Lehrstuhl für Bildverarbeitung
Institute of Imaging & Computer Vision

A Lapped Directional Transform for Spectral Image Analysis and Its Application to Restoration and Enhancement

Til Aach and Dietmar Kunz

in: Signal Processing. See also BIBTEX entry below.

BIBTEX:

```
@article{AAC00a,
author = {Til Aach and Dietmar Kunz},
title = {A Lapped Directional Transform for Spectral Image Analysis and
   Its Application to Restoration and Enhancement},
journal = {Signal Processing},
        publisher       = {SPIE},
volume = {80},
number = {11},
year = {2000},
pages = {2347--2364}}
```


document created on: December 20, 2006
created from file: sphdl65coverpage.tex
cover page automatically created with CoverPage.sty
(available at your favourite CTAN mirror)

RWTH AACHEN
RHEINISCH-WESTFÄLISCHE TECHNISCHE HOCHSCHULE AACHEN

ELSEVIER

Signal Processing 80 (2000) 2347–2364

SIGNAL PROCESSING

www.elsevier.nl/locate/sigpro

A lapped directional transform for spectral image analysis and its application to restoration and enhancement☆

Til Aach[a,*], Dietmar Kunz[b,1]

[a]*Institut für Signalverarbeitung und Prozessrechentechnik, Medizinische Universität zu Lübeck, Ratzeburger Allee 160, D-23538 Lübeck, Germany*

[b]*Philips GmbH Forschungslaboratorien, Postfach 50 01 45, D-52085 Aachen, Germany*

Received 6 August 1999

Dedicated to Professor H.D. Lüke on the occasion of his 65th birthday

Abstract

We describe a new real-valued lapped transform for 2D-signal and image processing. Lapped transforms are particularly useful in block-based processing, since their overlapping basis functions reduce or prevent block artifacts. Our transform is derived from the modulated lapped transform (MLT), which is a real-valued and separable transform. Like the discrete cosine transform, the MLT does not allow to unambiguously identify spatial orientation from modulus spectra or spectral energy. This is in marked contrast to the complex-valued discrete Fourier transform (DFT). The new lapped transform is real valued, and at the same time allows unambiguous detection of spatial orientation from spectral energy. Furthermore, a fast and separable algorithm for this transform exists. As an application example, we investigate the transform's performance in anisotropic spectral approaches to image restoration and enhancement, and compare it to the DFT.

Zusammenfassung

Wir beschreiben eine neue reellwertige überlappende Transformation für 2D-Signal- und Bildverarbeitung. Überlappende Transformationen sind besonders geeignet für blockweise Verarbeitung, da sie Blockartefakte reduzieren oder vermeiden. Die neue Transformation wird aus der reellwertigen und separierbaren Modulierten Überlappenden Transformation (MLT) abgeleitet. Im Gegensatz zur Diskreten Fourier-Transformation (DFT) erlauben die MLT wie auch die Diskrete Cosinus-Transformation (DCT) keine eindeutige Detektion von Orientierung allein auf der Basis spektraler Energieverteilung. Die neue überlappende Transformation ist reellwertig, und ermöglicht eindeutige Orientierungsdetektion aus der spektralen Energieverteilung. Ein schneller, separierbarer Algorithmus für die Transformation existiert. Die Transformation wird in spektralen Ansätzen zur Bildrestauration und Verbesserung eingesetzt, und mit der DFT verglichen.

☆Paper presented in part at the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), Phoenix, AZ, March 1999.

* Corresponding author. Tel.: + 49-451-3909-556; fax: + 46-451-3909-555.

E-mail address: aach@isip.mu-luebeck.de (T. Aach).

[1] Now with Department of Imaging and Media Technology, Cologne University of Applied Sciences, Betzdorfer Str. 2, D-50679 Köln, Germany.

0165-1684/00/$ - see front matter
PII: S0165-1684(00)00122-5

Résumé

Nous décrivons dans cet article une transformation réelle à recouvrement nouvelle pour le traitement de signaux 2D et d'images. Les transformations à recouvrement sont particulièrement utiles en traitement par blocs, car leurs fonctions de base à recouvrement réduisent ou empêchent les artefacts de bloc. Notre transformation est dérivée de la transformation à recouvrement modulée (MLT), qui prend des valeurs réelles et est séparable. De même que la transformation discrète en cosinus (DCT), la MLT ne permet pas d'identifier sans ambigüité l'orientation spatiale à partir des spectres d'amplitude ou de l'énergie spectrale. Ceci contraste avec la transformation de Fourier discrète (DFT), qui prend des valeurs complexes. La transformation à recouvrement nouvelle prend des valeurs réelles, et permet aussi une détection sans ambigüité de l'orientation spatiale à partir del'énergie spectrale. De plus, un algorithme rapide et séparable existe pour cette transformation. Pour donner un exemple d'application, nous étudions les performances de cette transformation dans des approches spectrales anisotropes de restauration et de rehaussement d'image, et les comparons à celles de la DFT.

Keywords: Lapped transform; Spatial orientation; Energy spectrum; Image restoration; Spectral amplitude estimation

1. Introduction

Calculation of block or short-space spectra from images and image reconstruction from processed block spectra are standard operations in many image processing tasks. Prime examples are image compression [14] and restoration by spectral amplitude estimation [3,6,21]. These methods rely on the ability of the spectral transform to concentrate signal energy into only a few coefficients. As shown in [1,2] for image restoration and enhancement, the performance of block spectra-based algorithms can be considerably improved if perceptually important information, like oriented lines and edges, can be detected and processed in a special manner. In this respect the discrete Fourier transform (DFT) is particularly advantageous since the presence of local orientation within a block results in concentration of spectral energy along the line perpendicular to spatial orientation and passing through the origin [2,11]. The drawback of the DFT, however, is that in order to avoid spurious high-frequency artifacts, each block must be windowed prior to transformation. Perfect reconstruction then requires overlapping blocks [16]. When using standard windows like the triangular Bartlett window [21] or the cosine-shaped Hanning window [2], the blocks overlap by half the block size in each dimension. Each pixel is hence part of four blocks, resulting in a fourfold increased data volume.

Unlike the DFT, the discrete cosine transform (DCT) does not require windowing. Perfect reconstruction of the unprocessed image hence avoids the necessity of overlapping blocks and redundancy. Additionally, DCT basis functions approximate eigenvectors of a first-order Markov process with reasonably high intersample correlation [8,15]. The DCT is therefore widely used in image compression techniques (JPEG, MPEG). The DCT basis functions, however, do not generally decay smoothly to zero at the block boundaries. They thus exhibit discontinuities when thought of as being padded by zeros beyond the boundaries. Independent processing of the DCT spectra of adjacent blocks therefore generates the well-known block artifacts.

One possibility to avoid block artifacts in transform-based processing is to use a window which moves pixel by pixel, as described for the DCT in [33,34]. The redundancy of this approach, however, is even higher than that of the above windowed DFT methods. In this paper, we therefore consider so-called lapped transforms like the lapped orthogonal transform (LOT) [27] or the modulated lapped transform (MLT) [24] to reduce or even prevent block artifacts. Essentially, these real-valued transforms yield a non-redundant image representation using overlapping basis functions. Fast algorithms for these transforms exist [24,25]. In addition, lapped transforms generally exhibit good energy compaction performance for the mentioned first-order Markov data model [10]. These transforms are therefore particularly attractive for compression, for instance of medical images [12].

Unlike the DFT, the discussed real-valued transforms do not allow unambiguous detection of oriented structures from spectral energy concentration. The reason for this is that the DCT is derived as the Fourier transform of vertically and horizontally mirrored blocks [22]. On the one hand, this avoids the spurious high-frequency artifacts of the DFT. On the other hand, spatial orientation is mirrored, too. A given spatial orientation can therefore not be distinguished from its mirrored counterpart when looking at spectral energy distribution.

More generally, this observation results from the fact that in 2D-signal processing, DCT, LOT and MLT are each separable into two 1D-transforms, like the DFT. Regarding a block transform as a decimated filter bank [9], the filters are separable and real valued. Their modulus transfer functions are therefore fourfold-symmetric about the origin, as shown in Fig. 1 for the DCT, and in Fig. 2 for the MLT. Hence, a filter is equally sensitive for two orientations, unless it is horizontally or vertically oriented.

The complex-valued Fourier basis functions do not exhibit this symmetry (see Fig. 3). Accordingly, complex-valued lapped transforms have been developed, albeit not explicitly with the aim of unambiguous orientation detection in mind. In [35],

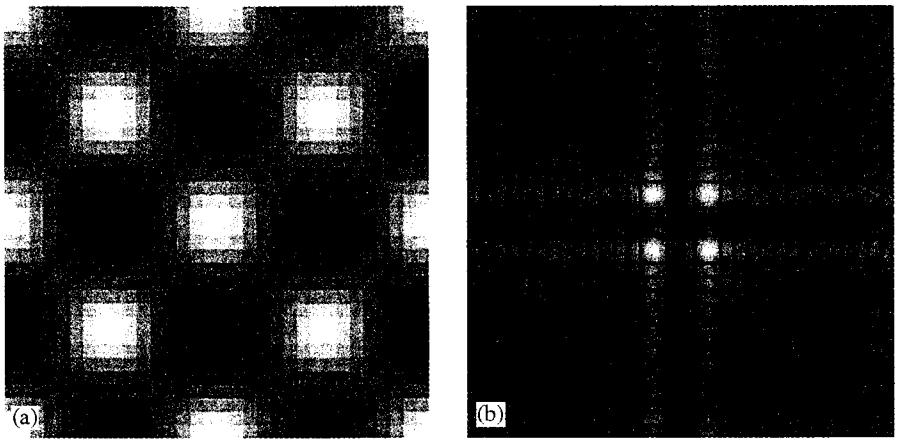

Fig. 1. DCT basis function of size 32×32 pixels for the spatial frequency coefficient (4,4) (a), and quasi-continuous centered log-modulus transfer function (b). The transfer function was calculated by using the DFT after zero-padding the basis function to a size of 256×256. The filter corresponding to the basis function clearly is equally sensitive to two different orientations. The spectral sidelobes are caused by the discontinuities at the block boundary.

