TRANSFORM-BASED DENOISING AND ENHANCEMENT IN MEDICAL X-RAY IMAGING

Til Aach
Institute for Signal Processing, Medical University of Lübeck,
Ratzeburger Allee 160, D-23538 Lübeck, GERMANY
Tel: +49 451 3909 556; fax: +49 3909 555
e-mail: aach@isip.mu-luebeck.de

ABSTRACT

Physical and clinical constraints of the imaging process often degrade the quality of medical images. A typical case is low-dose X-ray image acquisition, where the signal-to-noise ratio is limited due to X-ray quantum noise (photon-limited imaging). This paper discusses and quantitatively evaluates spectral transform-based methods for denoising and enhancement of such images. We examine algorithms based on the DFT, the Modulated Lapped Transform (MLT) and a new lapped transform termed the Lapped Directional Transform (LDT).

1 INTRODUCTION

The quality of non-invasively acquired images of inner parts of the human body is often limited by constraints on the imaging process. One example is live X-ray imaging, called X-ray fluoroscopy. This imaging modality is used to provide guidance during examinations or interventions. The resulting prolonged exposure of humans to harmful X radiation allows to apply only low dose rates, i.e. low dose per frame. Noise in the images observed on a CRT monitor is then mainly caused by the very low number of X-ray quanta contributing to the imaging process at each pixel. Known as "quantum noise" (QN), this noise is filtered by the imaging system's transfer function, and is therefore spatially coloured [6]. Furthermore, as QN is Poisson distributed, its noise variance is signal dependent [6]. An example image from a still human phantom sequence [1] is shown in Fig. 1. This sequence consists of 40 frames of size 512×512 pixels with eight bit resolution.

In this contribution we discuss and evaluate spectral amplitude estimation-based methods for denoising and enhancement of such images. Estimation of short time amplitude spectra is well-established in speech restoration, since it is well adapted to the short time stationarity of speech. Motivations to transfer this approach to filtering images [9] is their short-space stationarity, and the fact that the spatial correlations of QN are easily taken into account in the spectral domain [3, 4, 6].

[1] A phantom consists of parts of a human body cast in resin. The phantom images were acquired by an X-ray imaging system under clinically relevant settings, and allow representative measurements to be made.

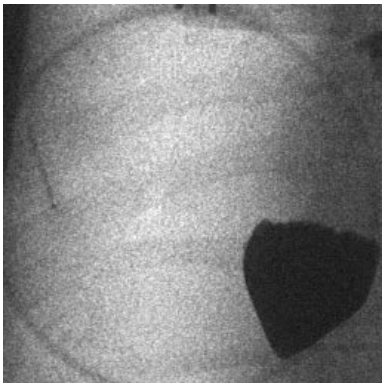

Figure 1: Part of size 256×256 pixels from a phantom image sequence.

In the following, we first review approaches to spectral amplitude estimation for images. We discuss several transforms, viz. the DFT, the Modulated Lapped Transform (MLT) [10], and a new lapped transform termed the Lapped Directional Transform (LDT) [8, 5]. We then compare their performances quantitatively using the phantom data.

2 SPECTRAL AMPLITUDE ESTIMATION

In spectral amplitude estimation, the input images are first subjected to a short space spectral analysis by means of a block transform. In each block spectrum, every noisy coefficient $G(k)$ is then attenuated depending on its observed signal to noise ratio (SNR) $r^2(k)$ according to

$$\hat{F}(k) = G(k) \cdot h(r(k)), \text{ where } r^2(k) = |G(k)|^2 / N(k) \tag{1}$$

with $N(k)$ denoting the k-th coefficient of the (estimated) noise power spectrum (NPS), and $\hat{F}(k)$ the estimate of the noisefree coefficient $F(k)$. In our algorithms, the *attenuation function* $h(r)$ was derived invoking a minimum mean square approach, yielding [6, 1]

$$h(r) = 1 + \lambda \cdot \exp\left\{-\frac{r^2(k)}{\alpha}\right\}^{-1} \tag{2}$$

The parameter λ depends on the (unknown) signal power as well as on $N(k)$, and on the a priori probability that $G(k)$ represents noise only (see [1, 5]). Here, it is regarded as a parameter which is preset manually. Similarly, α is a weighting factor like the one used in generalized Wiener filters, which is also preset. The processed block spectra are then subjected to an inverse transform, and the restored image is assembled from the filtered blocks.

2.1 DFT-Based Spectral Image Restoration

The DFT is perhaps the most common transform for spectral analysis. Before applying the DFT, each block must be multiplied by a smooth window to avoid spectral distortions. Perfect reconstruction from block spectra then requires that the blocks overlap. At the same time, overlap prevents block artifacts from occurring. In our algorithm, we use a cosine shaped Hanning window, which requires that the blocks overlap by half in each dimension to permit perfect reconstruction by adding the overlapping blocks.

