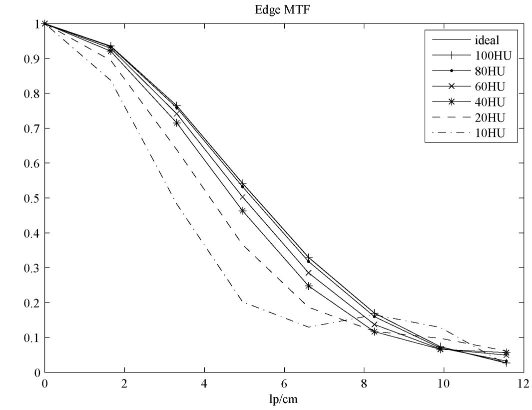

(b) Corr - ATR

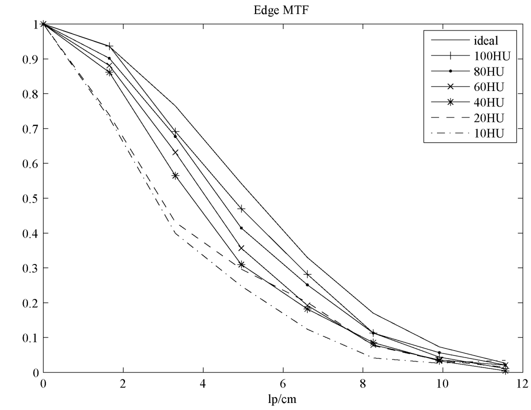

(c) Grad - DWT

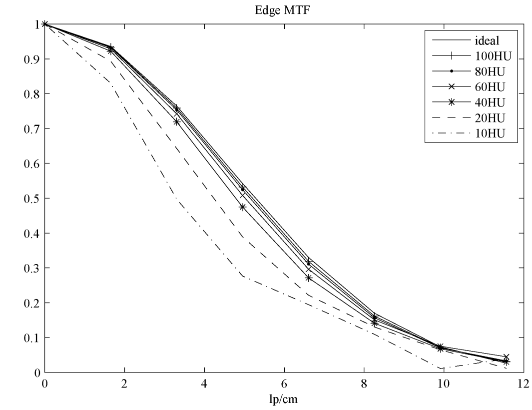

(d) Corr - DWT

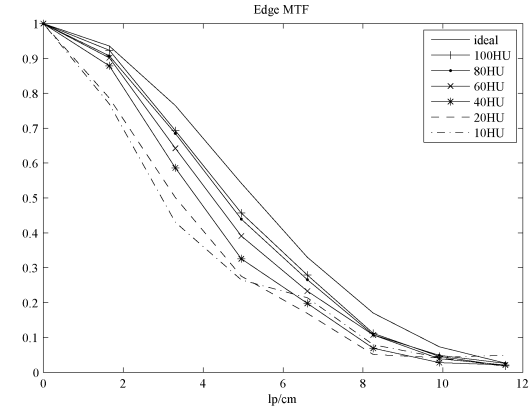

(e) Grad - SWT

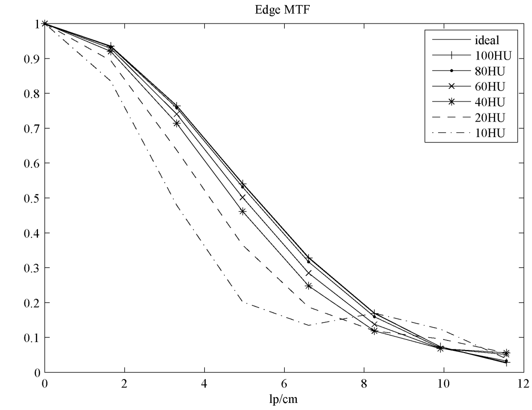

(f) Corr - SWT

Fig. 9. MTF for varying contrast at the edge using the CDF9/7 wavelet. Comparison of correlation coefficient approach (Corr) and gradient approximation (Grad) in combination with different wavelet transformations.

3) Evaluation of Edge-Preservation: In our first test, the influence of the noise suppression method to the MTF is evaluated with regard to the contrast of the edge. We used phantom images, as described above, reconstructed with the S80 kernel, with varying contrasts at the edge (10, 20, 40, 60, 80 and $100\,\mathrm{HU}$). The noise suppression method is performed for the first three decomposition levels using a CDF9/7 wavelet. In all cases a continuous weighting function is utilized, as presented in Equation (17). The MTF is computed for the modified noise-free images and compared to the MTF of the ideal image, without modifications, reconstructed from the complete set of projections. The results of this test are illustrated in Fig. 9, allowing a comparison of the different wavelet transformation methods and the *Corr* and *Grad* approaches for similarity computation. Ideally, the noise reduction methods do not influence the MTF in any respect. Specifically, the edge is not blurred. If the corresponding MTF falls below the original ideal curve, this indicates that the edge is smoothed. Alternatively, the MTF raises if some frequencies are amplified. As seen in Fig. 9 the *Corr* method leads to better edge detection in comparison to the *Grad* approach for all cases. This can be explained by the better statistical properties of the similarity evaluation based on correlation coefficients between pixel regions. More values are included in the correlation computations and, therefore, the results are more reliable. As expected, the approximated gradients are more sensitive to noise. For all methods we can see that decreasing edge contrast results in decreasing MTF. This clearly shows that decreasing CNR lowers the probability that the edge can be perfectly detected. However, one can see that with increasing contrast, the MTF gets closer to the ideal MTF. In the case of the *Corr* method the difference to the ideal MTF, even for a contrast of 60 HU, is very small. The *Grad* approach, in contrast, does not reach the ideal MTF even for an edge contrast of 100 HU. One can also observe that the performances for the three different wavelet computation methods are quite similar. The two nonreducing transformations give slightly better results in case of the *Corr* method, at least for higher contrasts. In combination with the *Grad* method, ART and SWT slightly outperform DWT. The redundant information included in nonreducing wavelet transformations, such as ATR and SWT, smooths the edge detection results. The similarity is evaluated for all coefficients. The reconstruction from the weighted redundant data, therefore, leads to smoothed results. On the other hand, the additional lowpass filtering orthogonal to the highpass filtering direction, in the case of DWT and SWT, improves the edge detection results. Altogether, this explains why SWT, which combines both positive aspects, gives best results.

An even better comparison of the results can be obtained regarding the ρ_{50} values. This is the resolution for which the MTF reaches a value of 0.5. In Fig. 10, ρ_{50} is plotted against the contrast of the edge for the different methods. This time, three different wavelets (Haar, Db2 and CDF9/7) are compared. Two different convolution kernels (S80 and B40) were used for image reconstruction (see MTFs in Fig. 8). Using a smoothing kernel changes the image resolution, as well as the noise characteristics. From Fig. 10 it can be seen that the resolution in the original image using the B40 kernel is lower than for the S80 kernel. In addition to that, the noise level is also lower (see next section on noise evaluation) using the B40 kernel. Due to the better signal-to-noise-level in the input images the edges can be better preserved when using B40. All other effects are similar for both cases. First of all, we can see that the clear differences between the *Corr* and *Grad* methods decrease when using the Db2 and the Haar wavelet. The results of the *Grad* approach get better with decreasing length of the wavelet filters. More specifically, the better the highpass filter of the wavelet is in spatially localizing edges, the better the results of the *Grad* method. For the Haar wavelet, we can see that ρ_{50} even exceeds the ρ_{50} value of the ideal image. This can be attributed to the discontinuity of the wavelet, which can lead to rising higher frequencies during noise suppression.