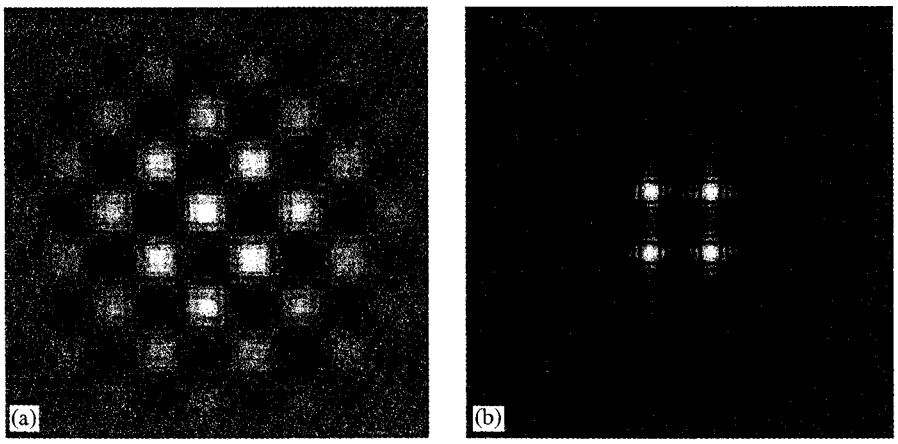

Fig. 2. MLT basis function of size 64×64 pixels for coefficient (4,4) (a), and quasi-continuous log-modulus transfer function (b), calculated as in Fig. 1. As before, the filter corresponding to the basis function is equally sensitive to two different orientations. Spectral sidelobes are reduced by the smooth window.

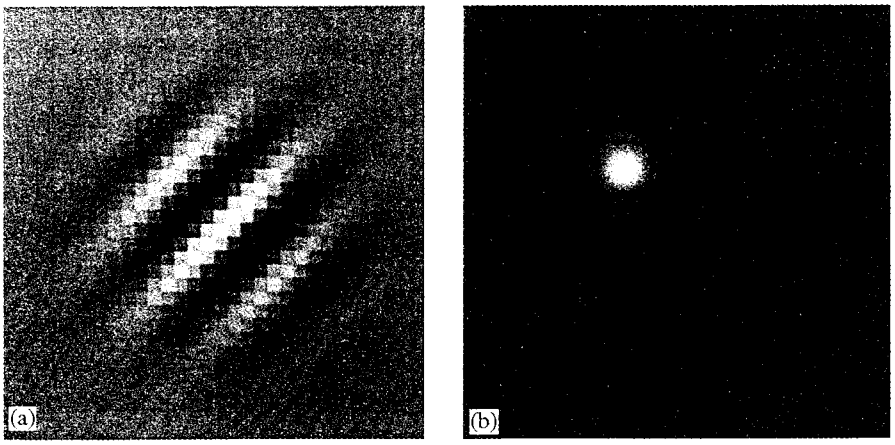

Fig. 3. Real part of windowed DFT basis function of size 32×32 pixels for coefficient (4,4) (a), and quasi-continuous log-modulus transfer function (b), calculated as in Fig. 1. Because of the complex basis function, the spectrum exhibits no symmetry about the origin.

a complex lapped transform (CLT) is developed by extending the LOT. It is used for motion estimation by evaluating cross correlation in the spectral domain. The specific aim is to avoid artifacts from the DFT-based circular convolution without zero padding the input data. Like the windowed DFT, the 1D-CLT is redundant by a factor of two. A fast algorithm, based on the FFT, is given. The 2D-CLT is separable, and redundant by a factor of four. Approximate LOT-coefficients can be obtained from real and imaginary part of the CLT.

A complex extension of the MLT for 1D-signals is described in [26]. Termed the modulated complex lapped transform (MCLT), it improves performance of spectral subtraction and echo cancellation algorithms. Unlike the MLT, the MCLT results in a diagonal matrix when direct and inverse transform are cascaded [26]. Time-domain aliasing cancellation in the overlapping blocks, as required by the MLT, is therefore not needed. While time-domain aliasing from unprocessed blocks cancels exactly, this is generally not the case for independently processed MLT-transformed blocks. The resulting artifacts limit the performance of e.g. echo cancellation. The MCLT avoids these artifacts. As applied to 1D-signals, it is redundant by a factor of two. A fast algorithm based on the DCT-IV and the discrete sine transform-IV is also provided.

In this paper, we derive a new lapped transform for 2D-signals which permits unambiguous orientation detection from modulus or energy spectra. Unlike the CLT and MCLT, our transform is real valued, but not separable. The central idea is to perform two MLTs with different but in some sense complementary basis functions, and to combine the resulting spectra. The transform results in only a twofold increased data volume, and is hence less redundant than the windowed DFT and the CLT in [35]. The fact that the basis functions of the new transform – named the lapped directional transform (LDT) – are not separable is no serious drawback since the LDT is computed from the separable MLTs by a simple butterfly operation. Another consequence of using MLTs is that the basis functions are smooth and overlapping, thus avoiding block artifacts. A first description of the LDT was presented in [20].

In the remainder of this paper, we first review the MLT, and then derive the LDT. We then describe applications of the LDT to the restoration and enhancement of noisy images by anisotropic spectral amplitude estimation before finishing with conclusions.

2. Modulated lapped transform

The 1D-MLT decomposes a signal into overlapping blocks of $2L$ samples [24]. The overlap between adjacent blocks is L samples so that each sample belongs to two blocks.

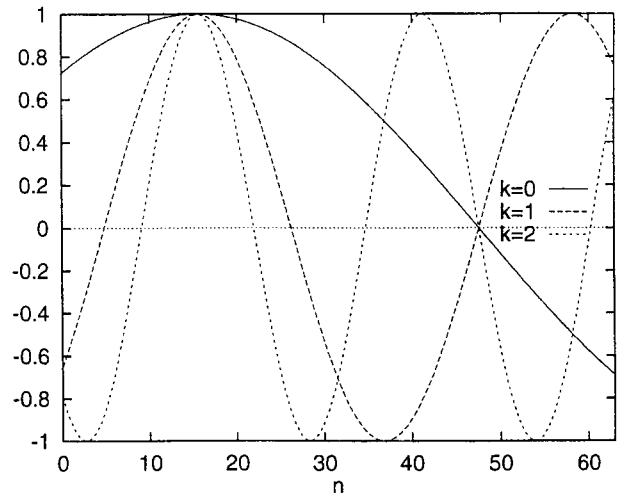

Fig. 4. Sinusoid components of 1D-MLT basis functions for $k = 0,1,2$, and $L = 32$.

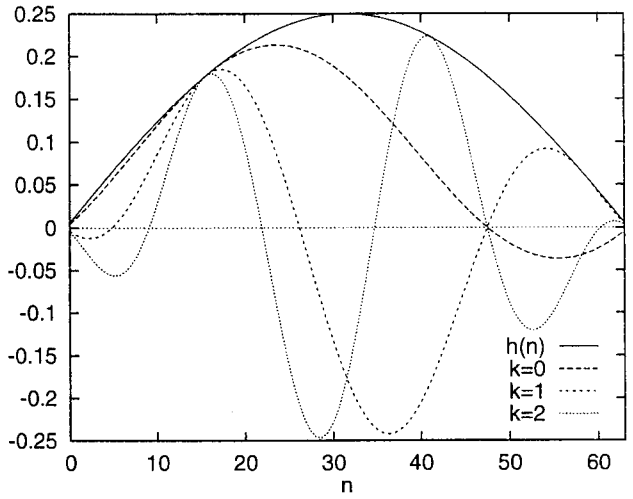

Fig. 5. Modulated 1D-MLT basis functions and modulating window $h(n)$.