A downside of the DFT is that the large block overlaps increase the data volume to be processed fourfold. On the other hand, it is well known that the DFT performs well with respect to spatially oriented patterns: in the spectral domain, the energy of oriented structures is concentrated along the line perpendicular to the spatial orientation which passes through the origin [7]. As oriented patterns generally consist of visually important lines or edges, their spectral components should be subjected to less attenuation than other ones, or even be enhanced. To detect orientation in our spectral approach, we calculate a 2×2 inertia-like matrix from each block spectrum. The eigenvalues of this matrix indicate whether or not orientation is present, while the line along which spectral energy concentrates best is determined by an eigenvector [2, 5]. Coefficients close to or lying on this line are then attenuated more carefully than other ones, where the underlying assumption is that these coefficients contribute to signal even when their observed SNR is low. The attenuation function hence becomes a family of attenuation functions, which depends not only on the SNR, but also on the angle δ between the detected line and a processed coefficient's location (Fig. 2). Moreover, a tunable anisotropic enhancement filter can be applied directly in the spectral domain to enhance these coefficients contributing to spatial orientation (see Fig. 3).

2.2 MLT-Based Spectral Image Restoration

To avoid the increase in the processed data volume caused by the DFT-window, the Discrete Cosine Transform (DCT) could be used alternatively [3, 1]. Perfect reconstruction then does not anymore require overlapping blocks. The filtered images, however, tend to show the block raster. To prevent both data redundancy and block artifacts, we therefore investigate here the use of

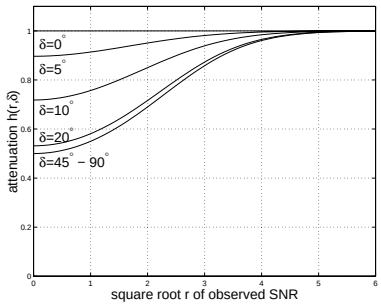

Figure 2: Orientation-adaptive attenuation functions for $\lambda = 1$, $\alpha = 5$. The attenuation function for $\delta = 45° - 90°$ correponds to the non-adaptive attenuation function $h(r)$ in Eq. (2).

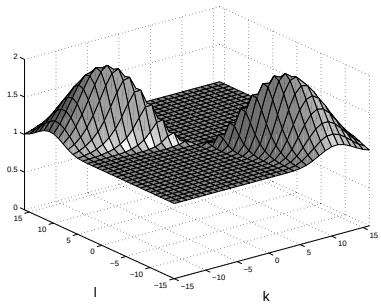

Figure 3: Anisotropic adaptive enhancement filter for a DFT block size of 32×32 pixels. An ideal orientation of $45°$ was assumed.

the MLT. Like the DCT, the MLT is a real-valued separable transform. Its basis functions, however, possess two distinct advantages over DCT basis functions [10]:

First, when decomposing an image into blocks, the basis functions overlap by half their size without creating redundancy. More specifically, an image is decomposed by the MLT into blocks of size $2L \times 2L$, which corresponds to the size of the MLT basis functions. The overlap between neighbouring blocks is L in each dimension. The block samples are transformed into a local spectrum of size $L \times L$. Perfect reconstruction is achieved when adding the overlapping blocks obtained after the inverse transform ("time domain alias cancellation").

Secondly, the MLT basis functions decay smoothly to zero toward their boundaries (see Fig. 4). Together with the intrinsic overlap, the MLT therefore avoids both data redundancy and block artifacts.

Unlike the DFT (and like the DCT), the MLT does not allow to unambiguously identify spatial orientation from modulus or energy spectra. The reason for this is the separability of the real-valued basis functions, resulting in fourfold symmetric modulus transfer functions when viewing the basis functions as filters. Unless oriented vertically or horizontally, a MLT basis function is hence sensitive to two different orientations, what can also be appreciated from Fig. 4.

2.3 LDT-Based Spectral Image Restoration

To combine the advantages of the MLT with the orientation-detection capabilities of the DFT, we have developed a new lapped transform called the Lapped Directional Transform (LDT). The LDT basis functions are still real-valued, but, to circumvent the constraint of fourfold symmetric modulus transfer functions, they are not separable [8, 5]. The modulus transfer functions of the LDT basis functions are then symmetric with respect to the origin of the spectral coordinates. To create the LDT basis functions, the MLT and a slightly modified MLT denoted MLT′ are required. The basis functions of the MLT′ are obtained from MLT basis functions by shifting the oscillating sinusoids by $\pi/2$ in each spatial dimension, i.e. by replacing cosine terms by sine terms. This is also illustrated in Fig. 4. From a MLT basis function and MLT′ basis function of identical spectral indices k, two LDT basis functions are formed by the taking the sum and difference of these (see Fig. 4). Clearly, each basis function is sensitive to only one orientation. Note also that, by definition, the integer-valued spatial frequency indices of the MLT differ from the actual frequency of the basis functions by one half. We have therefore shifted the spatial frequency indices of the LDT by one half in each dimension, giving noninteger indices, which, however, correspond again to the frequencies of the sinusoids. Although not separable, the LDT can hence be calculated from two real-valued separable lapped transforms, for which fast algorithms exist. Invoking two lapped transforms implies doubling the processed data volume. To invert the LDT, each pair of LDT basis functions is straightforwardly reconverted to MLT and MLT′ basis functions. After applying both inverse MLT and MLT′, we take the average of the resulting two blocks [8]. Alternatively, one might discard either the MLT or MLT′ coefficients, and use only the remaining coefficients for the inverse transform [5].

