

Signal-Adapted Analytic Wavelet Packets in Arbitrary Dimensions

Bächle, Matthias

Institute of Industrial Inform. Techn.

Karlsruhe Institute of Technology

Karlsruhe, Germany

<https://orcid.org/0000-0001-7228-299X>

Schambach, Maximilian

Institute of Industrial Inform. Techn.

Karlsruhe Institute of Technology

Karlsruhe, Germany

<https://orcid.org/0000-0002-4927-266X>

Puente León, Fernando

Institute of Industrial Inform. Techn.

Karlsruhe Institute of Technology

Karlsruhe, Germany

fernando.puente-leon@kit.edu

Abstract—The analytic wavelet packet transform, based on the dual-tree approach, represents a complex-valued extension of the wavelet packet transform. A generalization to multiple dimensions can be realized using fully separable wavelet trees, but this restricts the possible subband combinations. To overcome these limitations, we present a flexible framework to calculate N-D analytic wavelet packets with configurable decomposition structures and filter types. By introducing a new subband notation for the nodes of the N-D wavelet binary tree, both anisotropic and isotropic decomposition structures can be realized. Based on this subband notation, a full frame in N dimensions is defined and combined with an optimal basis selection, which we generalized to arbitrary dimensions, to find signal-adapted decomposition structures. As a multi-dimensional example, the framework is applied to the compression and denoising of a 4D light field. The results are evaluated in terms of the PSNR and SSIM and compared with the discrete cosine transform.

Index Terms—analytic wavelet packets, signal-adapted filtering, signal compression, signal denoising

I. INTRODUCTION

The Discrete Wavelet Transform (DWT) is well known to sparsely represent structured, piecewise continuous signals and has been a powerful tool in signal processing for decades. Extensions of the transform to two-dimensional signals (most importantly, images) have been studied since the late 80's [1] and successfully applied to various tasks, ranging from signal analysis [2] to compression [3] and denoising [4], to name a few. However, with recent advances in sensor technology, computational power and data storage capacities, high-dimensional data has become more available and feasible to process and analyse. For example, 3D tomography data (computer tomography, magnetic resonance tomography, and others) are widely used in medical imaging. In the computer vision community, high-dimensional data such as 3D hyperspectral images [5], 4D light fields [6], and 5D hyperspectral light fields [7] have been attracting attention. For these high-dimensional signals, sparse or signal-adapted bases become increasingly important for data analysis, processing and storage.

While the Wavelet Packet Transform (WPT) [13] allows a more flexible signal representation than the DWT, the transform is also not shift-invariant, due to the downsampling in the filter bank. Therefore, the Analytic Wavelet Packet Transform (AWPT) was introduced simultaneously by Selesnick [8] and

Weickert et al. [9], combining the dual-tree approach, presented by Kingsbury [10] and Selesnick et al. [11], with the real-valued WPT. Using the dual-tree extension, the resulting basis functions are designed to be nearly analytic, resulting in an improved shift-invariance. However, a detailed formulation of an N-D extension of the AWPT has not yet been discussed.

To this end, we present a framework and reference implementation for the N-D AWPT with fully configurable decomposition structures and dimension-specific filter types. The source code of our implementation is made publicly available [12]. Because different kind of data has its own inherent characteristics, we generalize the decomposition structure and the dual-tree approach such that the filter types and frequency subbands can be chosen independently for each dimension. Furthermore, a frame is calculated by the new decomposition method, which is used to determine a signal-adapted basis employing the search principle introduced by Coifman and Wickerhausen [13]. Finally, the framework is applied to the compression and denoising of 4D light field data.

II. N-D ANALYTIC WAVELET PACKETS

One-dimensional wavelet packets are based on a cascaded filter bank and can be represented by a binary tree, where every node stands for a different frequency subband. Every intermediate node is split by a lowpass and a bandpass filter, with subsequent downsampling by a factor 2, into its two children. If only the lowpass signal is further split, the resulting structure is called the DWT. When the bandpass signal is also split, the DWT is extended to the more general WPT. The complex extensions of the DWT and the WPT are called the Dual-Tree Complex Wavelet Transform (DT-CWT) and the AWPT, respectively. The presented framework allows arbitrary decomposition structures based on an N-D binary tree representation of the wavelet packets. In the following section, the necessary definitions of the N-D binary tree, the nodes and their relationships are introduced. Furthermore, a detailed description of the algorithm that performs the decomposition based on an N-D recursive depth-first search is given.