Let $s(n)$, $n = 0,\ldots,2L-1$ denote the $2L$ samples in a block. The real-valued MLT basis functions are based on (co-)sine waves being symmetric about $(L-1)/2$ and anti-symmetric about $(3L-1)/2$, as shown in Fig. 4. Hence, there are only L different functions satisfying both symmetry constraints. These functions are modulated by a sine-shaped window function with period $4L$ which is symmetric about the block centre $(2L-1)/2$. Denoting the window by $h(n)$, it is given by[2]

$$h(n) = \sqrt{\frac{2}{L}}\sin\left[\frac{\pi}{2L}\left(n+\frac{1}{2}\right)\right], \quad n = 0,\ldots,2L-1. \tag{1}$$

The transform is then defined by a $(L, 2L)$-matrix with entries

$$w(k,n) = h(n)\cdot\cos\left[\frac{\pi}{L}\left(n-\frac{L-1}{2}\right)\left(k+\frac{1}{2}\right)\right],$$
$$k = 0,\ldots,L-1,\ n = 0,\ldots,2L-1. \tag{2}$$

Note that with this definition, there is an offset of $\frac{1}{2}$ between the spectral index k of $w(k,n)$ and the actual frequency $k+\frac{1}{2}$ of the cosine. These basis functions are depicted in Fig. 5. The basis functions $p(k,l,n,m)$ of the 2D-MLT obey

$$p(k,l,n,m) = w(k,n)\cdot w(l,m), \tag{3}$$

where we have assumed that the basis functions are defined over a quadratic support of $2L\times 2L$ samples. The indices k,l denote spatial frequencies, and n,m spatial coordinates. The MLT coefficients $S(k,l)$ for a 2D-signal block $s(n,m)$ are then given by

$$S(k,l) = \sum_{n=0}^{2L-1}\sum_{m=0}^{2L-1} s(n,m)p(k,l,n,m) \tag{4}$$

for $k,l = 0,\ldots,L-1$. When processing images, 2D localized spectra can thus be obtained by cascading a 1D-MLT along rows and another 1D-MLT along columns, or vice versa. An example basis function was shown in Fig. 2.

3. Lapped directional transform (LDT)

For the LDT we need another MLT, where we exchange positions of the symmetry and anti-symmetry constraint in each basis function. The anti-symmetry constraint is now positioned at $(L-1)/2$, and the symmetry constraint at $(3L-1)/2$. We denote this transform – which corresponds to

[2] These basis functions differ by a sign from the ones in [24]. The underlying reasoning, however, remains unchanged.

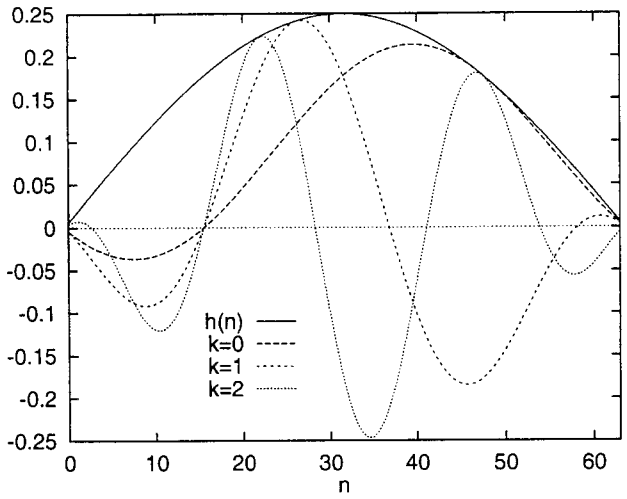

Fig. 6. Basis functions of 1D-MLT′ for $k = 0,1,2$ and modulating window $h(n)$.

a time-inverted MLT except for a sign – by MLT′. The basis functions of the 1D-MLT′ are given by

$$w'(k,n) = h(n)\cdot\sin\left[\frac{\pi}{L}\left(n-\frac{L-1}{2}\right)\left(k+\frac{1}{2}\right)\right],$$

$$k = 0,\ldots,L-1,\ n = 0,\ldots,2L-1. \tag{5}$$

Examples of these basis functions are shown in Fig. 6.

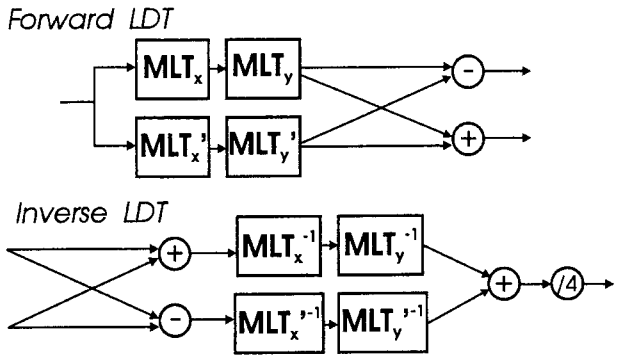

Fig. 7. Block diagram of the LDT and its inverse. The forward transform is calculated from MLT (Eq. (4)) and MLT′ (Eq. (7)), each of which is applied as a cascade of 1D-transforms along rows and columns. LDT coefficients are obtained from sums and differences of the MLT coefficients and MLT′ coefficients.

The basis functions of the 2D-MLT′ are then given by the product

$$p'(k,l,n,m) = w'(k,n)\cdot w'(l,m) \tag{6}$$

and MLT′-coefficients of the 2D-signal $s(n,m)$ are calculated by

$$S'(k,l) = \sum_{n=0}^{2L-1}\sum_{m=0}^{2L-1} s(n,m)\cdot p'(k,l,n,m). \tag{7}$$

Taken alone, neither of the local spectra $S(k,l)$ and $S'(k,l)$ can distinguish, based on spectral energy, between mirrored orientations. The central idea of our new transform is to compute both MLT and MLT′, and to obtain new spectra from the sum and the difference of these, as shown in Fig. 7. We will show that taken together, spectral energy then distinguishes between mirrored orientations, thus permitting unambiguous orientation detection from spectral energy distribution.

For every block of $2L\times 2L$ samples $s(n,m)$ of an image, there are $L\times L$ coefficients each for the MLT and the MLT′. We combine these to LDT-coefficients $Y(i,j)$ by

$$Y(i,j) = \begin{cases} S(i-1/2,j-1/2) - S'(i-1/2,j-1/2) & \text{if } i>0, j>0,\\ S(i-1/2,-j-1/2) + S'(i-1/2,-j-1/2) & \text{if } i>0, j<0,\\ Y(-i,-j) & \text{if } i<0, \end{cases} \tag{8}$$

where $i,j = -L+\frac{1}{2}, -L+\frac{3}{2},\ldots,\frac{-1}{2},\frac{1}{2},\ldots,L-\frac{3}{2}, L-\frac{1}{2}$. Note that, compared to the previous equations, we have here shifted the spatial frequency index of each dimension by $\frac{1}{2}$, resulting in noninteger indices i,j. Unlike in (2), the indices i,j of $Y(i,j)$ do hence correspond again to the spatial frequencies of the sinusoids (see (10)). Moreover, the above defined local spectrum is symmetric with respect to the origin rather than to $(\frac{1}{2},\frac{1}{2})$. Formally, this spectrum consists of $4L^2$ real coefficients, which due to the symmetry result in $2L^2$ independent coefficients. In total, this representation therefore contains twice as many coefficients as the original image.

A closer look at the LDT's basis functions reveals its ability to identify orientation. Writing the inner product to calculate the LDT-coefficients from $s(n,m)$ as the sum

$$Y(i,j) = \sum_{n=0}^{2L-1} \sum_{m=0}^{2L-1} q(i,j,n,m)s(n,m), \tag{9}$$

we obtain for the LDT basis function $q(i,j,n,m)$ for $i,j > 0$

$$\begin{aligned} &q(i,j,n,m) \\ &= w(i-\tfrac{1}{2},n)w(j-\tfrac{1}{2},m) - w'(i-\tfrac{1}{2},n)w'(j-\tfrac{1}{2},m) \\ &= h(n)h(m)\cdot\left\{\cos\left[\frac{\pi}{L}i\left(n-\frac{L-1}{2}\right)\right]\right. \\ &\qquad \times\cos\left[\frac{\pi}{L}j\left(m-\frac{L-1}{2}\right)\right] \\ &\qquad \left.-\sin\left[\frac{\pi}{L}i\left(n-\frac{L-1}{2}\right)\right]\sin\left[\frac{\pi}{L}j\left(m-\frac{L-1}{2}\right)\right]\right\} \\ &= h(n)h(m)\cos\left[\frac{\pi}{L}\left(i\left(n-\frac{L-1}{2}\right)+j\left(m-\frac{L-1}{2}\right)\right)\right]. \end{aligned} \tag{10}$$

The last line of (10) is an orientated cosine wave of spatial frequency i,j, which is multiplied by the separable 2D-window function $h(n)h(m)$. A similar relation holds for $i > 0, j < 0$. These directional basis functions are shown in Fig. 8 in comparison to their non-directional counterparts from the MLT and MLT′. Unlike the MLT and MLT′ basis functions, the 2D-LDT basis functions are not separable into real-valued 1D-basis functions. The modulus transfer function of a LDT basis function is symmetric about the origin, but not fourfold symmetric (Fig. 9). The LDT therefore allows to