4) Evaluation of Noise Reduction: The same phantom images are used for evaluating the noise reduction rate. The use of simulations has the advantage that we have an ideal, noise-free image. Therefore, noise in the images can be clearly separated from the information by computing the differences from the ideal image. The effect of the noise reduction algorithm can be evaluated by comparing the amount of noise in the noise-suppressed images to that in the average of the input images. We used two different regions, each 100×100 pixels, and computed the standard deviation of the pixel values in the difference images. The first region was taken from a homogeneous area. Here the achievable noise reduction rate of the different approaches can be measured. The second region was chosen at an edge because the performance near the edges differs for the various approaches. Sometimes a lower noise reduction rate is achieved near higher contrast edges. Therefore, it is interesting to compare the noise reduction rates at edges for different contrasts. Furthermore, the noise reduction rates are evaluated for the two different reconstruction kernels (S80 and B40).

In the homogeneous image region, no noticeable changes are observed when the contrast of the objects is changed. Therefore, the measurements in cases of 100, 60 and 20 HU are averaged. Table I presents the noise reduction rates (percentage values) measured in the homogeneous image region. The first clear observation is that the noise suppression for the *Corr* method is much higher than that for the *Grad* method. The computation of correlation coefficients between pixel regions taken from the approximation images leads to smoother similarity measurements. This is also noticeable regarding the weighting matrices in Fig. 5 in comparison to Fig. 4. An interesting observation is that, for the *Grad* method, the noise reduction rates do not vary for the different wavelets. In contrast to that, when using the *Corr* approach, slightly increased noise suppression can be achieved for longer reaching wavelets. By increasing the length of the wavelet filters, larger pixel regions are used for the similarity computations. This avoids the case where noisy homogeneous pixel regions are accidentally detected as correlated. In contrast, the fact that the approximated gradient vectors in noisy homogeneous pixel regions can sometimes point to the same direction cannot be reduced by using longer reaching wavelets. The comparison

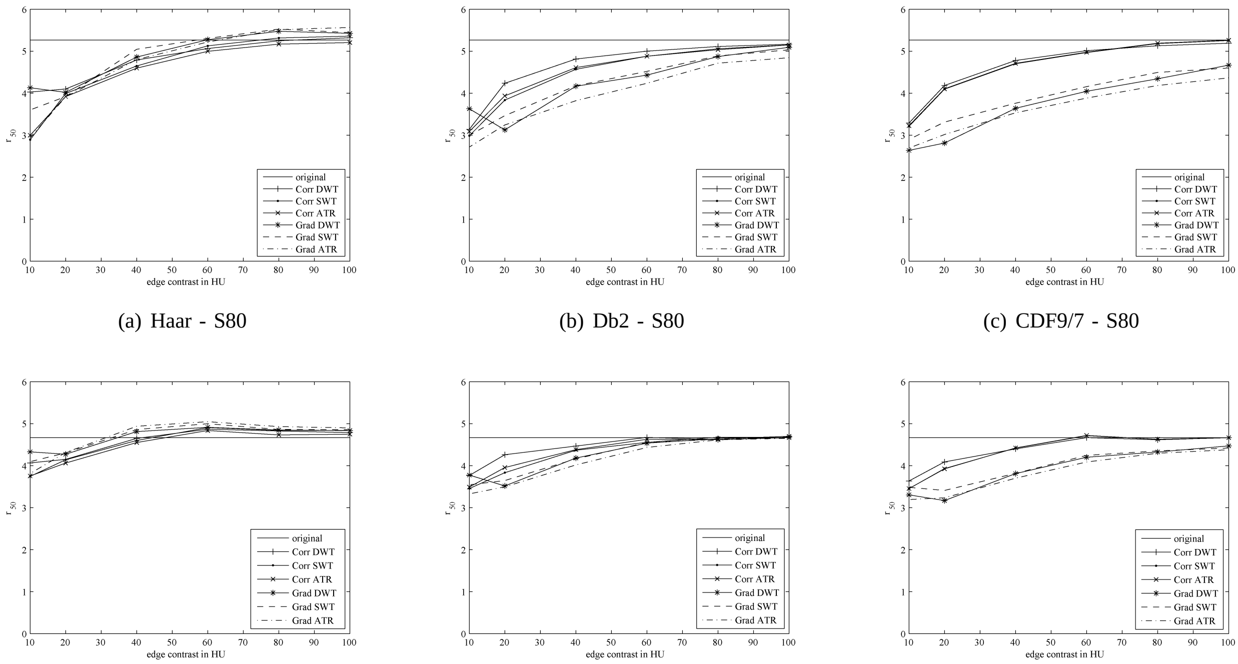

(a) Haar - S80 (b) Db2 - S80 (c) CDF9/7 - S80

(d) Haar - B40 (e) Db2 - B40 (f) CDF9/7 - B40

Fig. 10. The ρ_{50} values in dependence on contrast at the edge for different methods and wavelets.

TABLE I
PERCENTAGE NOISE REDUCTION IN A HOMOGENEOUS IMAGE REGION.

	Grad			**Corr**		
S80	ATR	DWT	SWT	ATR	DWT	SWT
Haar	26.9	22.9	26.0	42.1	39.2	40.7
Db2	27.4	22.9	26.3	46.2	44.9	45.7
CDF9/7	27.6	23.2	26.5	48.2	47.9	48.1
B40	ATR	DWT	SWT	ATR	DWT	SWT
Haar	26.6	22.0	25.4	38.9	36.3	38.3
Db2	26.1	22.5	25.9	43.5	42.2	43.2
CDF9/7	27.0	22.7	26.2	45.5	44.9	45.4

TABLE II
PERCENTAGE NOISE REDUCTION RATES IN AN EDGE REGION.

		Grad			**Corr**		
S80		ATR	DWT	SWT	ATR	DWT	SWT
100HU	Haar	25.4	21.2	24.1	38.4	35.3	37.0
	Db2	25.2	21.6	24.1	40.0	39.0	39.6
	CDF9/7	25.9	21.5	24.5	36.0	35.6	36.0
60HU	Haar	26.5	22.1	25.2	40.1	37.9	38.9
	Db2	26.6	21.6	24.8	42.1	41.1	41.8
	CDF9/7	27.0	21.7	25.4	39.0	38.0	38.9
20HU	Haar	26.6	21.7	25.1	40.7	38.1	39.4
	Db2	26.6	22.1	25.4	43.8	42.3	43.3
	CDF9/7	27.1	22.4	25.4	43.2	42.5	43.1
B40		ATR	DWT	SWT	ATR	DWT	SWT
100HU	Haar	22.9	19.8	22.4	33.0	31.1	32.6
	Db2	21.8	19.3	22.0	34.8	34.1	34.7
	CDF9/7	22.3	19.2	22.2	29.4	28.8	29.7
60HU	Haar	25.3	21.2	24.0	35.8	33.6	35.0
	Db2	24.4	19.6	23.1	37.7	36.5	37.4
	CDF9/7	25.3	19.9	23.7	32.7	31.7	32.6
20HU	Haar	25.1	20.3	24.0	36.0	34.1	35.4
	Db2	24.8	20.3	24.2	39.0	37.1	38.8
	CDF9/7	25.7	21.0	24.4	36.9	35.7	36.9

of the three wavelet transformation methods shows that DWT again has the lowest noise suppression capability, while SWT and ATR perform comparably. This shows that nonreducing wavelet transformations are better for noise suppression due to their inherent redundancy. All these observations can be made for both convolution kernels. The difference is, that in the images with lower noise level, due to the reconstruction with a smoothing kernel like the B40, the noise reduction rate is approximately 3 percent points in the case of the *Corr* method and less than 1 percent point in the case of the *Grad* method below the noise reduction rate in the more noisy images reconstructed with the S80.