A. Decomposition structure notation

To uniquely define a binary tree, a set of leaves and a parent-child relationship is needed. Therefore, a unique identifier for




Figure 1. Distribution of the frequency domain of a 2D signal with its associated node identifiers.

the leaves has to be introduced. In the one-dimensional case, a node is uniquely identified by the tuple

$$\mathbf{b} = (s, k) \in \mathbb{N}^2, \quad \text{with } s \geq 0, k \in [0, 2^s - 1], \quad (1)$$

where s denotes the level of the node and k denotes the subband within this level. Throughout this paper, $[\cdot, \cdot]$ denotes a subset of \mathbb{N} . For a signal to be reconstructible, only full binary trees are admissible, i.e. binary trees where every node has exactly zero or two children. Given a parent node (s, k) , its two children are naturally identified as

$$(s + 1, 2k), \quad \text{lowpass child}, \quad (2a)$$

$$(s + 1, 2k + 1), \quad \text{bandpass child}. \quad (2b)$$

This describes that the children are one level deeper than their parent and that a higher subband number represents a higher center frequency in the associated level. This naturally leads to the subband center frequency $f_c(s, k) = kf_N/2^s$, where f_N denotes the Nyquist frequency. An extension of the node identifier to N dimensions is now realized using the set

$$\mathcal{B} = \{\mathbf{b}_i, i \in [1, N]\}, \quad (3)$$

consisting of N tuples \mathbf{b}_i , which each describe a level and subband per dimension, as defined in (1). In this N-D extension, a node $\tilde{\mathcal{B}}$ is called a child of \mathcal{B} , if for at least one dimension i , $\tilde{\mathbf{b}}_i$ is a child of \mathbf{b}_i and for all remaining $j \neq i$, $\mathbf{b}_j = \tilde{\mathbf{b}}_j$. An example of this notation, using an arbitrary frequency resolution of a 2D signal, is shown in Fig. 1.

Using the node identifier and the parent-child relationship, the decomposition structure of the N-D WPT can be described by a binary tree, which is uniquely identified by the set

$$\mathcal{T} = \{\mathcal{B}_m, m \in [1, M]\}, \quad (4)$$

where the nodes \mathcal{B}_m identify the M leaves of the tree. A decomposition containing redundant nodes can be associated with an extension of \mathcal{T} by additional intermediate nodes, which will be discussed in detail in Section III.

B. Decomposition algorithm

The decomposition algorithm is recursively designed, starting its execution with the root of the binary tree. Each recursive call consists of mainly four steps:

- 1) Check if the current node $\tilde{\mathcal{B}}$ is a leaf of the binary tree. If so, stop the recursion.

- 2) Determine in which dimension the data has to be filtered in the current recursive call.
- 3) Filter the input data of the current recursive call along the determined dimension.
- 4) Start two new recursive calls for the lowpass and bandpass child which resulted from the filtering step.

By testing every element of \mathcal{T} for its accessibility from the current node $\tilde{\mathcal{B}}$, it can be decided whether $\tilde{\mathcal{B}}$ is a leaf or not. For this purpose, an ancestor-descendant relationship between two nodes has to be defined. The node $\tilde{\mathcal{B}}$ is a leaf if \mathcal{T} does not contain any descendants of $\tilde{\mathcal{B}}$. Two nodes \mathcal{B} and $\tilde{\mathcal{B}}$, containing the tuples $\mathbf{b}_i = (s_i, k_i)$ and $\tilde{\mathbf{b}}_i = (\tilde{s}_i, \tilde{k}_i)$, respectively, are defined to have an ancestor-descendant relationship

$$\mathcal{B} \succ \tilde{\mathcal{B}}, \quad \mathcal{B} \text{ is ancestor of } \tilde{\mathcal{B}}, \quad (5)$$

if all of the relations

$$\forall_{i \in [1, N]} : \tilde{s}_i \geq s_i, \quad (6a)$$

$$\exists_{i \in [1, N]} : \tilde{s}_i > s_i, \quad (6b)$$

$$\forall_{i \in [1, N]} : k_i = \lfloor \tilde{k}_i \cdot 2^{s_i - \tilde{s}_i} \rfloor \quad (6c)$$

hold. Whereas the first two conditions can be derived from the fact that only deeper levels can be reached by splitting a node, the third condition reflects that the descendant subband \tilde{k}_i is consistent with the subband k_i of the ancestor node. To this end, we use the inverse of the parent-child relationship (2), which has been extended to cover multiple generations. In the following, the set of all descendants of $\tilde{\mathcal{B}}$ in \mathcal{T} is denoted by $\tilde{\mathcal{T}}_d = \{\mathcal{B} \in \mathcal{T} \mid \tilde{\mathcal{B}} \succ \mathcal{B}\} \subset \mathcal{T}$.