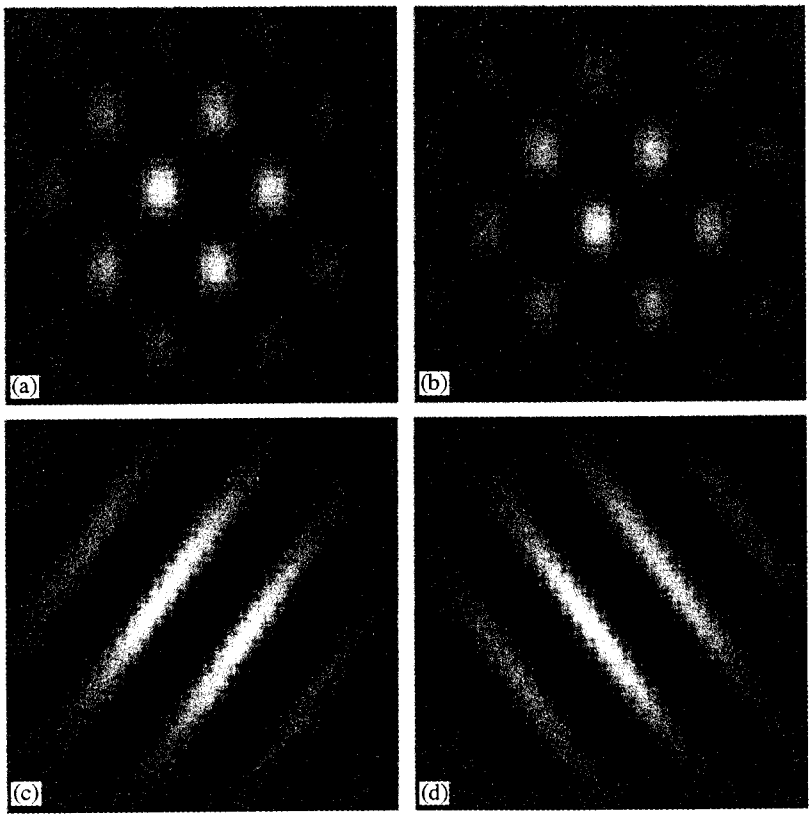

Fig. 8. Example 2D-MLT and LDT basis functions of size 64 × 64 pixels, i.e. $L = 32$: (a) basis function for coefficient (3,2) of the MLT, (b) of the MLT′, (c) LDT basis function (3.5,2.5), (d) LDT basis function (3.5, − 2.5).

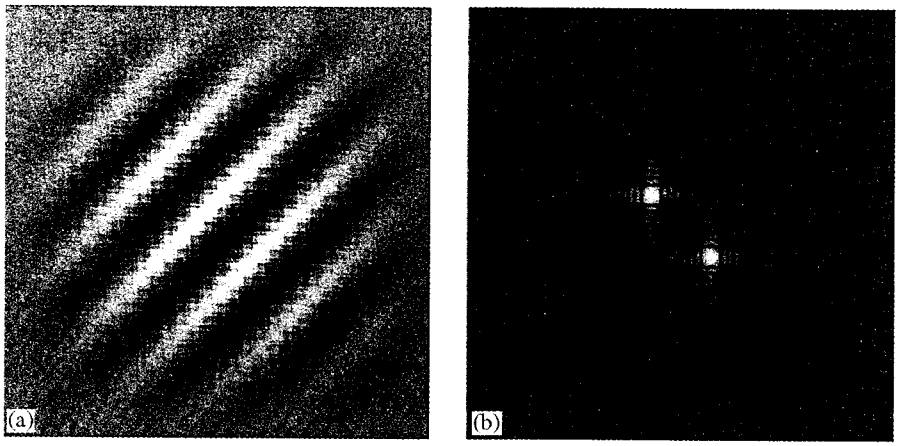

Fig. 9. LDT basis function for coefficient (4,4) and $L = 32$ (a), and its quasi-continuous log-modulus transfer function.

uniquely identify orientation from energy or absolute spectra.

The LDT can easily be inverted since both MLT and MLT′ coefficients can be recovered from the LDT coefficients as follows:

$$S(k,l) = \tfrac{1}{2}[Y(k+\tfrac{1}{2}, l+\tfrac{1}{2}) + Y(k+\tfrac{1}{2}, -l-\tfrac{1}{2})], \quad (11)$$

$$S'(k,l) = \tfrac{1}{2}[-Y(k+\tfrac{1}{2}, l+\tfrac{1}{2}) + Y(k+\tfrac{1}{2}, -l-\tfrac{1}{2})]. \quad (12)$$

In principle, one of these data sets – $S(k,l)$ or $S'(k,l)$ – is enough for perfect image reconstruction. Therefore, when LDT-based spectral processing is to be combined with MLT-based encoding, it suffices to calculate either MLT-coefficients $S(k,l)$ by (11) or MLT′-coefficients $S'(k,l)$ by (12) from the processed LDT-coefficients, which are then fed to the codec. In this case, the redundancy of the LDT does not affect the compression stage (cf. for the MCLT in audio processing [26]).

The drawback of this approach to reconstruction is that orientation of the remaining basis functions is ambiguous, which may produce artifacts when weighted by a large coefficient. In the following, we therefore reconstruct the processed image with respect to the oriented basis functions by calculating both the inverse MLT and inverse MLT′, and taking the average of these. This is also illustrated in Fig. 7. The final division by four accommodates the division by two in Eqs. (11) and (12), and the averaging operation.

4. Image restoration by local spectral analysis

4.1. Spectral amplitude estimation

We have applied the LDT within an algorithm for image restoration, which is based on analysing local spectra of the image data using overlapping windows. The central idea is to estimate the amplitudes of the spectral coefficients of a block from the noisy observations $Y(k,l)$, and to reconstruct the restored data from these estimates. Spectral analysis is carried out by block transforms like DFT, DCT, MLT, or by filter banks. Known as spectral amplitude (or magnitude) estimation, this approach is widely used in speech processing [13,32], and more recently also in image processing [3,5,6,21]. Estimation exploits the following properties of spectral transforms: first, we view the transform coefficients as uncorrelated. This assumption is motivated by the similarity of Fourier-type transforms to the Fourier series expansion, which decorrelates the analysed signal if the analysis interval exceeds the correlation length of the transformed signal [31]. Furthermore, each transform coefficient is a weighted sum of the intensities in the processed block. Invoking the central limit theorem, we regard the coefficients of real transforms as Gaussian distributed, and of complex transforms as complex Gaussian distributed (cf. for speech [17,28]). The coefficients can then be assumed to be statistically independent, so that the amplitude estimate can be determined separately for each coefficient. Starting

from a minimum mean-square-error (MMSE) approach, and denoting an uncorrupted coefficient by $F(k,l)$, it is estimated by $\hat{F}(k,l) = E[F(k,l)|Y(k,l)]$. Since a spectral transform compresses the energy of correlated image samples into only a few coefficients [30], we assume for each block that the uncorrupted image signal appears in only a fraction of the transform coefficients, with the rest representing noise. The conditional expectation can then be expressed as

$$\hat{F}(k,l) = \sum_{i=0}^{1} E[F(k,l)|Y(k,l), H_i(k,l)] \times P[H_i(k,l)|Y(k,l)]. \quad (13)$$

Here, $H_1(k,l)$ denotes the hypothesis that the observation $Y(k,l)$ contains signal and noise, and $H_0(k,l)$ the hypothesis that $Y(k,l)$ contains noise only (null hypothesis). The probability $P(H_i(k,l)|Y(k,l))$ is the conditional probability of hypothesis $H_i(k,l)$ given $Y(k,l)$. Clearly, we have $E[F(k,l)|Y(k,l), H_0(k,l)] = 0$. Furthermore, without any additional prior knowledge about $F(k,l)$ except that it is corrupted by additive zero mean Gaussian noise, we have $E[F(k,l)|Y(k,l), H_1(k,l)] = Y(k,l)$. Eq. (13) then reduces to $\hat{F}(k,l) = Y(k,l) \cdot P[H_1(k,l)|Y(k,l)]$. The noisy observation is hence attenuated the more, the lower its conditional probability is to contain signal. As the attenuation is a probability, it is always non-negative and real valued.

Applying Bayes' theorem to the conditional probability $P[H_1(k,l)|Y(k,l)]$, and inserting the above Gaussian assumptions, we obtain the following estimation rule (for details see [1,3,6]):

$$\hat{F}(k,l) = Y(k,l)\left[1 + \lambda(k,l)\exp\left\{-\frac{r^2(k,l)}{\alpha}\right\}\right]^{(-1)} = Y(k,l)f[r(k,l)], \quad (14)$$

where $f[r(k,l)]$ is denoted the attenuation function. The attenuation function depends on the so-called instantaneous signal-to-noise ration (SNR) $r^2(k,l)$, which is defined by

$$r^2(k,l) = \frac{|Y(k,l)|^2}{N(k,l)}, \quad (15)$$

$N(k,l)$ is the noise variance of the (k,l)th spectral coefficient. The parameter α is a weighting factor similar to the one used in a generalized Wiener filter (cf. [23]). The parameter $\lambda(k,l)$ is obtained to

$$\lambda(k,l) = \frac{R(k,l)}{N(k,l)} \cdot \frac{P_0(k,l)}{1 - P_0(k,l)}, \quad (16)$$

where $R(k,l)$ is the variance of coefficients containing both signal and noise. Finally, $P_0(k,l)$ is the unconditional a priori probability for $H_0(k,l)$. Verbalizing $\lambda(k,l)$, it is the signal-plus-noise-to-noise ratio weighted by the ratio $P_0(k,l)/[1 - P_0(k,l)]$.