Table II lists the noise reduction rates achieved in the edge region, again using the two different convolution kernels. Here, the results are compared for three different contrasts at the edge. Most of the observations we made for the homogeneous image region are also valid for the edge region. Our *Corr* approach clearly outperforms the *Grad* method. The DWT shows the lowest noise suppression, whereas ART and SWT are comparable. In the case of the *Grad* method, we can again observe that nearly no differences between the different wavelets can be obtained. Generally, we can see that with decreasing contrast at the edge, more noise in the local neighborhood of the edge can be removed. The reason for this is that the lower the contrast, the lower the influence of the edge to the correlation analysis. However, one difference between the two similarity computation methods becomes clear. For the *Grad* approach the increment in noise suppression with decreasing contrast at the edge is quite similar for all wavelets.

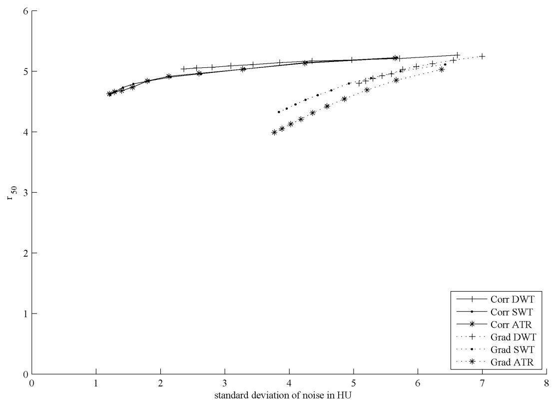

Fig. 11. Noise-Resolution-Tradeoff: Comparison of high-contrast resolution and standard deviation of noise in homogeneous image region for different denoising methods using Db2 wavelet. The power p within the weighting function (17) is used for varying the amount of noise suppression.

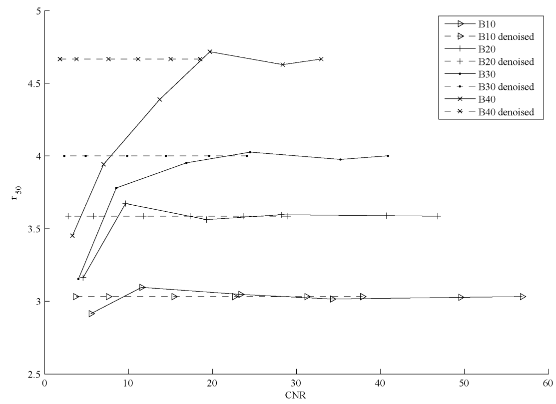

Fig. 12. Noise-Resolution-Tradeoff: ρ_{50} polotted against CNR for different reconstruction kernels. Denoising configuration: 3 level SWT with CDF9/7 wavelet and *Corr* method.

This does not hold for the *Corr* approach. Here, we can see that by increasing the spatial extension of the wavelet filters, the difference between the noise suppression rate at 100 HU increases in comparison to 20 HU. This means that for higher contrast, more noise close to edges remains in the image if longer reaching filters are utilized. The reason for this is that the size of the pixel regions used for the correlation computations are adapted to the filter lengths of the wavelets. This is needed in order to ensure that all coefficients, which include information of an edge, are included in the similarity computations, as already remarked during the discussion of Fig. 4. The effect is that edges with contrast high above the noise level dominate the correlation computation, as long as they occur within the pixel region. As a result, nearly no noise is removed within a band around the edge. The width of the stripe depends on the spatial extension of the wavelet filters.

5) Noise-Resolution-Tradeoff: Within the last two sections we presented a very detailed, contrast dependent evaluation of noise and resolution. For easier comparison of the different denoising approaches, noise-resolution-tradeoff curves are plotted in Fig. 11. The phantom described in section III-A.1 with an edge-contrast of 100 HU, reconstructed with the S80 kernel, was used for this experiment. The ρ_{50} values are plotted against the standard deviation of noise, measured within a homogeneous image region. The *Corr* and *Grad* method in combination with DWT, SWT and ATR are compared, all using the Db2 wavelet and 3 decomposition levels. The power p within the weighting function (17) was used for varying the amount of noise suppression. The 10 points within each curve correspond to the powers $p = \{5.0, 4.5, 4.0, 3.5, 3.0, 2.5, 2.0, 1.5, 1.0, 0.5\}$ from left to right. In summary the following obervations can be made:

- SWT and DWT show better edge-preservation than ATR at the same noise reduction rate in combination with the *Grad* method.
- The *Corr* method clearly outperforms the *Grad* method in all cases.
- There is nearly no difference between the different wavelet transformations if the *Corr* approach is used.

In a second test, the influence of the reconstruction kernel to the noise-resolution-tradeoff was evaluated. Different reconstruction kernels can be selected in CT, always leading to a noise-resolution-tradeoff. Smoothing reconstruction kernels implicate lower noise power, but also lower image resolution. As we have already seen during the discussion of noise and resolution in the last two sections the reconstruction kernel also influences the results of the denoising method. Therefore, we compared the noise-resolution-tradeoff for different reconstruction kernels (see Fig. 8) with and without the application of the proposed denoising method. We used again the phantom images described in section III-A.1 with varying contrasts c, reconstructed with B10, B20, B30 and B40 kernel. We then compared the contrast-to-noise ratio ($\mathrm{CNR} = c/\sigma$) and resolution (ρ_{50}) of the original and denoised images. We used a 3 level SWT with CDF9/7 wavelet and the *Corr* method for the comparison shown in Fig. 12. The dashed lines correspond to the original and the solid lines to the denoised images. Each line consists of five points corresponding to the contrasts (10, 20, 40, 60, 80 and $100\,\mathrm{HU}$) devided by the respective standard deviation of noise σ measured in a homogeneous image region. Ideally the denoising procedure would only increase the CNR without lowering resolution. This would mean that the solid lines are just shifted to the right in comparison tho the corresponding dashed lines. The observed behavior, however, was more complex and corroborates the results presented in the previous sections:

- The sharper the kernel (high resolution, low CNR), the higher the improvement in CNR that can be achieved by applying the proposed method.
- The smoother the kernel (low resolution, high CNR), the better the edge-detection and thus the preservation of resolution in the denoised image.

The new insight we gained from this analysis is that we can achieve better results with respect to image resolution and CNR using a sharper reconstruction kernel in combination with our proposed method than using a smoothing reconstruction kernel. For example, we can achieve higher resolution and higher CNR for the same input data if the sharper B30 kernel is used in combination with our filter than using the smoother

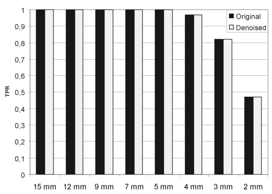

(a) 10 HU

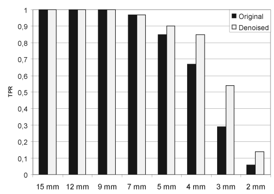

(b) 5 HU

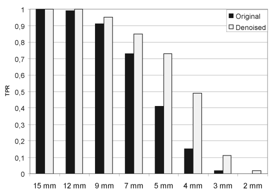

(c) 3 HU

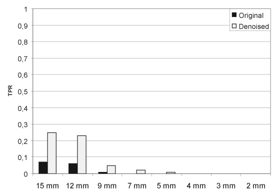

(d) 1 HU

Fig. 14. Comparison of true-positive rates for different objects of noisy (original) and denoised LCP.

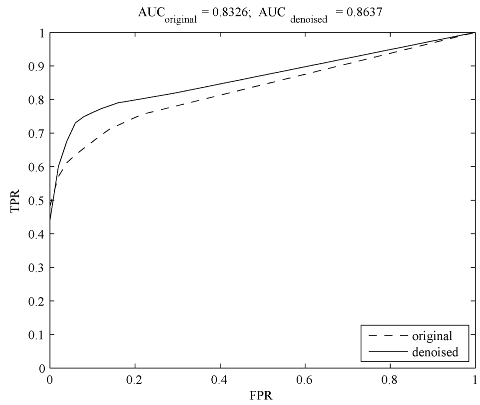

Fig. 15. ROC curves resulting from human observer study. Comparison between noisy (original) and denoised results.