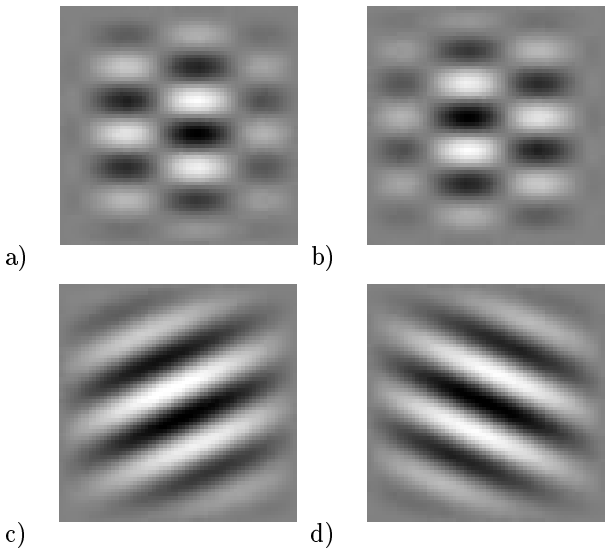

Figure 4: Examples of 2D MLT and MLT′ basis function, and the pair of LDT basis functions calculated from these. a) basis function for coefficient $(1, 3)$ of the MLT, b) of the MLT′, c) LDT basis function $(1.5, 3.5)$, d) LDT basis function $(1.5, -3.5)$. Note the shift of the LDT frequency indices by 1/2.

Due the use of pairs of complementary MLT basis functions, the LDT energy spectra are unambiguously related to spatial orientation. Orientation can hence be detected in the same manner as in DFT energy spectra using the inertia matrix approach, and the spectral amplitude estimation algorithm adapted accordingly.

3 The Noise Model

To calculate the observed SNR $r^2(k)$ needed to determine the attenuation in Eqs. (1) and (2), estimates of the spectral noise power are necessary. In fluoroscopy imaging, quantum noise generally outweighs system noise sources (*quantum-limited imaging*, [6]). Each spectral noise power coefficient $N(k)$ then depends on two quantities: its spatial frequency index k, and the signal intensity I. We separate the influence of these two quantities according to [3]

$$N(k) = \sigma^2(I) \cdot P(k) \tag{3}$$

where $P(k)$ is the normalized spectral noise power describing the *shape* of the distribution of spectral noise power over k. The shape is scaled by the absolute noise power $\sigma^2(I)$ to obtain $N(k)$. The dependence of the noise power σ^2 on the signal intensity I is determined by Poissonian distribution of QN as well as by the (potentially nonlinear) conversion curve from X-ray intensity to signal intensity. We consider both the noise power function $\sigma^2(I)$ and the shape $P(k)$ to be known for images acquired by a given system with known acquisition parameter settings. Fig. 5 shows $\sigma^2(I)$ for the phantom image sequence in Fig. 1, as well as for processing results below. For the original, these data were estimated by first calculating the average of all 40 images, and subtracting the average frame from all others. The resulting images contain almost only noise. To estimate the noise power function, the intensity range in the average frame was divided into intervals of 4 quantization steps. Estimation of the noise power function $\sigma^2(I)$ for each interval is then straightforward. To estimate $P(k)$, each pixel m in the difference sequence was normalized by its estimated standard deviation $\sigma(I(m))$, which is easily obtained from the noise power function after crossreferencing the pixel position m with the average frame to obtain $I(m)$. Estimation of $P(k)$ from the normalized noise sequence again is straightforward. Of course, the same transform (DFT, MLT, or LDT) and window function, if required, and the same basis function size (32×32 in our experiments) as later in the spectral amplitude estimation algorithm must be used. Let us finally note that during processing, the noise power function is avail-

able as a LUT which is accessed via the DC transform coefficient, i.e. via the mean intensity in a block.