In the second step, the dimension in which the data has to be filtered is determined. It is then tested for all descendants $\mathcal{B} \in \tilde{\mathcal{T}}_d$ in which dimension i their level s_i is deeper than the ancestor level \tilde{s}_i . Similarly, for the filtering in dimension i to be allowed, *all* descendants need to have a deeper level $s_i > \tilde{s}_i$. Otherwise, at least one descendant would not be reachable from the current node's children anymore. Summarized, the sets of the necessary and allowed dimensions

$$\mathcal{I}_{\text{nec}} = \{i \in [1, N] \mid \exists_{\mathcal{B} \in \tilde{\mathcal{T}}_d} : s_i > \tilde{s}_i\}, \quad (7a)$$

$$\mathcal{I}_{\text{allow}} = \{i \in [1, N] \mid \forall_{\mathcal{B} \in \tilde{\mathcal{T}}_d} : s_i > \tilde{s}_i\}, \quad (7b)$$

are defined. Consequently, the next dimension to filter is chosen from the intersection of the necessary and allowed dimensions

$$\mathcal{I} = \mathcal{I}_{\text{nec}} \cap \mathcal{I}_{\text{allow}}. \quad (8)$$

If there are no redundant nodes in \mathcal{T} , the intersection \mathcal{I} contains at least one element. In case \mathcal{I} contains multiple elements, simply the first element is chosen, due to the convolution being commutative.

In the third step, the actual filtering of the data is carried out. Given the input data $c[n_1, \dots, n_N]$ of the current recursive call, the lowpass and the bandpass child are calculated via

$$c_{\text{LP}}[\dots, n_i, \dots] = \left(c[n_1, \dots, n_N] *_{i} g_{\text{LP}}[n] \right) \Big|_{n=2n_i}, \quad (9a)$$

$$c_{\text{BP}}[\dots, n_i, \dots] = \left(c[n_1, \dots, n_N] *_{i} g_{\text{BP}}[n] \right) \Big|_{n=2n_i}, \quad (9b)$$

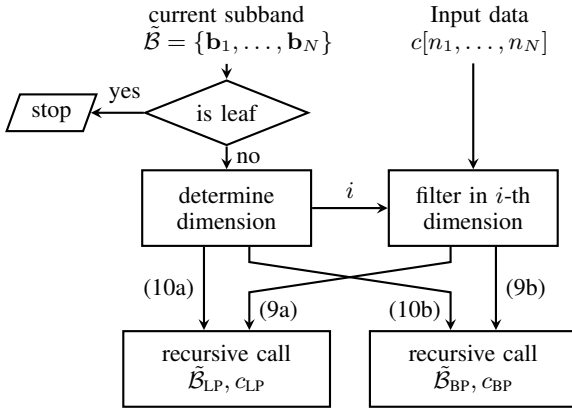


Figure 2. Flowchart of the decomposition algorithm.

where $*_i$ denotes a convolution in the i -th dimension. Equation (9) describes the filtering step as a convolution along the i -th dimension using the finite impulse responses g_{LP} , g_{BP} and a subsequent downsampling by a factor of 2. Consequently, the number of samples of the data in the i -th dimension is halved.

In the last step of the algorithm, the next recursive calls are invoked using the children

$$\tilde{\mathbf{b}}_{LP} = \{\tilde{\mathbf{b}}_1, \dots, (\tilde{s}_i + 1, 2\tilde{k}_i), \dots, \tilde{\mathbf{b}}_N\}, \quad (10a)$$

$$\tilde{\mathbf{b}}_{BP} = \{\tilde{\mathbf{b}}_1, \dots, (\tilde{s}_i + 1, 2\tilde{k}_i + 1), \dots, \tilde{\mathbf{b}}_N\} \quad (10b)$$

together with the filtered data (9), since each recursive call requires both the input data and a (new) current node.