In practice, the precise values of $R(k,l)$ and $P_0(k,l)$ are not known. A common simplification is to regard $\lambda(k,l)$ as a free parameter balancing noise reduction and signal preservation, which does not vary over (k,l). As derived in the [1], an alternative is to determine λ as a function of the significance of a hard decision between the hypotheses H_0 and H_1. The shape of $f(r)$ for different values of λ is shown in Fig. 10. As $f(r)$ increases monotonically over r, an observed coefficient is attenuated the more, the smaller its instantaneous SNR r^2.

Fig. 11 illustrates this algorithm. After decomposing the input image into blocks and transforming these, the absolute value of each coefficient is taken in the lower branch. For each coefficient, the instantaneous SNR is then calculated, and the attenuation determined. After multiplication according to Eq. (14), the restored image is reconstructed.

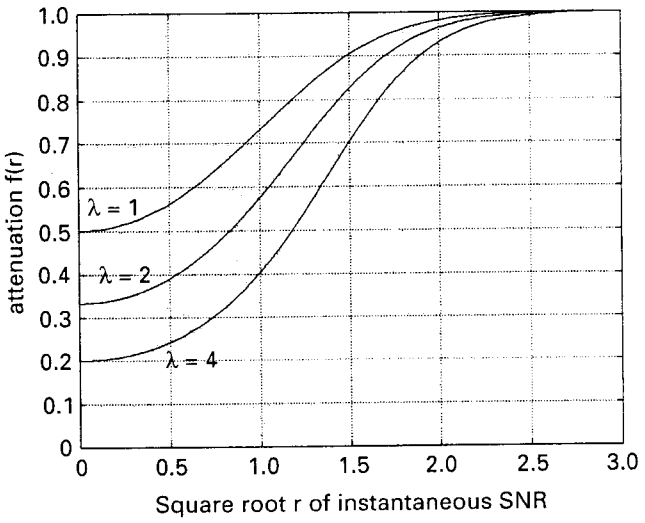

Fig. 10. Illustration of the noise attenuation function $f(r)$ for different values of λ, and $\alpha = 1$.

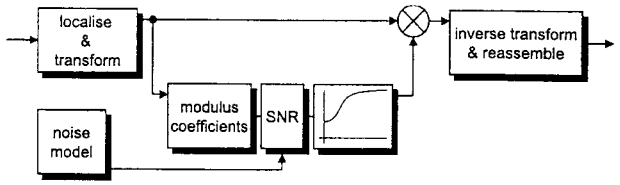

Fig. 11. Block diagram of spectral amplitude estimation.

The box labelled 'noise model' stores the spectral noise variance $N(k,l)$, which is assumed to be known for each coefficient.

4.2. Anisotropic spectral amplitude estimation

So far, the restoration algorithm is isotropic in the sense that the same attenuation function is applied to all coefficients, regardless of their position (k,l). This is a consequence of assuming the parameter λ to be independent of (k,l). With respect to oriented patterns like lines and edges, however, the performance of spectral amplitude estimation can be enhanced if orientation is detected, and the attenuation function adapted for coefficients which are likely to contribute to orientation. More specifically, we extend our algorithm to take into account the prior knowledge that a block can exhibit an arbitrary orientation. When using the LDT or the DFT, spectral coefficients contributing to the local orientation concentrate along a line perpendicular to the spatial orientation which passes through the origin. The central idea is to apply less attenuation to coefficients along these coefficients, with this behaviour the more pronounced, the more prominent the local orientation is [2].

4.2.1. Orientation detection

An approach to detect local orientation in the Fourier domain was described in [11]. The local Fourier energy spectrum is analysed by means of a 2×2 inertia matrix, the eigenvectors of which determine the two axes along which energy concentration is strongest (local orientation) and least in a least-squares sense, respectively. The corresponding eigenvalues measure how well the concentration is pronounced. By means of Parseval's theorem, this approach was transformed into the spatial domain. This resulted in the same algorithm for orientation detection as in [18], where orientation detection was formulated as the detection of the direction in which spatial intensity variance is least.

For our purposes, the spectral-domain formulation in [11] is suited best, where we have carried out the following modifications with respect to the calculation of the inertia matrix:

- Calculation of the inertia matrix is based on the instantaneous SNR $r^2(i,j)$ rather than on the coefficients $Y(i,j)$ alone. This increases robustness against noise. Additionally, potential noise anisotropies, which may occur, e.g. in images recorded in interlaced mode, are normalized out.
- The inertia matrix is rotated and normalized such that its eigenvalues range between a maximum of $+1$ for the local orientation axis, and a minimum of -1 for the axis perpendicular to local orientation. These values are reached for optimal concentration of energy along the local orientation axis.
- To weight all spectral coefficients equally independent of their distance from the origin, all coefficients are 'projected' onto the unit circle. In other words, spectral energy is thought to be concentrated on the unit circle.

In a first step, we hence calculate the following matrix A:

$$A=\begin{pmatrix}\sum_{i,j}\frac{i^2}{i^2+j^2}r^2(i,j) & \sum_{i,j}\frac{ij}{i^2+j^2}r^2(i,j)\\ \sum_{i,j}\frac{ij}{i^2+j^2}r^2(i,j) & \sum_{i,j}\frac{j^2}{i^2+j^2}r^2(i,j)\end{pmatrix}, \tag{17}$$

where the sums extend over all sites in a block spectrum. The division of each matrix entry by i^2+j^2 is responsible for the projection of $r^2(i,j)$ onto the unit circle. With respect to a 'genuine' inertia matrix, A is rotated by 90°. To obtain eigenvalues between -1 and $+1$, another matrix B is calculated from A according to

$$B=\frac{2}{\mathrm{Trace}(A)}A-I=\begin{pmatrix}b_1 & b_2\\ b_2 & -b_1\end{pmatrix}, \tag{18}$$

where b_1 and b_2 are easily identified. I is the 2×2 identity matrix. The eigenvalues of B are

$$d_1=\sqrt{b_1^2+b_2^2}=d,\quad d_2=-d. \tag{19}$$

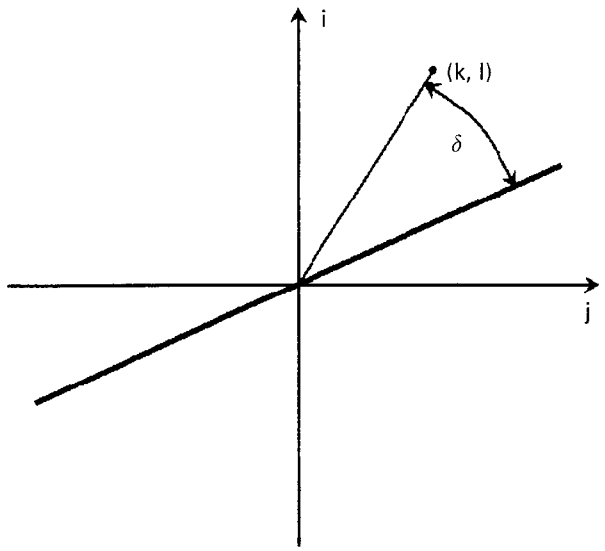

Fig. 12. 2D spatial frequency domain for each block, where (k,l) are the coordinates of an observed coefficient $Y(k,l)$. The bold line is the orientation axis, while δ is the angle between the orientation axis and the vector pointing to (k,l).

The eigenvectors of B point parallel to and perpendicular to the local orientation axis, respectively. In the ideal case that all non-zero coefficients are restricted to the orientation axis, we have $d = 1$, whereas an ideally isotropic distribution of the $r^2(i,j)$ results in $d = 0$.

It is not necessary to determine the eigenvectors explicitly. We rather evaluate for each site (k,l) in the spectral domain the expression

$$g(k,l) = \frac{(k,l)\cdot B\cdot(k,l)^{\mathrm{T}}}{k^2+l^2} = d\cdot\cos(2\delta), \tag{20}$$

where δ is the angle between (k,l) and the local orientation, which needs not to be known explicitly (see Fig. 12, and Appendix A). Clearly, the maximum of $g(k,l)$ is equal to the eigenvalue $d_1 = d$, which is reached if (k,l) is situated on the orientation axis. As mentioned above, the maximum d is the larger, the more prominent the local orientation is.

4.2.2. Anisotropic attenuation

The orientation information is used to modulate the attenuation depending on the processed coefficient's position relative to the orientation axis. The statistical vehicle to do so is to decrease the noise-only probability $P_0(k,l)$ for locations k,l close to the detected orientation, thus implying that spectral coefficients close to or along the orientation axis contribute to signal even when their instantaneous SNR is low. To this end, we define a 'selectivity' function

$$M(k,l) = \frac{\max[g(k,l),0]^8}{d^7} \tag{21}$$

which ranges between zero and one, and increases sharply when k,l is close to the orientation axis. The exponentiation by eight serves to make $M(k,l)$ sufficiently selective. The maximum operation is necessary since $g(k,l)$ also takes negative values (see (20)), which would become positive after exponentiation, hence introducing unwanted mirror orientations. For coefficients lying on the orientation axis, the normalization by d^7 ensures $M(k,l) = d$. Assuming $P_0(k,l)$ to decrease sharply when approaching the orientation axis, we model

$$\frac{P_0(k,l)}{1-P_0(k,l)} = 1 - M(k,l). \tag{22}$$

This model is inserted into expression (16) for $\lambda(k,l)$, which now varies over (k,l) if orientation is detected. From (14), (16) and (22), we obtain for the attenuation function $f(r)$

$$f[r(k,l)] = \left[1 + \lambda_0(1 - M(k,l))\exp\left\{-\frac{r^2(k,l)}{\alpha}\right\}\right]^{(-1)}, \tag{23}$$

where $\lambda_0 = R(k,l)/N(k,l)$ is assumed to be constant. The attenuation function in Fig. 10 then fans out to an angle-dependent family of attenuation functions, which is illustrated in Fig. 13. Alternatively or additionally, one could assume the parameter α to be variable according to

$$\alpha = \alpha(k,l) = \frac{\alpha_0}{1 - M(k,l)} \tag{24}$$

which can be motivated by the 'signal-equivalent' approach in [7,19].