4 Results

Fig. 5 shows the noise power functions measured from the phantom image sequence in Fig. 1 after processing by the DFT-based, MLT-based and LDT-based algorithms. In the DFT-based and LDT-based algorithms, the orientation-dependent attenuation functions in Fig. 2 were applied, but no enhancement. Parameters were $\alpha = 5$ and $\lambda = 1$. The measured noise power functions are rather close to each other, with the DFT performing best, and LDT and MLT slightly worse. Quantitatively, at a grey level of 90, the orientation-adaptive DFT and LDT algorithms attenuate the noise power by 3.6 dB and 3 dB, respectively, while the MLT algorithm results in 2.9 dB attenuation. To relate noise attenuation to loss of signal, we also calculated the MSE between the average frame of each processed sequence and the average frame of the original sequence. In terms of PSNR, the MSE is 51.9 dB, 52.7 dB, and 52 dB for the DFT-based, LDT-based and MLT-based algorithm, respectively. In total, orientation adaptive DFT and LDT algorithms perform best, while the MLT-based algorithm is some 0.7 dB worse. Applying directional edge enhancement (Fig. 3) in the DFT-based and LDT-based algorithms reduces the achieved noise attenuation by 0.1-0.3 dB. A processing result is given in Fig. 6.

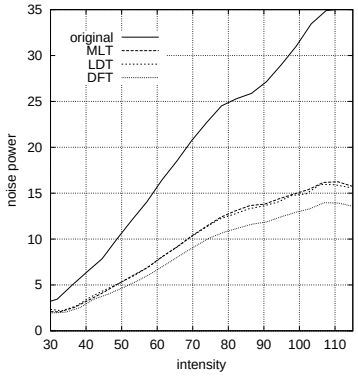

Figure 5: Measured noise power curves before and after filtering for the phantom image sequence. The (almost) linear rise of noise power over intensity is caused by the Poissonian nature of QN. The dynamic range of the data is $[20, 130]$.

5 CONCLUSIONS

We have compared the performances of DFT, MLT and a new lapped transform termed the LDT in local spectral analysis algorithms for filtering medical images. Results reveal that transforms with good energy compression of oriented structures, i.e. the DFT and the LDT, are at an advantage. Compared to the MLT, the downside of using these transforms is the increase in the data volume to be processed, which is fourfold for the DFT

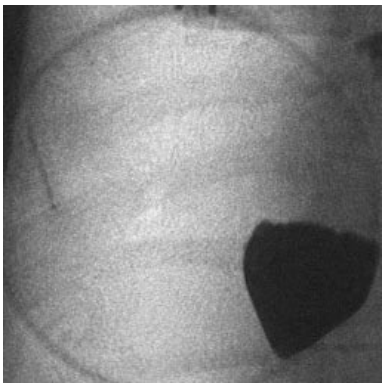

Figure 6: Fig. 1 filtered by the LDT-based algorithm with orientation-adaptive attenuation.

in connection with a standard 2D Hanning window, and twofold for the LDT. Also, spectral orientation detection and adaptation adds to the computational load. Doing so, however, allows more noise reduction while preserving visually important lines and edges better, or even to enhance these.

References

[1] T. Aach, "Spectral transform-based nonlinear restoration of medical images: Algorithms and a comparative evaluation", SPIE Vol. 3646, pp. 270–280, 1999.

[2] T. Aach, D. Kunz, "Anisotropic spectral magnitude estimation filters for noise reduction and image enhancement", Proc. IEEE ICIP-96, Lausanne, Switzerland, Sept. 16–19 1996, pp. 335–338.

[3] T. Aach, D. Kunz, "Spectral estimation filters for noise reduction in x-ray fluoroscopy imaging", Proc. EUSIPCO-96, Trieste, Italy, Sept. 10–13 1996, pp. 571–574.

[4] T. Aach, D. Kunz, "Spectral amplitude estimation-based x-ray image restoration: An extension of a speech enhancement approach", Proc. EUSIPCO-98, Rodos, Greece, Sept. 5–9 1998, pp. 323–326.

[5] T. Aach, D. Kunz, "A lapped directional transform for spectral image analysis and its application to restoration and enhancement", Signal Processing (to appear, 2000).

[6] T. Aach, U. Schiebel, G. Spekowius, "Digital image acquisition and processing in medical x-ray imaging", Journ. Electro. Im. 8:7–22, 1999.

[7] J. Bigün, G. H. Granlund, "Optimal orientation detection of linear symmetry", Proc. IEEE 1st ICCV, London, UK, June 1987, pp. 323–326.

[8] D. Kunz, T. Aach, "Lapped directional transform: A new transform for spectral image analysis", Proc. ICASSP-99, Phoenix, AZ, USA, Mar. 15–19 1999, pp. 3433–3436.

[9] J. S. Lim, "Image restoration by short space spectral subtraction", IEEE T-ASSP 28(2):191–197, 1980.

[10] H. S. Malvar, "Lapped transforms for efficient transform/subband coding", IEEE T-ASSP 38(6):969–978, 1990.