Figure 2 shows a summary of the proposed algorithm. After evaluating the stopping criterion, the next filtering dimension is determined, followed by the actual filtering step and the new recursive calls for the child nodes.

C. Dual-tree extension to multiple dimensions

So far, the data is only filtered with real-valued finite impulse response filters which suffer from a lacking shift-invariance (due to the downsampling) and oscillations in the time-frequency domain, even if the signal energy is constant within a frequency subband. To overcome these drawbacks, an additional tree, whose filters build a Hilbert pair with the corresponding filters of the original WPT tree, has been introduced by Kingsbury [14]. A generalization of the so-called dual-tree approach to multiple dimensions leads to the multi-dimensional complex wavelet functions

$$\psi^C[n_1, \dots, n_N] = \psi_1^C[n_1] \cdots \psi_N^C[n_N], \quad (11)$$

with

$$\psi_i^C[n_i] = \psi_i^{\text{Re}}[n_i] + j\psi_i^{\text{Im}}[n_i], \quad (12)$$

where $\psi_i^{\text{Re}}[n_i]$ and $\psi_i^{\text{Im}}[n_i]$ denote the resulting wavelet functions of the original and its dual tree of the i -th dimension, respectively. For further reference on the filter design to obtain the analytic basis functions $\psi_i^C[n_i]$, see [8], [10], [15]. Note that Eq.(11) implies the usage of separable filter kernels without limitation to isotropic decomposition structures. By

replacing some of the complex wavelets $\psi_i^C[n_i]$ with their complex conjugates, the support of the wavelet in the corresponding frequency domain can be adjusted.

Equation (11) can be implemented using the multiplication principle to calculate 2^N WPT trees with all possible combinations in $\{\psi_1^{\text{Re}}, \psi_1^{\text{Im}}\} \times \{\psi_2^{\text{Re}}, \psi_2^{\text{Im}}\} \times \cdots \times \{\psi_N^{\text{Re}}, \psi_N^{\text{Im}}\}$. Depending on the desired support of the complex wavelet ψ^C , the AWPT of the input signal \mathbf{X} is obtained by the linear combination

$$\text{AWPT}(\mathbf{X}) = \sum_{n=0}^{2^N-1} a_n \cdot \text{WPT}_n(\mathbf{X}) \quad (13)$$

of the individual trees WPT_n . To reduce the memory requirements, only basis functions with support in the positive frequency domain are used. For this, the coefficients a_n are defined by the sequence $(a_n)_{n=0}^{2^N-1}$ through

$$a_0 = 1, \quad a_n = j \cdot a_{n \bmod m[n]}, \quad (14)$$

where $m[n] = 2^{\lfloor \text{ld } n \rfloor}$ and \bmod denotes the modulo operation. For every dimension i individually, our framework offers the possibility to choose if a real or complex-valued wavelet should be used, i.e. whether the wavelet $\psi_i^C[n_i] = \psi_i^{\text{Re}}[n_i]$ or $\psi_i^C[n_i] = \psi_i^{\text{Re}}[n_i] + j\psi_i^{\text{Im}}[n_i]$ is used. Using real-valued wavelets reduces the number of necessary trees and the computational complexity at the cost of a missing shift-invariance.

III. SIGNAL-ADAPTED FRAME REPRESENTATION

For the signal synthesis, only the leaves of the tree are needed. If the intermediate nodes are also kept, a redundancy is introduced, leading to a frame (instead of a basis). The frame can then be used to determine an optimal, signal-adapted basis. As mentioned in Section II, the frame can be defined using an extended set $\mathcal{T}_{\text{Frame}} \supset \mathcal{T}$ that contains both the leaves and all intermediate nodes. Each node in $\mathcal{T}_{\text{Frame}}$ is calculated by the decomposition algorithm and evaluated using the Shannon entropy as the cost functional. Finally, by selecting the best nodes in terms of entropy, we remove all redundant nodes and find the decomposition structure \mathcal{T}_{SA} , which we call the Signal-Adapted WPT (SA-WPT) in the real-valued case and the SA-AWPT in the complex-valued case.