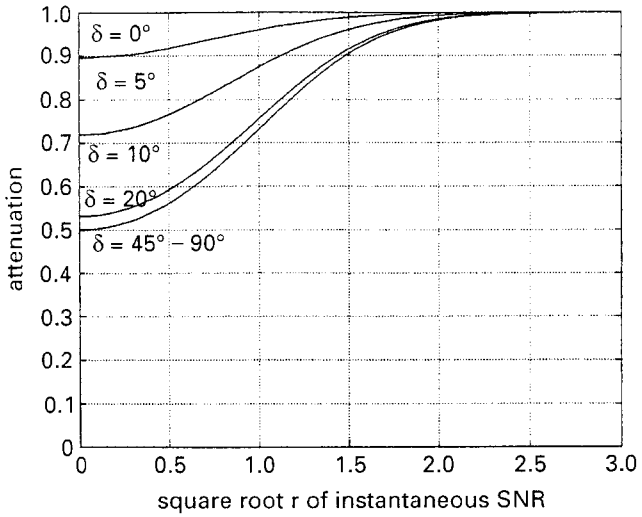

Fig. 13. Orientation-dependent family of attenuation functions as given by Eq. (23), with $\lambda_0 = 1.0$ and $\alpha = 1$. Depending on the angle of a coefficient's position relative to the dominant orientation the attenuation is reduced. For this plot, ideal orientation, i.e. $d = 1$ was assumed.

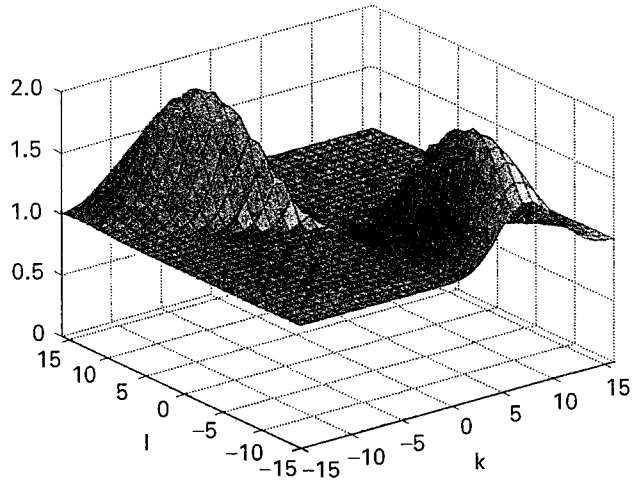

Fig. 14. Anisotropic transfer function for selective enhancement as a function of (k, l), plotted for a DFT block spectrum with $K = 32$, and $-K/2 < (k, l) \leqslant K/2$. For each block spectrum, the passband is tuned to the detected local orientation. For this plot, an ideal orientation of $-60°$ and $\beta = 1$ were assumed.

4.2.3. Enhancing orientation

In addition to adaptive attenuation, spectral orientation information can also be used to selectively enhance the detected oriented patterns. Generally, sharpness can be improved by bandpass or highpass filtering. Isotropic filtering, however, increases noise considerably. Confining the enhancement to coefficients contributing to local orientation allows to keep noise amplification moderate due to the limited number of coefficients involved and their relatively good SNR. A radial selection of spatial frequencies that are relevant for appreciation of image sharpness can be achieved by the bandpass filter

$$H(k,l) = \beta \sin^2\left(\sqrt{2}\pi \frac{\sqrt{k^2+l^2}}{K}\right), \tag{25}$$

where the block spectrum is of size $K \times K$. The parameter β allows to control the amount of amplification. Selectivity to orientation is obtained by combining this isotropic bandpass filter with the selectivity function $M(k,l)$ from (21) by

$$H_a(k,l) = 1 + M(k,l)H(k,l). \tag{26}$$

The adaptivity of $M(k,l)$ to d ensures that the enhancement applied depends on the prominence of the detected orientation. An instance of this filter is illustrated in Fig. 14.

As the LDT behaves like the DFT with respect to local orientation, the LDT can be used within this framework in the same manner as the DFT. The algorithm's ability to preserve or even emphasize visually important structures remains untouched.

5. Results

Fig. 15 shows the original 'boats' image and a noisy version, where zero mean Gaussian noise with a variance of $\sigma^2 = 35^2 = 1225$ was added, yielding a PSNR of 17.25 dB. The size of all images used in this section is 512×512 pixels. From the processed images, a 32-pixel wide stripe along each image edge was removed before reproduction here, since these contain edge artifacts due to the finite block size. In LDT-based processing, the size of the basis functions was chosen to 32×32 pixels. Correspondingly, the blocksize in DFT-based processing was 32×32 pixels, where a Hanning window was applied to each block, requiring an overlap of 16 pixels in each dimension.

The spectral noise power $N(k,l)$ needed to determine the instantaneous SNR $r^2(k,l)$ in (15) for each

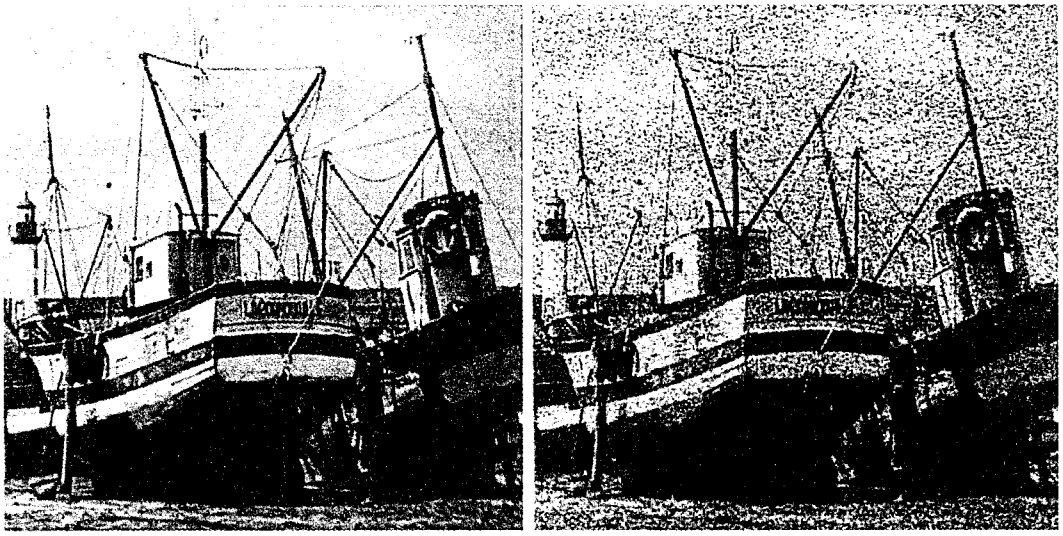

Fig. 15. Original 'boats' image (left) and its noisy version with a PSNR of 17.25 dB.

Fig. 16. Restoration of the noisy 'boats' image using DFT-based (left, PSNR = 21.5 dB) and LDT-based (right, PSNR = 21dB) anisotropic amplitude estimation. Parameters: $\lambda_0 = 1$, $\alpha = 4$, $\beta = 0$.

coefficient was measured from 40 images containing only realizations of the unprocessed additive noise by estimating the expected value $E[|Y(k,l)|^2|H_0]$ given that only noise is observed. This was carried out individually for each transform. For the DFT, the blocksize and the window function were the same as used for processing the images. Similarly, a basis function size of 32×32 pixels was used for the LDT. Since the added noise was nearly white, the estimated spectral noise power $N(k,l)$ did not vary strongly over (k,l), yielding about $N(k,l) = \sigma^2 \cdot 128$ for the windowed DFT, and about $N(k,l) = \sigma^2$ for the LDT.

Fig. 16 shows the restoration results obtained by anisotropic spectral amplitude estimation using the DFT (left) and the LDT (right). The parameters of the attenuation function were experimentally determined to $\lambda_0 = 1$ and $\alpha = 4$. Here, no orientation enhancement was applied, i.e. $\beta = 0$ in (25) and (26). The results of simultaneous noise reduction and orientation enhancement using $\beta = 1$ are given in Fig. 17.