B10 kernel without denoising.

B. Low-Contrast-Detectability

In addition to the quantitative evaluation of noise and resolution we performed a human observer study to test how the low-contrast-detectability is influenced by the application of our proposed method.

1) Data and Experiment: For our experiments we used reconstructions from a simulated cylindrical water phantom ($r = 14.5\,\mathrm{cm}$), with four blocks of embedded cylindrical objects with different contrasts ($10, 5, 3, 1\,\mathrm{HU}$) and different sizes ($15, 12, 9, 7, 5, 4, 3, 2\,\mathrm{cm}$ diameter). A reconstructed slice from this phantom is shown in Fig. 13(a). We simulated and reconstructed 10 noisy realizations of this phantom, all at the same dose level ($30\,\mathrm{mAs}$), leading to an average pixel noise in the homogeneous water region of $\sigma = 4.3\,\mathrm{HU}$. One noisy example slice is shown in Fig. 13(b). In addition to this, 20 noisy phantoms where some (95 in sum) of the embedded objects were missing were simulated and reconstructed using the same scanning and reconstruction parameters. For all 30 images the corresponding denoised images (with approx. 44% noise reduction, leading to $\sigma = 2.4\,\mathrm{HU}$ in average) were computed. We used 3 decomposition levels of SWT in combination with CDF9/7 wavelet together with the *Corr* method. As an example, in Fig. 13(c) the denoised image of Fig. 13(b) can be seen.

For easier accomplishment and evaluation of the experiment we developed a proprietary evaluation tool for low-contrast-detectability. This tool showed the images from a list in randomized order to the human observer. The observer then had to select which objects he can detect by mouse click. All 47 observers evaluated 40 images, 10 original noisy images where all objects were present, the 10 corresponding denoised, 10 noisy images where some objects were missing, and again the 10 corresponding denoised.

2) Results and Discussion: In a first step we evaluated the average true-positive rate (TPR) achieved for the different objects. The performance of detecting objects of different size and contrast was compared between the noisy and denoised images. We computed the average TPR for all 32 objects from all noisy images and all observers and compared it to the average from all denoised images and all observers. In Fig. 14 the TPR is plotted for all objects of different contrasts and sizes. The closer the TPR is to 1 the better the object was correctly judged to be visible in average. The clear result is that all objects were judged to be as well or even better detectable in the denoised images in comparison to the noisy originals. The corresponding false-positive rates (FPR) are all below 0.03 and in average below 0.005 for both noisy and denoised images. In Fig. 14(a) the TPR for the $10\,\mathrm{HU}$ objects can be seen, where no clear difference between the noisy and denoised objects is visible. In case of the $5\,\mathrm{HU}$ and $3\,\mathrm{HU}$ objects (see Fig. 14(b) and 14(c)) a clear difference can be seen. If objects with a TPR above 0.5 are said to be detectable, two more objects ($5\,\mathrm{HU}, 3\,\mathrm{mm}$ and $3\,\mathrm{HU}, 5\,\mathrm{mm}$) are detectable in the denoised images than in the noisy ones. The TPR of the $3\,\mathrm{HU}, 4\,\mathrm{mm}$ object is also very close to 0.5. The $1\,\mathrm{HU}$ objects were nearly never detected in the noisy images, but in the denoised at least the 15 and $12\,\mathrm{mm}$ objects were detected correctly in more than 20% of the cases.

In a second step, we computed the ROC-curves for the noisy and denoised cases based on a thresholding approach as described in detail in [38]. We firstly computed the average detection rate (number of positive votes that object is visible / number of overall votes for this object) for each single image and object from all observers. Then, a sliding threshold was applied for the noisy and denoised cases separately. All objects with a detection rate above a certain threshold were set to be detected and then the corresponding FPR and TPR was calculated, leading to the curves shown in Fig. 15. In addition to the curves the area under the curve (AUC) was computed for the noisy and denoised case. The AUC improved from 0.8326 in case of the noisy to 0.8637 in case of the denoised samples.

C. Comparison with Adaptive Filtering of Projections

1) Data and Description: Fig. 16 shows a comparison of the proposed method to a projection based adaptive filtering,

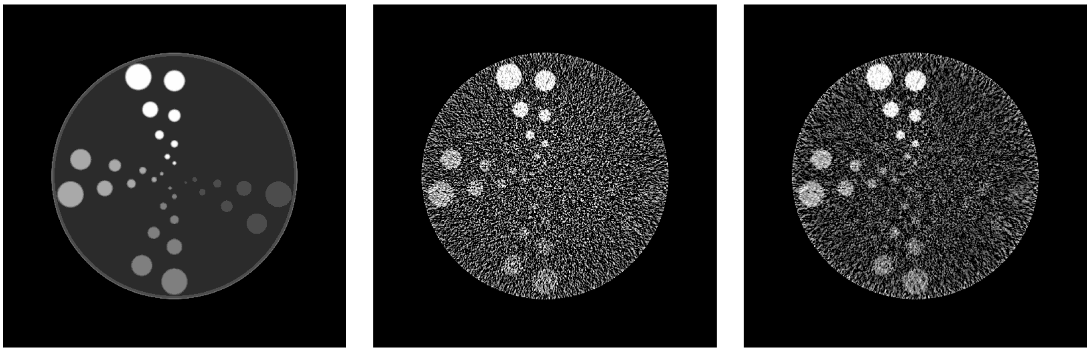

(a) no noise (b) noisy (c) denoised

Fig. 13. Low-contrast-phantom (LCP) used for human observer study: (a) ideal noise-free, (b) one noisy example and (c) corresponding denoised image using 3 levels of SWT, CDF9/7 wavelet and the *Corr* method. Display options: $c = 5$,$w = 12$.

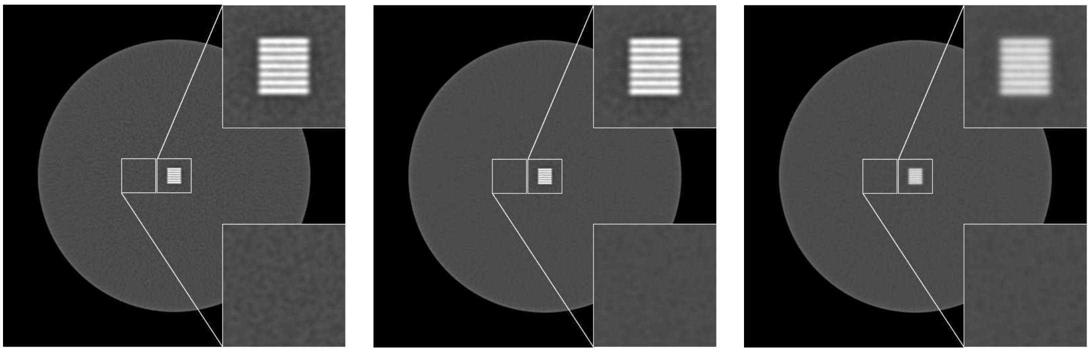

(a) original: $\sigma = 11.1\,\mathrm{HU}$ (b) proposed method: $\sigma = 6.0\,\mathrm{HU}$ (c) adaptive filtering of projections: $\sigma = 6.0\,\mathrm{HU}$