First, the frame is defined. For each dimension i , the deepest level L_i has to be specified, resulting in $M_i = 2^{L_i+1} - 2$ tuples per dimension if all intermediate nodes are considered. By combination of all these tuples, using the multiplication principle, we obtain the frame

$$\mathcal{T}_{\text{Frame}} = \{\mathbf{b}_{1,1}, \dots, \mathbf{b}_{1,M_1}\} \times \cdots \times \{\mathbf{b}_{N,1}, \dots, \mathbf{b}_{N,M_N}\}, \quad (15)$$

containing $M = M_1 \cdot M_2 \cdots M_N$ nodes.

Consequently, the set \mathcal{I} , containing the next possible dimensions to filter, may be empty, because for every dimension i , there exists a descendant node that violates $s_i > \tilde{s}_i$ in (7b). Hence, if \mathcal{I} is empty, for every $i \in \mathcal{I}_{\text{nec}}$, an additional recursion is started, i.e. the algorithm steps (9) and (10) are repeated for all i . Since the additional recursions are being executed independently and the same nodes may be accessible by different recursions, many redundant calculations would be

performed. To avoid these calculations, a globally accessible set, containing the nodes that have already been calculated, is used in order to define an additional stopping criterion, namely: stop the recursion if the current node has already been calculated by another, independent recursion. Apart from that, the decomposition is performed analogously to Section II-B.

After the decomposition of the signal, every node $\mathcal{B} \in \mathcal{T}_{\text{Frame}}$ represents a subspace of the frequency domain and contains the filtered data. By calculating the entropy

$$H(\mathcal{B}, \mathbf{c}) = - \sum_{k=1}^K \frac{|c_k|^2}{E_0} \ln \frac{|c_k|^2}{E_0}, \quad (16)$$

the sparsity of the data in these subspaces can be evaluated [13]. Here, $\mathbf{c} \in \mathbb{C}^K$ denotes the node's vectorized filtered data and E_0 the energy of the input signal. In the one-dimensional case, the optimal basis selection is performed by using a recursion starting with the root of the tree. In each recursive call, if the node is not a leaf, the cost is calculated as the minimum of its entropy and the summed costs of its children. In case the node is a leaf, its cost is equal to its entropy. Generalizing this approach to arbitrary dimensions, there are now $2N$ direct children, i.e. children with *exactly* one \mathbf{b}_i fulfilling Eq. (2). The cost $J(\mathcal{B})$ of a node \mathcal{B} is therefore calculated using its direct children $\mathcal{B}_{i,\text{BP}}, \mathcal{B}_{i,\text{LP}}$, via

$$J(\mathcal{B}) = \min \{H(\mathcal{B}), J_{\text{desc}}\}, \quad \text{with} \quad (17a)$$

$$J_{\text{desc}} = \min_{i \in [1, N]} \{J(\mathcal{B}_i^{\text{BP}}) + J(\mathcal{B}_i^{\text{LP}})\}. \quad (17b)$$

Similarly to the decomposition, redundant calculations are avoided by using a set containing the nodes which have already been calculated. In addition, the recursion is used to create the set \mathcal{T}_{SA} . If the current node $\tilde{\mathcal{B}}$ is a leaf or $J(\tilde{\mathcal{B}}) = H(\tilde{\mathcal{B}})$ in (17a), a temporary set $\tilde{\mathcal{T}} = \{\tilde{\mathcal{B}}\}$, which contains only the current node, is returned by the recursive call. Otherwise, the minimizer i^* of (17b) is used to find the best direct children. In this case, the returned temporary set $\tilde{\mathcal{T}}$ is calculated as the union $\tilde{\mathcal{T}}_{i^*}^{\text{LP}} \cup \tilde{\mathcal{T}}_{i^*}^{\text{BP}}$, with $\tilde{\mathcal{T}}_{i^*}^{\text{LP}}$ and $\tilde{\mathcal{T}}_{i^*}^{\text{BP}}$ being the temporary sets returned by the lowpass and bandpass child of the best direct children, respectively. Finally, the set \mathcal{T}_{SA} , defining the signal-adapted decomposition, is obtained by the set $\tilde{\mathcal{T}}$ returned by the initial call of the root of the tree.

IV. LIGHT FIELD COMPRESSION AND DENOISING

As a high-dimensional example, we apply the proposed framework to 4D light fields. In the so-called plane-plane parametrization, a 4D light field $L(u, v, s, t)$ is parametrized by its angular and spatial coordinates (u, v) and (s, t) , respectively. In this parametrization, a light field can be thought of as a multi-view collection of images: for fixed (u, v) , the subapertures $I_{uv}(s, t) = L(u, v, s, t)$ are regular 2D images.

