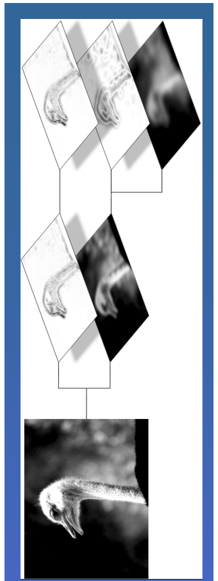

Tutorial presented at: 6th International Symposium on Signal Processing and its Applications (ISSPA) 2001

Wavelets - from Theory to Applications

Claudia Schremmer
University of Mannheim / Germany

November, 2001

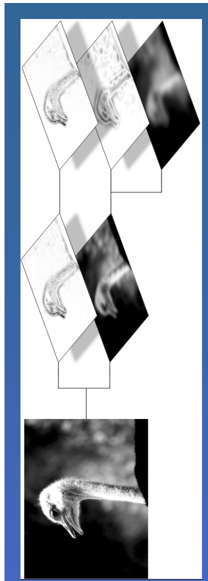

Overview (I)

- Part I: Wavelets
  - 1.1 Historic Outline
  - 1.2 The Wavelet Transform
  - 1.3 Multiscale Analysis
  - 1.4 Transformation Based on the Haar Wavelet
- Part II: Implementation Issues
  - 2.1 Wavelets in Multiple Dimensions
  - 2.2 Signal Boundary
  - 2.3 Painting the Time-scale Domain
  - 2.4 Lifting

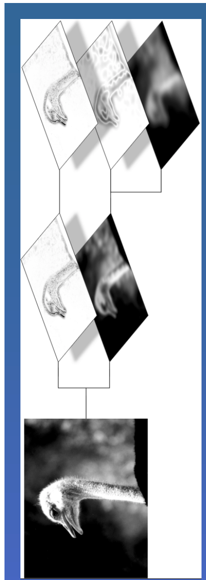

Overview (II)

- Part III: Applications of Wavelets in Multimedia
  - 3.1 JPEG2000
  - 3.2 Layered Wavelet Video Coding
- Part IV: Java Applets for demonstration

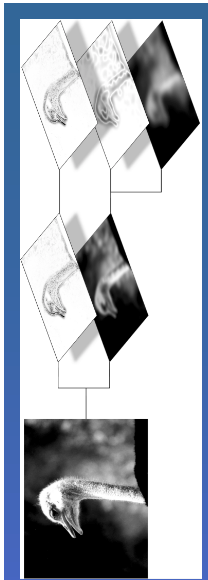

1.1 Historic Outline

- Wavelet theory combines pure and applied mathematics, physics, computer science, engineering, etc.
- **1981.** Morlet: kept the number of oscillations within a window constant, varying the width of the window.

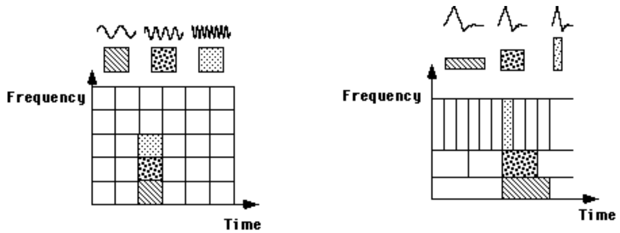

- **1985.** Grossmann: discrete wavelet transform is reversible.
- **1985.** Meyer: prove of existence of orthogonal wavelets.
- **1986.** Mallat and Meyer: multiscale analysis
- **1992.** Daubechies: orthog. wavelets with compact support.
- **Since then.** Wavelet analysis evolved from a mathematical curiosity to a major foundation of signal processing algorithms.