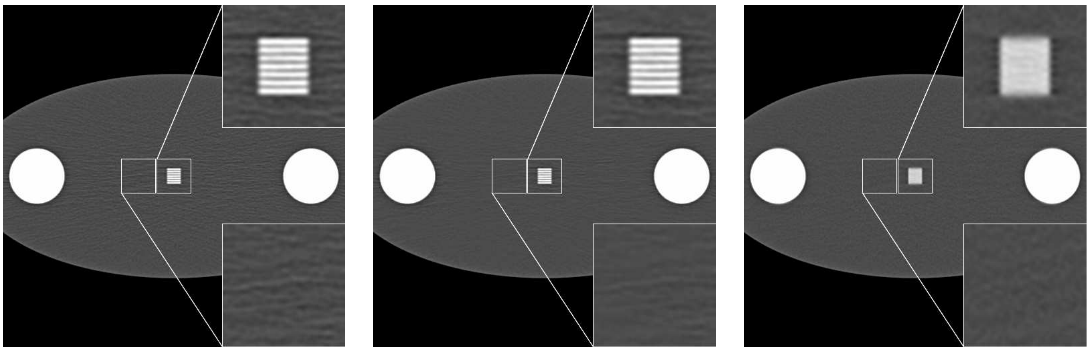

(d) original: $\sigma = 19.0\,\mathrm{HU}$ (e) proposed method: $\sigma = 10.2\,\mathrm{HU}$ (f) adaptive filtering of projections: $\sigma = 10.2\,\mathrm{HU}$

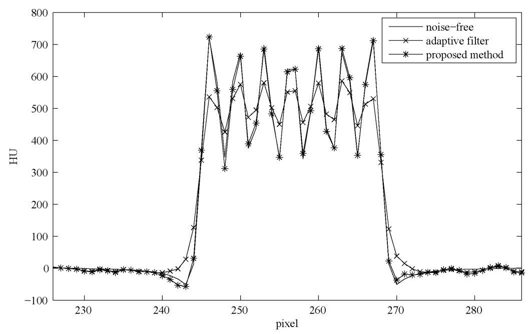

(g) vertical lineplot through noise-free, (b) and (c)

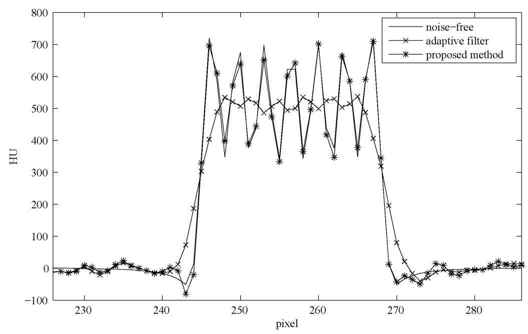

(h) vertical lineplot through noise-free, (e) and (f)

Fig. 16. Comparison of proposed method to adaptive filtering of projections. The reconstruction without noise suppression are displayed in (a) and (d). The proposed wavelet based noise reduction method was applied in (b) and (e) and the adaptive filtering of the projections is shown in (c) and (f). Image resolution of the filtered images is compared at the same noise reduction rates. In (g) and (h) the corresponding vertical lineplots through the center of the two phantoms are compared between the noise-free, adaptive-filtered and wavelet denoised images. Display options: $c = 200$ and $w = 1000$.

which is used in clinical practice [39]. The 2D-projections are filtered with a linear filter of fixed spatial extension. Then, a weighted sum of the filtered and original noisy projections is computed based on the attenuation at a respective position. The higher the attenuation, the higher the noise power and, therefore, the stronger the smoothing being performed. This method, like most other noise reduction methods based on filtering the projections, has the goal to reach nearly constant noise variance over all projections in order to reduce directed noise and streak artifacts.

For the comparison we used reconstructions from two simulated elliptical phantoms, one homogeneous water phantom ($r = 10\,\mathrm{cm}$) and one eccentric water phantom ($a = 15\,\mathrm{cm}$ and $b = 7.5\,\mathrm{cm}$). In the center of both phantoms a line-pattern with $6\,\mathrm{lp/cm}$ was embedded at a contrast of $1000\,\mathrm{HU}$. In the eccentric phantom two additional cylindrical objects ($r = 2\,\mathrm{cm}$) are embedded. All reconstructions to a pixel grid of 512×512 with FOV of $250\,\mathrm{cm}$ were performed using the B40 kernel. In Fig. 16(a) and 16(d) the original noisy phantoms reconstructed from the complete set of projections are shown. We measured the standard deviation of noise in homogeneous regions in north, south, west and east direction around the center resulting in an average noise of $\sigma = 11.1\,\mathrm{HU}$ in the homogeneous, and $\sigma = 19.0\,\mathrm{HU}$ in the eccentric water phantom. We applied both denoising methods to achieve the same average noise reduction rate, leading to $\sigma = 6.0\,\mathrm{HU}$ in the homogenous and $\sigma = 10.2\,\mathrm{HU}$ in the eccentric case, and compared resolution. For our proposed method we used 3 levels of SWT together with CDF9/7 wavelet and the *Corr* method.

2) Results and Discussion: Fig. 16(f) shows that directed noise pointing out the direction of highest attenuation is reduced and a remarkable noise suppression can be achieved by adaptively filtering the projections. However, it can also be noticed that structures orthogonal to the direction of highest attenuation visibly lose spatial resolution. In comparison to this the wavelet based filtering method preserves structures much better and no blurring effects are visible. This can be seen well in the detailed vertical lineplots through the line-pattern, shown in Fig. 16(g). Although the same average noise reduction rate was obtained, the streak artifacts could not be completely removed using the wavelet approach. This is the strength of the adaptive filtering method. In contrast to this, the adaptive filtering method does not perform well if nearly homogeneous objects are present. The goal of the adaptive filtering of the projections is to achieve nearly constant noise variance over all projections. If the noise variance is already very similar in all projections, the adaptive filtering does nothing at all, or loses resolution in all directions. This can be seen well in Fig. 16(c). Here, the wavelet based method, as shown in Fig. 16(b), can again achieve a high noise reduction rate without loss of resolution. The detailed vertical lineplots are again shown in Fig. 16(h). Nevertheless, we want to emphasize that the noise suppression based on the projections is a processing step prior to reconstruction, while the proposed method is a post-processing step, thus making the combination of the two methods possible

D. Clinically-Acquired Data

1) Data and Experiment: In order to test the noise reduction method with respect to its practical usability, the application of the algorithm on clinically-acquired data is indispensable. Noise reduction methods are particularly critical in their application to low contrast images. Thus, images predominantly including soft tissue are well suited for performance assessment. Theoretically, as already discussed, the higher the contrast of edges, the higher the probability that the edge can be detected and preserved. If the application of the method with specific parameter settings leads to good results in slices with soft tissue, the use for higher contrast regions will not be critical. Therefore, we used a thoraco-abdomen scan (see examples in Fig. 18), acquired at a Siemens Sensation CT-scanner, for the clinical evaluation. The reconstruction of slices at a FOV of $38\,\mathrm{cm}$ with a thickness of $3\,\mathrm{mm}$ was performed with a B40 kernel, which is one of the standard kernels for this body region.

For our clinical experiment, we computed 12 noise-suppressed images from the same input images with different configurations. We used three different wavelet transformation methods (ATR, DWT and SWT) in combination with two different wavelets (Haar and CDF9/7). Furthermore, we used these configurations together with the *Corr* and *Grad* methods for similarity computation. The resulting noise-reduced images and the average of the input images, which corresponds to the reconstruction from all projections, were compared by a radiologist. All images correspond to the same dose level. For simple comparison, we developed a proprietary evaluation tool. A randomized list of comparisons between image pairs can be performed with this tool. Within each comparison, an image pair is shown to the radiologist. The initial position of the two images is also randomized. However, the positions of the two images can be easily switched by the radiologist, in order to facilitate the detection of even very small differences between the images. The radiologist decides if there is one preferred image (clear winner) or both images are judged of equal quality with respect to some predefined evaluation criteria.