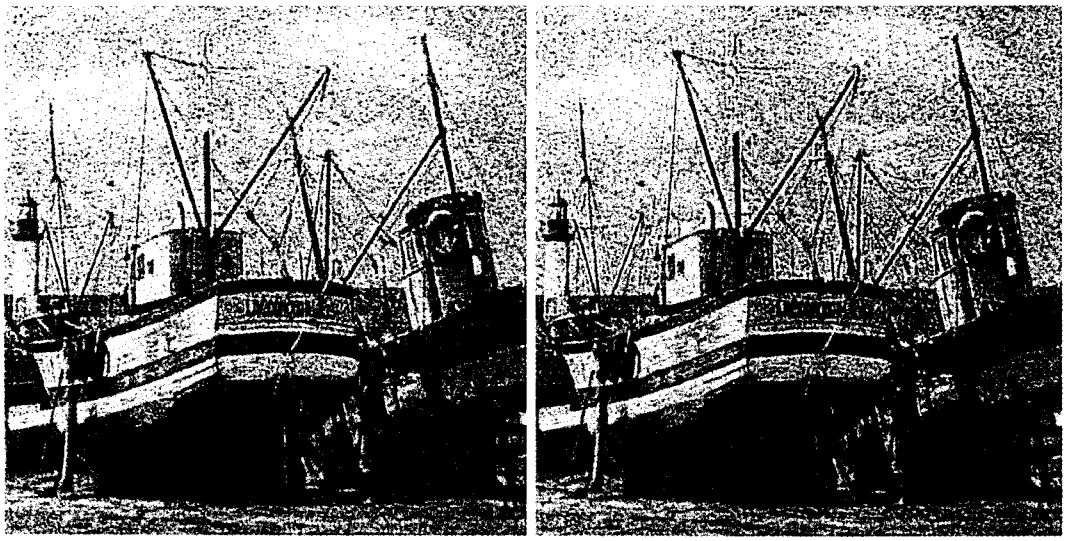

Fig. 17. Restoration and enhancement of the noisy 'boats' image using DFT-based (left, PSNR = 21.2 dB) and LDT-based (right, PSNR = 20.5 dB) anisotropic amplitude estimation and orientation enhancement ($\beta = 1$).

Qualitatively, LDT-based and DFT-based versions of the algorithm perform similarly. For a quantitative comparison, we applied the above processing to 40 different realizations of the noisy image. After processing, the remaining noise was then estimated by averaging these into one frame, and subtracting it from each processing result. To assess the loss of detail information relative to the original image caused by noise reduction, we have subtracted the averaged frame from the original image (which was given on the left-hand side of Fig. 15), and calculated the mean-square error.

Table 1 summarizes the measured noise reduction performances of both DFT- and LDT-based versions of our algorithm for the cases of applying anisotropic noise reduction only as shown in Fig. 16, and applying anisotropic noise reduction and orientation enhancement as shown in Fig. 17. In addition, quantitative results for isotropic noise reduction as described in Section 4.1 are also included. To obtain the latter results, a constant λ of $\lambda = 1$ was used.

In Table 2, the measured error energies between original and averaged processing results can be found. They prove a slightly better ability of the anisotropic algorithm to preserve details. Visually, the effects of anisotropic processing are stronger than implied by these measurements, since it affects predominantly the important edge and line information. Of course, the impact of anisotropic processing is strongest when combined with orientation enhancement.

Table 1
PSNR in dB which remains after processing the image in Fig. 15. Method 1: isotropic noise reduction with $\lambda(k,l) = 1$, $\alpha = 4$, $\beta = 0$. Method 2: anisotropic noise reduction with $\lambda_0(k,l) = 1$, $\alpha = 4$, $\beta = 0$. Method 3: anisotropic noise reduction and orientation enhancement with $\lambda_0(k,l) = 1$, $\alpha = 4$, $\beta = 1$

	Method 1	Method 2	Method 3
LDT	21	21	20.5
FFT	21.5	21.5	21.2

Table 2
Estimated detail loss of noise reduction captured by the mean square error MSE(Δ) between original image and the average of 40 processed realizations of the noisy image. The table entries are calculated as $10\log[255^2/\mathrm{MSE}(\Delta)]$ dB

	Method 1	Method 2
LDT	33.2	33.4
FFT	33	33.2

We compared our approach to the FIR/median-hybrid (FMH) filter which was introduced by Nieminen et al. [29], and which we found earlier to possess outstanding detail preservation properties

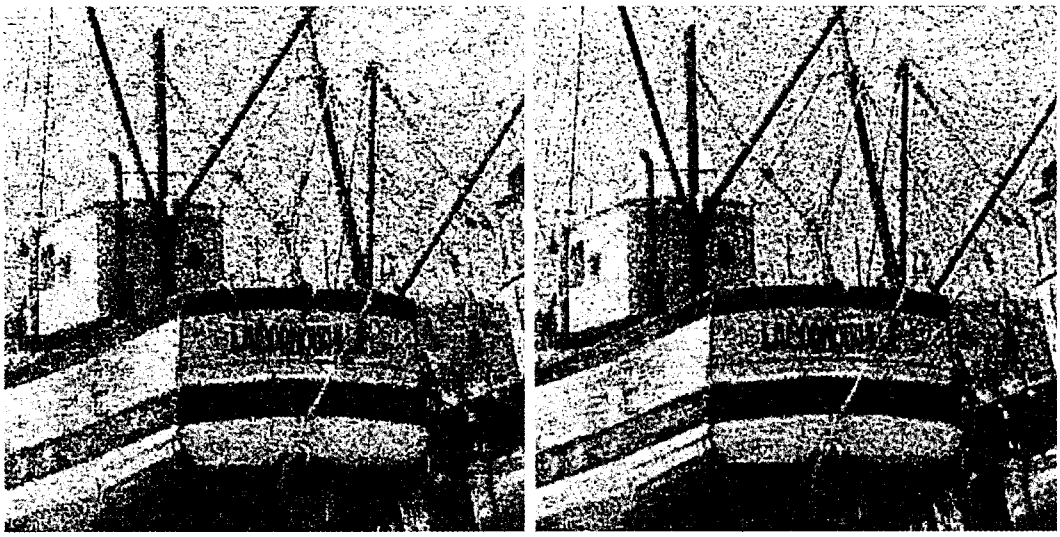

Fig. 18. Left: Central portion of the noisy 'boats' image processed by a 7×7 FMH filter. Right: Central portion of the LDT-processing result from Fig. 16.

Fig. 19. Original 'lena' image (left) and its noisy version with a PSNR of 20.2 dB.

(cf. [4]). The processing result is given in Fig. 18. Quantitatively, we measured a PSNR of 21 dB for the FMH filter, while its detail loss as measured in Table 2 was 32 dB, indicating a somewhat higher loss of detail than for our algorithm. Qualitatively, the different detail preservation properties of the FMH-filter and the anisotropic spectral amplitude estimator are best noticeable when comparing oriented structures like the rigging in Fig. 18. Also, the FMH-filtered result shows a needle-like, more grainy noise texture.

The degree of noise reduction achieved depends on the parameter settings. While the mentioned settings reduce noise by 3–4 dB, more reduction is possible by increasing the values for λ and α, albeit at the cost of more loss of image detail. As an example, another test image is shown in Fig. 19. Here, zero mean Gaussian noise with a variance of $\sigma^2 = 625$ was added (PSNR = 20.2 dB). Fig. 20 shows LDT-based processing results for anisotropic noise reduction and anisotropic noise reduction combined with orientation enhancement. The

Fig. 20. Processing results for the 'lena' image using the LDT. Left: Anisotropic noise reduction ($\lambda_0 = 1.5, \alpha = 5, \beta = 0$).PSNR: 24.8 dB. Right: Anisotropic noise reduction and orientation enhancement ($\lambda = 1.5, \alpha = 5, \beta = 1$). PSNR: 24.3 dB.

Table 3
PSNR in dB which remains after processing the image in Fig. 19. Method 1: isotropic noise reduction with $\lambda(k,l) = 1.5$, $\alpha = 5$, $\beta = 0$. Method 2: anisotropic noise reduction with $\lambda_0(k,l) = 1.5$, $\alpha = 5, \beta = 0$. Method 3: anisotropic noise reduction and orientation enhancement with $\lambda_0(k,l) = 1.5$, $\alpha = 5$, $\beta = 1$

	Method 1	Method 2	Method 3
LDT	25	24.8	24.3
FFT	25.6	25.5	25.2

Table 4
Estimated detail loss of noise reduction for processing of the noisy image in Fig. 19. The measurements were obtained in the same manner as in Table 2

	Method 1	Method 2
LDT	34.5	34.9
FFT	33.9	34.3

parameters were set to $\lambda_0 = 1.5$ and $\alpha = 5$ for anisotropic noise reduction. The orientation enhancement parameter was set to $\beta = 1$, as before. Quantitative measurements can be found in Tables 3 and 4, indicating a noise reduction by 4–5 dB.

6. Conclusions

We have derived a new real-valued lapped transform termed the lapped directional transform. Unlike other real-valued standard transforms as DCT or MLT, the new transform uses directional basis functions like the complex-valued DFT, hence providing for an unambiguous relationship between orientation in the spatial domain and spectral energy distribution. The LDT basis functions are real valued, and not separable. They are also overlapping, and decay smoothly to zero toward the block boundaries. The LDT is therefore unambiguously orientation selective, and free of block and leakage artifacts. As it doubles the data volume rather than quadrupling it, it is also less redundant than the DFT in connection with standard data windows. Although the LDT's basis functions are not separable, it is calculated by invoking a separable transform, viz., the MLT. Since fast algorithms exist for the MLT, so do fast algorithms for the LDT.