Using the proposed framework, we apply the 4D DWT, DT-CWT, (A)WPT, and SA-(A)WPT to a light field compression and denoising. For all wavelet-based transforms, Daubechies-3 and Daubechies-5 filters [15] are applied in the (u, v) and (s, t) dimensions, respectively. These filters were chosen because they show a good trade-off between filter length and possible

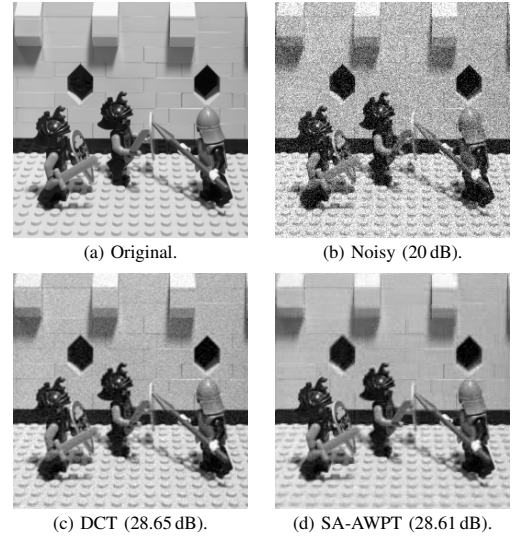


Figure 3. Central views of the original, noisy and denoised light fields.

tree depth (which is determined by the size of the data and the filter length). For the analytic transforms, starting from level two, the Q-shift-2 filters [10] are applied in order to obtain (approximately) analytic wavelet functions. As a reference, we will compare these methods to the 4D Discrete Cosine Transform (DCT). Since it would be out of the scope of this paper, we do not perform a full-scale statistical analysis of the compression and denoising performance but merely use a single example light field to demonstrate the potential of multi-dimensional analytical wavelet packets. To this end, we use the *Lego Knights* light field of the well-known Stanford Light Field Archive [16], which was converted to grayscale and downsampled to a size of $(16, 16, 512, 512)$. The central view $I_{8,8}(s, t)$ of the original light field is shown in Fig. 3 (a).

A. Compression

Sparse representations are well suited for signal compression. Most straightforwardly, the transformed signal's smallest coefficients are thresholded to zero and the resulting signal energy is measured. In a sparse representation, the signal energy is concentrated in only a few large coefficients. To investigate the sparsity of the different transforms, we apply hard thresholding to the transformed light field for different inverse compression ratios $\alpha = N_{\text{nz}}/N$, where N_{nz} is the number of non-zero elements and N is the total number of elements of the light field. The results in terms of the compressed signal's Peak Signal-to-Noise Ratio (PSNR) for the different transforms are shown in Fig. 4. Overall, the SA-WPT performs best, while the DCT performs slightly better for very high compression ratios. It can be observed that the regular WPT performs worst for most compression ratios, which is to be expected since most of the signal energy is contained in low frequencies. Due to the low frequency characteristic of the light field's angular components, the wavelet approach cannot unfold its full potential for high compression ratios when compared to the DCT, which is well suited for nearly constant signals.




Figure 4. Compression results of the example light field for various bases.

B. Denoising

To investigate the denoising capabilities of the different wavelet approaches, the original light field is superposed with uncorrelated white Gaussian noise (PSNR = 20 dB), as shown in Fig. 3 (b). Denoising is performed by applying a hard threshold in the transformed domain, where the threshold value is estimated using the commonly used VisuShrink estimator [17]. After thresholding, the signals are transformed back to the original domain and the resulting PSNR values are calculated by comparing the original light field and the denoised one. The results are shown in Fig. 5. Except for the DCT, all transforms were carried out using the proposed framework. While, again, the DCT performs quite well in terms of the denoised PSNR, the visual quality, measured using the Structural Similarity Index Metric (SSIM) [18], is inferior to those of the signal adapted and analytic wavelet approaches. This is in agreement with the visual impression of the denoising results as shown in Fig. 3 (c, d) for the DCT and SA-AWPT. While both denoising results show an almost identical PSNR value, the visual quality of the results is vastly different. It can be observed that the DCT successfully removes the high-frequency components mostly associated with noise resulting in a high overall PSNR value. However, due to the lack of high frequencies, the result appears less sharp, while low and mid-frequency noise is still clearly visible. On the other hand, the SA-AWPT is naturally more adapted to the underlying data, preserving edges whilst restoring homogenous areas of the light field nicely. However, the SA-AWPT denoising introduces some visible wavelet-related artefacts. Concluding, the potential of the (analytical) wavelet and wavelet packet transform for light field data has been demonstrated.