Three different quality criteria were evaluated separately in three consecutive tests:

- detectability of anatomical structures,
- noise in homogeneous image regions,
- noise in edge regions.

In each test, all possible image pairs were compared to each other. Altogether, 3×78 comparisons were performed. The outcome of these tests is shown in Fig. 17. The dark bars show the number of clear winners, normalized to the number of performed comparisons for one image. The corresponding light bars are the results of a score system. Three points are gained by a winning image and one point if two images are judged to be of equal quality. This value is again normalized, this time to the number of maximally reachable points, if the image won all comparisons.

2) Results and Discussion: In the first test (Fig. 17(a)), the detectability of anatomical structures was examined. Only in one case the anatomical structures were judged to be better

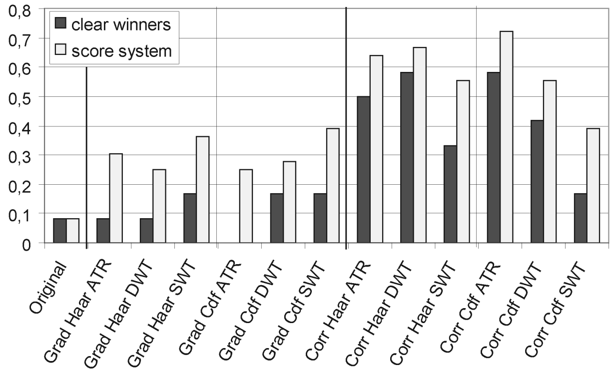

(a) Detectability of anatomical structures

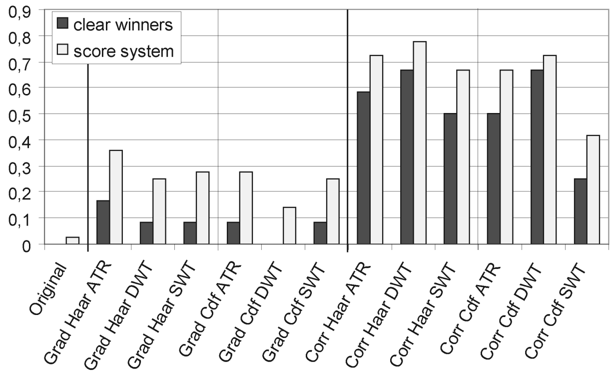

(b) Noise in homogeneous image region

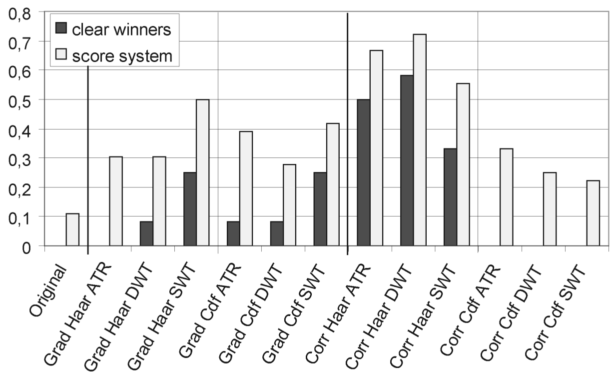

(c) Noise in edge region

Fig. 17. Results of the clinical evaluation - (a) Detectability of anatomical structures; (b) noise in homogeneous regions and (c) noise in edge regions was compared for different configurations.

detected in the original image than in the noise suppressed image. In all other direct comparisons of noise reduced images to the average of input images (here denoted as original), the processed images were chosen to be favorable. This shows that the anatomical structures are well preserved by the noise suppression method. The separation of information and noise is further improved because of the better signal-to-noise ratio. The comparison between the different configurations shows that our *Corr* method gives better edge detection results than the *Grad* approach. There is no clearly preferred wavelet basis or wavelet transformation.

In the second test (Fig. 17(b)), the treatment of noise in homogeneous image regions was analyzed. Here again, the *Corr* method gives much better visual results in all cases. There is nearly no difference between the Haar and the CDF9/7 wavelet.

In the last test (Fig. 17(c)), the noise in regions around edges was compared. This test reflects the results of the quantitative evaluation with phantom data. It shows that nearly no noise is removed in regions of edges if long reaching wavelets are used in combination with the *Corr* method. The results of the Haar wavelet are still judged better for the *Corr* method in comparison to the *Grad* approach.

The *Corr* method is clearly preferred considering the results of all three tests together. However, longer reaching wavelets lead to lower noise reduction around higher contrast edges. Therefore, a tradeoff between smoothness and spatial locality of the wavelet must be resolved.

E. Example Images

Two examples of noise suppression on clinically acquired data are shown in Fig. 18. Zoomed-in images from the abdomen (18(a)-18(c)) and thorax (18(d)-18(f)) are displayed. For denoising we used 3 levels of a Haar wavelet decomposition (SWT) in combination with the *Corr* method. The original images, which correspond to the reconstruction from the complete set of projections, are compared to the noise suppressed images. Additionally, the differences between the original and denoised images are shown. Noticeably, noise in homogeneous image regions is removed, while structures are well preserved.

In Fig. 19 two examples of a thorax-abdomen phantom acquired at a Siemens Definition dual-source CT (DSCT) scanner are shown. We used the same scan protocol ($100\,\mathrm{mAs}, 120\,\mathrm{kV}$, slice-thickness$=1.2\,\mathrm{mm}$) and reconstruction parameters (FOV $=350$, kernel = B30) for both source-detector systems. The image reconstructed from projections acquired at the first detector is denoted as A and the image from the second detector is denoted as B. The FOV ($26\,\mathrm{cm}$) of the second detector is smaller than that of the first detector. Therefore, the sinogram of the B-system is extended at the outer border with data from the A-system, as explained in detail in [26]. With this technique two images can be reconstructed at the full FOV. Inside the $26\,\mathrm{cm}$-FOV we have independent acquisitions from the two detectors. Consequently, noise within these regions can be assumed to be uncorrelated between the two images. Outside the FOV of $26\,\mathrm{cm}$ only parts of the sinogram derive from independent measurements due to the sinogram padding. Therefore, noise in this outer region is no longer perfectly uncorrelated. Evaluating the correlation during the *Corr* method or comparing the angle between the approximated gradient vectors in the *Grad* method still works in this outer region. However, only a lower noise reduction can be achieved because of the increasing correlation between A and B with increasing distance from the $26\,\mathrm{cm}$ radius. In Fig. 19(a) and 19(d) the average images of A and B are shown for two examples. The A and B images are then used as input to the proposed noise reduction method (3 levels SWT with Db2 wavelet and *Corr* method). The corresponding denoised results are shown in Fig. 19(b) and 19(e). For

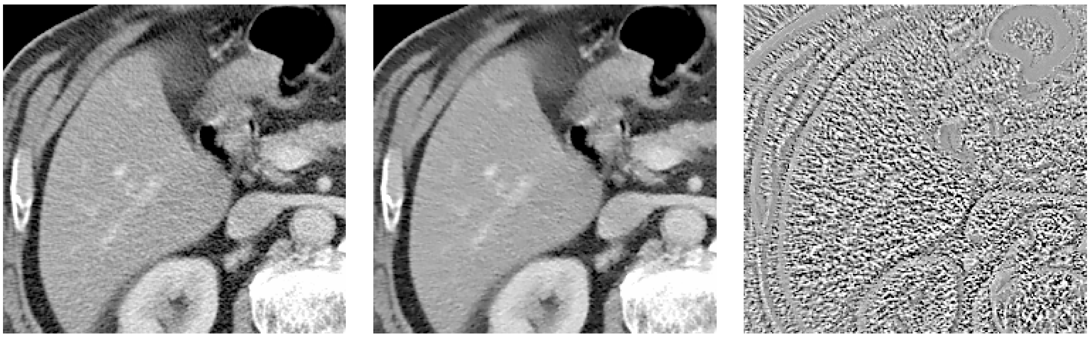

(a) original (b) denoised (c) difference

(d) original (e) denoised (f) difference

Fig. 18. Noise suppression in real clinical images from the abdomen (a)-(c) and thorax (d)-(f). Configuration: SWT, Haar wavelet, $l_{max} = 3$, $p = 1$, *Corr* method. Display options: $c = 50$, $w = 400$ for CT-images and $c = 0$,$w = 50$ for difference images.