These properties make the LDT especially suited for image processing and reconstruction by local spectral analysis. To illustrate this point, we have applied the LDT in an anisotropic spectral magnitude estimation algorithm for image restoration, where this orientation property was exploited to improve performance with respect to perceptually

relevant detail. Both LDT- and DFT-based versions of this algorithm performed comparably. If the processed coefficients are to be used in a subsequent encoding stage, the redundancy of the LDT can be circumvented by reconstructing either MLT-coefficients or MLT′-coefficients from the processed LDT-coefficients using (11) or (12). This is an advantage of the LDT over the DFT.

For future work, we intend to measure quantitatively the energy compaction of the LDT in a similar manner as done for other lapped transforms in [10]. With respect to orientation, a more basic question is how well transforms which are based on separable 1D-covariance models perform with respect to oriented structures which may not be characterized well by these models. We expect that such measurements will show that the LDT's orientation energy compaction allows better separation of lines or edges from noise than real-valued separable transforms like the DCT and MLT.

Appendix A

To show that $g(k,l)$ in (20) is equal to $d\cdot\cos(2\delta)$ in Fig. 12, let us define the unit vector $\boldsymbol{v}(k,l) = (k,l)/\sqrt{k^2+l^2}$. Then,

$$g(k,l) = \boldsymbol{v}(k,l)\cdot B\cdot \boldsymbol{v}(k,l)^{\mathrm{T}}. \tag{A.1}$$

Let us denote the normalized eigenvectors of B by $\boldsymbol{e}_1$ and $\boldsymbol{e}_2$, where $\boldsymbol{e}_1$ corresponds to the eigenvalue $d_1 = d$, and $\boldsymbol{e}_2$ to $d_2 = -d$ as given in (19). The eigenvector $\boldsymbol{e}_1$ defines the orientation axis. Since B is symmetric, the eigenvectors are orthonormal, and we can develop $\boldsymbol{v}(k,l)$ into

$$\boldsymbol{v}(k,l) = a\cdot\boldsymbol{e}_1 + b\cdot\boldsymbol{e}_2 \tag{A.2}$$

with $a = \cos(\delta)$ and $b = \sin(\delta)$. Inserting (A.2) in (A.1) yields

$$g(k,l) = d[\cos^2(\delta) - \sin^2(\delta)] = d\cdot\cos(2\delta) \tag{A.3}$$

References

[1] T. Aach, Spectral transform-based nonlinear restoration of medical images: algorithms and a comparative evaluation, in E.R. Dougherty, J.T. Astola (Eds.), Electronic Imaging: Nonlinear Image Processing X, Vol. 3646, SPIE, San Jose, USA, 23–29 January 1999, pp. 270–280.

[2] T. Aach, D. Kunz, Anisotropic spectral magnitude estimation filters for noise reduction and image enhancement, Proceedings IEEE International Conference on Image Processing (ICIP), Lausanne, Switzerland, 16–19 November, 1996, pp. 335–338.

[3] T. Aach, D. Kunz, Spectral estimation filters for noise reduction in x-ray fluoroscopy imaging, in: G. Ramponi, G.L. Sicuranza, S. Carrato, S. Marsi (Eds.), Proceedings European Signal Processing Conference (EUSIPCO), Trieste, Italy, 10–13 September 1996, pp. 571–574.

[4] T. Aach, D. Kunz, Multiscale linear/median hybrid filters for noise reduction in low-dose x-ray images, Proceedings IEEE International Conference on Image Processing (ICIP), Santa Barbara, CA, October 26–29, 1997, pp. 358–361.

[5] T. Aach, D. Kunz, Spectral amplitude estimation-based x-ray image restoration: an extension of a speech enhancement approach, in: A. Stouraitis, N. Kalouptsidis, S. Theodoridis, I. Pitas (Eds.), Proceedings European Signal Processing Conference (EUSIPCO), Rodos, Greece, 5–9 September 1998, pp. 5–9.

[6] T. Aach, U. Schiebel, G. Spekowius, Digital image acquisition and processing in medical x-ray imaging, J. Electron. Imaging (Special Section on Biomedical Image Representation) 8 (1) (1999) 7–22.

[7] J.F. Abramatic, L.M. Silverman, Nonlinear restoration of noisy images, IEEE Trans. Pattern Anal. Mach. Intell. 4 (2) (1982) 141–149.

[8] N. Ahmed, T. Natarajan, K.R. Rao, Discrete cosine transform, IEEE Trans. Comput. 23 (1974) 90–93.

[9] A.N. Akansu, R.A. Haddad, Multiresolution Signal Decomposition, Academic Press, Boston, 1992.

[10] A.N. Akansu, F.E. Wadas, On lapped orthogonal transforms, IEEE Trans. Signal Process. 40 (2) (1992) 439–443.

[11] J. Bigün, G.H. Granlund, Optimal orientation detection of linear symmetry, Proceedings IEEE First International Conference on Computer Vision, London, UK, June 1987, pp. 433–438.

[12] M. Breeuwer, R. Heusdens, P. Zwart, Overlapped transform coding of medical x-ray images, Proceedings Medical Imaging 1994: Image Capture, Formatting, and Display, Vol. 2164, SPIE, 1994.

[13] O. Cappé, Elimination of the musical noise phenomenon with the Ephraim and Malah noise suppressor, IEEE Trans. Speech Audio Process. 2 (2) (1994) 345–349.

[14] R.J. Clarke, Transform Coding of Images, Academic Press, London, 1985.

[15] R.J. Clarke, B. Tech, Relation between the carhunen loève and cosine transform, IEEE Proc. 128 (6) (1981) 359–360.

[16] R.E. Crochiere, A weighted overlap-add method of short-time Fourier analysis/synthesis, IEEE Trans. Acoust. Speech Signal Process. 28 (1) (1980) 99–102.

[17] Y. Ephraim, D. Malah, Speech enhancement using a minimum mean-square error short-time spectral amplitude estimator, IEEE Trans. Acoust. Speech Signal Process. 32 (6) (1984) 1109–1121.

[18] M. Kass, A. Witkin, Analyzing oriented patterns, Comput. Vision Graphics Image Process. 37 (1987) 362–385.

[19] H.E. Knutsson, R. Wilson, G.H. Granlund, Anisotropic nonstationary image estimation and its applications: Part I – restoration of noisy images, IEEE Trans. Commun. 31 (3) (1983) 388–397.

[20] D. Kunz, T. Aach, Lapped directional transform: a new transform for spectral image analysis, Proceedings IEEE International Conference on Acoustics Speech Signal Processing (ICASSP), Phoenix, AZ, 15–19 March 1999, pp. 3433–3436.

[21] J.S. Lim, Image restoration by short space spectral subtraction, IEEE Trans. Acoust. Speech Signal Process. 28 (2) (1980) 191–197.

[22] J.S. Lim, Two-Dimensional Signal and Image Processing, Prentice-Hall, Englewood Cliffs, NJ, 1990.

[23] J.S. Lim, A.V. Oppenheim, Enhancement and bandwidth compression of noisy speech, Proc. IEEE 67 (12) (1979) 1586–1604.

[24] H.S. Malvar, Lapped transforms for efficient transform/subband coding, IEEE Trans. Acoust. Speech Signal Process. 38 (6) (1990) 969–978.

[25] H.S. Malvar, Signal Processing with Lapped Transforms, Artech House, Norwood, MA, 1992.

[26] H. Malvar, A modulated complex lapped transform and its application to audio processing, Proceedings IEEE International Conference on Acoustics Speech and Signal Processing (ICASSP), Phoenix, AZ, 15–19 March 1999, pp. 1421–1424.

[27] H.S. Malvar, D.H. Staelin, The LOT: transform coding without blocking effects, IEEE Trans. Acoust. Speech Signal Process. 37 (4) (1989) 553–559.

[28] R.J. McAulay, M.L. Malpass, Speech enhancement using a soft-decision noise suppression filter, IEEE Trans. Acoust. Speech Signal Process. 28 (2) (1980) 137–145.

[29] A. Nieminen, P. Heinonen, Y. Neuvo, A new class of detail-preserving filters for image processing, IEEE Trans. Pattern Anal. Mach. Intell. 9 (1) (1987) 74–90.

[30] M. Unser, On the approximation of the discrete Karhunen–Loeve transform for stationary processes, Signal Processing 7 (1984) 231–249.

[31] H.L. van Trees, Decision, Estimation, and Modulation Theory, Part I, Wiley, New York, 1968.

[32] P. Vary, Noise suppression by spectral magnitude estimation – mechanism and limits, Signal Processing 8 (4) (1985) 387–400.

[33] L.P. Yaroslavsky, Local adaptive image restoration and enhancement with the use of DFT and DCT in a running window, in: M. Unser, A. Aldroubi, A.F. Laine (Eds.), Wavelet Applications in Signal and Image Processing IV, Vol. 2825, SPIE, 1997, pp. 1–13.

[34] L.P. Yaroslavsky, Local adaptive filtering in transform domain for image restoration, enhancement and target location, in: E. Wegner, L. Dimitrov (Eds.), Sixth International Workshop on Digital Image Processing and Computer Graphics, Vol. 3346, SPIE, 1997, pp. 2–17.

[35] R.W. Young, N.G. Kingsbury, Frequency-domain motion estimation using a complex lapped transform, IEEE Trans. Image Process. 2 (1) (1993) 2–17.