V. CONCLUSION

The presented framework for the N-D analytical wavelet packet transform unifies the discrete wavelet transform, the wavelet packet transform and their analytic counterparts. Due to the novel notation of the decomposition structure, flexible decompositions are possible, including isotropic, anisotropic and redundant representations. Using the proposed framework, a 4D light field compression and denoising was performed. In both cases, the signal adapted and analytic transforms showed promising results, improving the SSIM and/or PSNR over the standard wavelet and discrete cosine transform.




Figure 5. Denoising results of the example light field for various bases.

REFERENCES

- [1] S. G. Mallat, "A theory for multiresolution signal decomposition: the wavelet representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 7, pp. 674–693, 1989.
- [2] J.-P. Antoine, P. Carrette, R. Murenzi, and B. Piette, "Image analysis with two-dimensional continuous wavelet transform," *Signal Processing*, vol. 31, no. 3, pp. 241–272, 1993.
- [3] M. Antonini, M. Barlaud, P. Mathieu, and I. Daubechies, "Image coding using wavelet transform," *IEEE Transactions on Image Processing*, vol. 1, no. 2, pp. 205–220, 1992.
- [4] A. S. Lewis and G. Knowles, "Image compression using the 2-d wavelet transform," *IEEE Transactions on Image Processing*, vol. 1, no. 2, pp. 244–250, 1992.
- [5] X. Cao, T. Yue, X. Lin, S. Lin, X. Yuan, Q. Dai, L. Carin, and D. J. Brady, "Computational snapshot multispectral cameras: Toward dynamic capture of the spectral world," *IEEE Signal Processing Magazine*, vol. 33, no. 5, pp. 95–108, 2016.
- [6] E. H. Adelson and J. Y. Wang, "Single lens stereo with a plenoptic camera," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, pp. 99–106, 1992.
- [7] Z. Xiong, L. Wang, H. Li, D. Liu, and F. Wu, "Snapshot hyperspectral light field imaging," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [8] I. W. Selesnick, "On the dual-tree complex wavelet packet and m -band transforms," *IEEE Transactions on Signal Processing*, vol. 56, no. 6, pp. 2298–2310, 2008.
- [9] T. Weickert, C. Benjaminsen, and U. Kiencke, "Analytic wavelet packets—combining the dual-tree approach with wavelet packets for signal analysis and filtering," *IEEE Transactions on Signal Processing*, vol. 57, no. 2, pp. 493–502, 2009.
- [10] N. Kingsbury, "Design of q -shift complex wavelets for image processing using frequency domain energy minimization," in *IEEE International Conference on Image Processing*, 2003.
- [11] I. W. Selesnick, R. G. Baraniuk, and N. G. Kingsbury, "The dual-tree complex wavelet transform," *IEEE Signal Processing Magazine*, vol. 22, no. 6, pp. 123–151, 2005.
- [12] Institute of Industrial Information Technology, Karlsruhe Institute of Technology, "Public GitLab repositories," <https://gitlab.com/iit-public>, GNU GPLv3 License, 2020.
- [13] R. R. Coifman and M. V. Wickerhauser, "Entropy-based algorithms for best basis selection," *IEEE Transactions on Information Theory*, vol. 38, no. 2, pp. 713–718, 1992.
- [14] N. G. Kingsbury, "The dual-tree complex wavelet transform: a new technique for shift invariance and directional filters," in *IEEE Digital Signal Processing Workshop*, vol. 86, 1998.
- [15] I. Daubechies, *Ten lectures on wavelets*, ser. Regional Conference Series in Applied Mathematics. Philadelphia, Pa.: Society for Industrial and Applied Mathematics, 1992.
- [16] Computer Graphics Laboratory, Stanford University, "The (new) Stanford light field archive," <https://lightfield.stanford.edu/>, 2008.
- [17] D. L. Donoho and J. M. Johnstone, "Ideal spatial adaptation by wavelet shrinkage," *Biometrika*, vol. 81, no. 3, pp. 425–455, 1994.
- [18] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004.