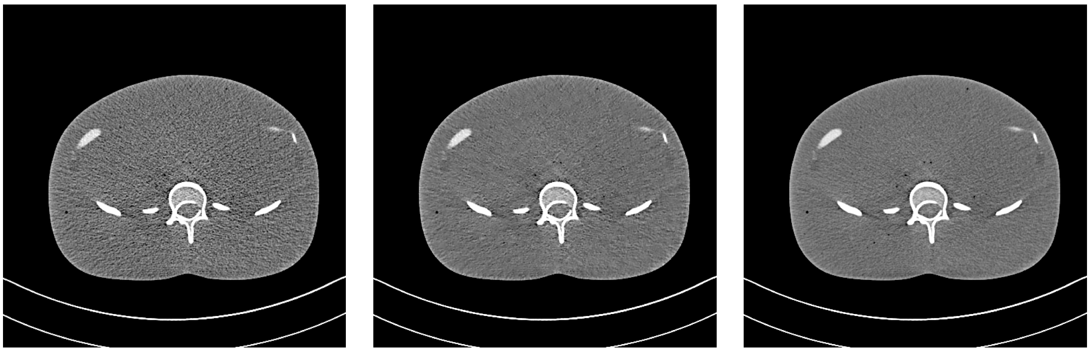

(a) $100\,\mathrm{mAs}$, $\sigma = 18.3\,\mathrm{HU}$ (b) $100\,\mathrm{mAs}$ denoised, $\sigma = 10.4\,\mathrm{HU}$ (c) $500\,\mathrm{mAs}$, $\sigma = 8.9\,\mathrm{HU}$

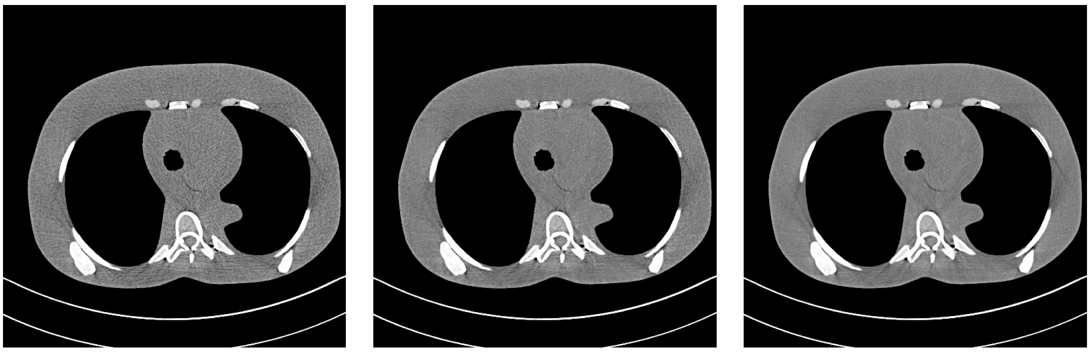

(d) $100\,\mathrm{mAs}$, $\sigma = 17.0\,\mathrm{HU}$ (e) $100\,\mathrm{mAs}$ denoised, $\sigma = 9.8\,\mathrm{HU}$ (f) $500\,\mathrm{mAs}$, $\sigma = 8.3\,\mathrm{HU}$

Fig. 19. Application of proposed method to Dual-Source-CT data: abdomen (a)-(c) and thorax (d)-(f). Configuration: SWT, Db2 wavelet, $l_{max} = 3$, $p = 1$, *Corr* method. Display options: $c = 50$, $w = 300$.

better comparison high-dose scans ($500\,\mathrm{mAs}$) are shown in Fig. 19(c) and 19(f). Within the overlapping FOV, where data from both detectors has been acquired, a noise reduction rate of approximately 43% was achieved. Due to the sinogram extension of the B-system with data from A, noise outside the FOV of $26\,\mathrm{cm}$ is no longer perfectly uncorrelated. Therefore, only a lower noise reduction of approximately 25% can be achieved in regions outside the overlapping FOV.

IV. Conclusions

In this paper, we have introduced a new, robust and efficient wavelet domain denoising technique for the suppression of pixel noise in CT-images. The separate reconstruction from disjoint subsets of projections allows the generation of images which only differ with respect to image noise but include the same information. We showed that correlation analysis based on the detail coefficients of the á-trous wavelet decomposition of the input images, as recently proposed by Tischenko, allows the separation of structures and noise, without assuming or estimating the underlying noise distribution. We extended the approach for the applicability with DWT and SWT. The quantitative and qualitative evaluation showed that comparable edge preservation, with only slightly lower noise reduction, can also be achieved with DWT at lower computational costs. More importantly, a second similarity measurement was introduced which makes use of correlation coefficients. This has lead to improved results with respect to edge preservation and noise suppression for all wavelet transformations. The performed human-observer study showed that the detectability of small low-contrast objects could be improved by applying the proposed method. In comparison to a commonly applied projection based algorithm, the proposed method achieved higher resolution at the same noise suppression. The evaluation on clinically-acquired CT data proves the practical usability of the methods.

Currently, we are working on the extension of the method to 3D. Improved results with respect to noise and resolution are expected, due to the more reliable correlation analysis in 3D. Further, we are planing more extended clinical tests, including additional human-observer studies, in order to investigate the potential dose reduction.

Acknowledgment

The authors would like to thank Naomi Knoess M. D. for evaluating the clinical data, and Holger Kunze, Karl Stierstorfer, Herbert Bruder, Martin Petersilka, Elli Angelopoulou and Laszlo Nyul for helpful suggestions and discussions.

References

[1] W. A. Kalender, H. Wolf, C. Suess, M. Gies, and W. Bautz, "Dose reduction in CT by on-line tube current control: Principles and validation on phantoms and cadavers," *European Radiology*, vol. 9, no. 2, pp. 323–328, 1999.

[2] H. Gress, H. Wolf, U. Baum, M. Lell, M. Pirkl, W. Kalender, and W. Bautz, "Dose reduction in computed tomography by attenuation based on-line modulation of tube current: Evaluation of six anatomical regions," *European Radiology*, vol. 10, no. 2, pp. 391–394, 2000.

[3] C. Suess and X. Y. Chen, "Dose optimization in pediatric CT: Current technology and future innovations," *Pediatric Radiology*, vol. 32, no. 10, pp. 729–734, 2002.

[4] I. Elbakri and J. Fessler, "Efficient and accurate likelihood for iterative image reconstruction in x-ray computed tomography," in *Medical Imaging 2003: Image Processing. Proceedings of the SPIE.*, M. Sonka and J. Fitzpatrick, Eds., vol. 5032, May 2003, pp. 1839–1850.

[5] J. Fessler, E. Ficaro, N. Clinthorne, and K. Lange, "Grouped-coordinate ascent algorithms for penalized-likelihoodtransmission image reconstruction," *IEEE Trans Med Imaging*, vol. 16, no. 2, pp. 166–175, April 1997.

[6] 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, no. 5, pp. 2505–2512, October 2004.

[7] P. La Rivire, J. Bian, and P. Vargas, "Penalized-likelihood sinogram restoration for computed tomography." *IEEE Trans Med Imaging*, vol. 25, no. 8, pp. 1022–1036, 2006.

[8] J. Hsieh, "Adaptive streak artifact reduction in computed tomography resulting from excessive x-ray photon noise," *Medical Physics*, vol. 25, no. 11, pp. 2139–2147, November 1998.

[9] 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," *Medical Physics*, vol. 28, no. 4, pp. 475–490, April 2001.

[10] O. Demirkaya, "Reduction of noise and image artifacts in computed tomography by nonlinear filtration of projection images," in *Medical Imaging 2001: Image Processing. Proceedings of the SPIE.*, M. Sonka and K. M. Hanson, Eds., vol. 4322, July 2001, pp. 917–923.

[11] G.-F. Rust, V. Aurich, and M. Reiser, "Noise/dose reduction and image improvements in screening virtual colonoscopy with tube currents of 20 mAs with nonlinear Gaussian filter chains," in *Medical Imaging 2002: Physiology and Function from Multidimensional Images. Proceedings of the SPIE.*, A. V. Clough and C.-T. Chen, Eds., vol. 4683, April 2002, pp. 186–197.

[12] H. Lu, X. Li, L. Li, D. Chen, Y. Xing, J. Hsieh, and Z. Liang, "Adaptive noise reduction toward low-dose computed tomography," in *Medical Imaging 2003: Physics of Medical Imaging. Proceedings of the SPIE.*, M. J. Yaffe and L. E. Antonuk, Eds., vol. 5030, June 2003, pp. 759–766.

[13] F. Catte, P.-L. Lions, J.-M. Morel, and T. Coll, "Image selective smoothing and edge-detection by nonlinear diffusion," *SIAM Journal on Numerical Analysis*, vol. 29, no. 1, pp. 182–193, 1992.

[14] C. Tomasi and R. Manduchi, "Bilateral filtering for gray and color images," in *IEEE International Conference on Computer Vision*, Bombay, India, 1998, pp. 836–846, http://www.cse.ucsc.edu/˜manduchi/Papers/ICCV98.pdf.

[15] B. B. Hubbard, *Wavelets: Die Mathematik der kleinen Wellen*. Basel, Switzerland: Birkhäuser Verlag, 1997.

[16] D. L. Donoho and I. M. Johnstone, "Ideal spatial adaptation by wavelet shrinkage," *Biometrika*, vol. 81, no. 3, pp. 425–455, 1994.

[17] Y. Xu, J. B. Weaver, D. M. Healy, and J. Lu, "Wavelet transform domain filters: A spatially selective noise filtration technique," *IEEE Transactions on Image Processing*, vol. 3, no. 6, pp. 747–758, July 1994.

[18] F. Faghih and M. Smith, "Combining spatial and scale-space techniques for edge detection to provide a spatially adaptive wavelet-based noise filtering algorithm," *IEEE Transactions on Image Processing*, vol. 11, no. 9, pp. 1062–1071, September 2002.

[19] A. Pizurica, W. Philips, I. Lemahieu, and M. Acheroy, "A versatile wavelet domain noise filtration technique for medical imaging," *IEEE Transactions on Image Processing*, vol. 22, no. 3, pp. 1062–1071, March 2003.

[20] O. Tischenko, C. Hoeschen, and E. Buhr, "An artifact-free structure-saving noise reduction using the correlation between two images for threshold determination in the wavelet domain," in *Medical Imaging 2005: Image Processing. Proceedings of the SPIE.*, J. M. Fitzpatrick and J. M. Reinhardt, Eds., vol. 5747, April 2005, pp. 1066–1075.

[21] A. Kak and M. Slanely, *Principles of Computerized Tomographic Imaging*. Society for Industrial and Applied Mathematics, July 2001, http://www.slaney.org/pct/pct-toc.html.

[22] T. Buzug, *Einf"uhrung in die Computertomographie*. Berlin Heidelberg: Springer-Verlag, 2004.

[23] K. Stierstorfer, A. Rauscher, J. Boese, H. Bruder, S. Schaller, and T. Flohr, "Weighted FBP - a simple approximate 3DFBP algorithm for multislice spiral CT with good dose usage for arbitrary pitch," *Physics in Medicine and Biology*, vol. 49, no. 11, pp. 2209–2218, June 2004.

[24] F. Natterer, *The Mathematics of Computerized Tomography*. Stuttgart: Wiley and B.G. Teubner, 1986.

[25] W. A. Kalender, *Computed Tomography*. Erlangen: Publicis Corporate Publishing, 2000.

[26] H. Bruder, K. Stierstorfer, C. McCullough, R. Raupach, M. Petersilka, M. Grasruck, C. Suess, B. Ohnesorge, and T. Flohr, "Design considerations in cardiac CT," in *Medical Imaging 2006: Physics of Medical Imaging. Proceedings of the SPIE.* , M. Flynn and J. Hsieh, Eds., vol. 6142, March 2006, pp. 151–163.

[27] A. Borsdorf, R. Raupach, and J. Hornegger, " Wavelet based Noise Reduction by Identification of Correlation ," in *Pattern Recognition (DAGM 2006), Lecture Notes in Computer Science* , K. Franke, K. Müller, B. Nickolay, and R. Schäfer, Eds., vol. 4174. Berlin: Springer, 2006, pp. 21–30.

[28] A. Oppelt, Ed., *Imaging Systems for Medical Diagnostics*, 2nd ed. Erlangen: Publicis Corporate Publishing, 2005.

[29] S. G. Mallat, "A theory for multiresolution signal decomposition: The wavelet representation," *IEEE Transactions on Pattern Analysis and Maschine Intelligence*, vol. 11, no. 7, pp. 674–693, 1989.

[30] G. Strang and T. Nguyen, *Wavelets and Filter Banks*. Wellesley-Cambridge Press, 1996.

[31] I. Daubechies, *Ten Lectures on Wavelets*. Philadelphia: Society for Industrial and Applied Mathematics, 1992.

[32] R. R. Coifman and D. L. Donoho, "Translation-invariant de-noising," in *Lecture Notes in Statistics: Wavelets and Statistics*, vol. 103, 1995, pp. 125–150.

[33] G. P. Nason and B. W. Silverman, "The stationary wavelet transform and some statistical applications," in *Lecture Notes in Statistics: Wavelets and Statistics*, vol. 103, 1995, pp. 281–300.

[34] S. Mallat and S. Zhong, "Characterization of signals from multiscale edges," in *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 7, July 1992, pp. 710–732.

[35] E. W. Weisstein, "Correlation coefficient," 2006, from MathWorld - A Wolfram Web Resource, http://mathworld.wolfram.com/CorrelationCoefficient.html.

[36] K. Stierstorfer, T. Flohr, and H. Bruder, "Segmented multiple plane reconstruction: A novel approximate reconstruction scheme for multi-slice spiral CT," *Physics in Medicine and Biology*, vol. 47, no. 4, pp. 2571–2581, July 2002.

[37] I. A. Cunningham and B. K. Reid, "Signal and noise in modulation transfer function determinations using the slit, wire, and edge techniques," *Medical Physics*, vol. 19, no. 4, pp. 1037–1044, July 1992.

[38] T. Fawcett, "An introduction to roc analysis," *Pattern Recognition Letters*, vol. 27, no. 8, pp. 861–874, 2006.

[39] H. Bruder, T. Flohr, O. Sembritzki, and K. Stierstorfer, "Adaptive filtering of projection data acquired by medical diagnostic device involves adding filtered acquired data to acquired data, determining quantitative ratio of filtered and acquired data," 2001, patent DE1043484.