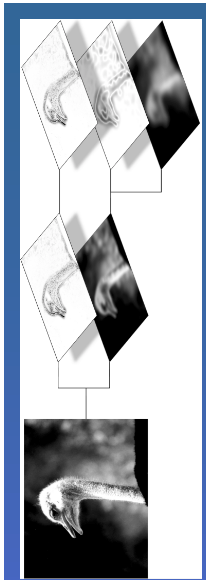

1.2 The Wavelet Transform

Definition. A wavelet is a function $\psi \in L_2(\mathbb{R})$ which meets the admissibility condition

$$0 < c_\psi := 2\pi \int_{\mathbb{R}} \frac{|\hat{\psi}(\omega)|^2}{|\omega|} d\omega < \infty,$$

where $\hat{\psi}$ denotes the Fourier transform of the wavelet ψ. The constant c_ψ designates the admissibility constant.

It follows that a wavelet integrates to zero:

$$0 = \hat{\psi}(0) = \int_{\mathbb{R}} \psi(t) e^{-2i\pi t0} dt = \int_{\mathbb{R}} \psi(t) dt$$

Thus, a wavelet has the same volume ‘above the x-axis’ as ‘below the x-axis’. This is where the name originates.

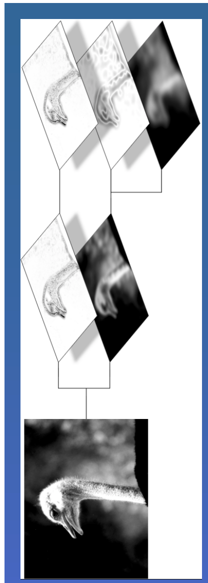

Example Wavelets (I)

- Haar wavelet

$$\psi(t) = \begin{cases} 1 & : \quad 0 \leq t < \frac{1}{2}, \\ -1 & : \quad \frac{1}{2} \leq t \leq 1, \\ 0 & : \quad \text{else.} \end{cases}$$

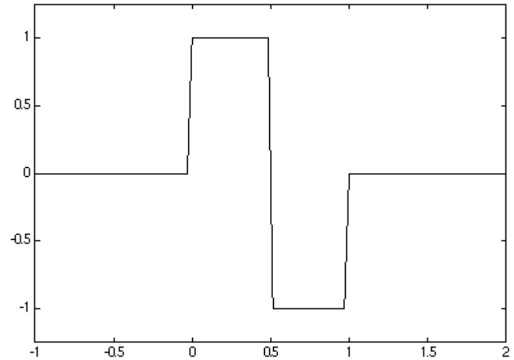

- Mexican Hat wavelet

$$\psi(t) = -\frac{d^2}{dt^2} e^{-t^2/2} = (1 - t^2) e^{-t^2/2}$$

which is the second derivative of a Gaussian.

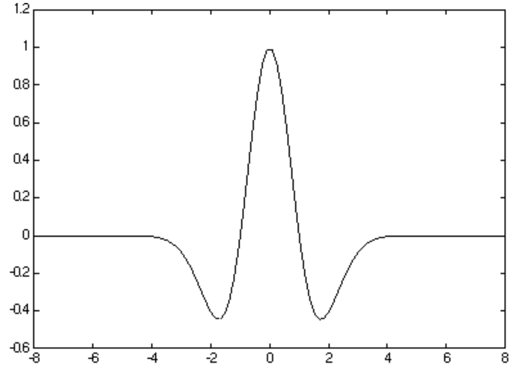

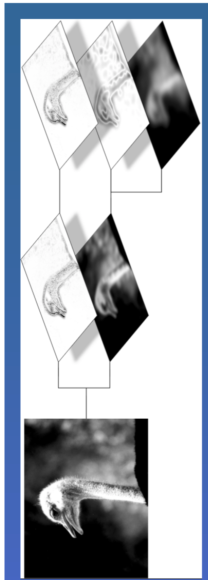

Example Wavelets (II)

- Morlet wavelet

  is defined via its Fourier transform:

$$\hat{\psi}(\omega) = e^{-2\pi^2(\omega-\omega_0)^2}$$

  and decomposes into two parts, a real and an imaginary one.

$$\psi_{\Re}(t) = \frac{1}{\sqrt{2\pi}} e^{-t^2/2} \cos(2\pi\omega_0 t)$$

$$\psi_{\Im}(t) = \frac{1}{\sqrt{2\pi}} e^{-t^2/2} \sin(2\pi\omega_0 t)$$

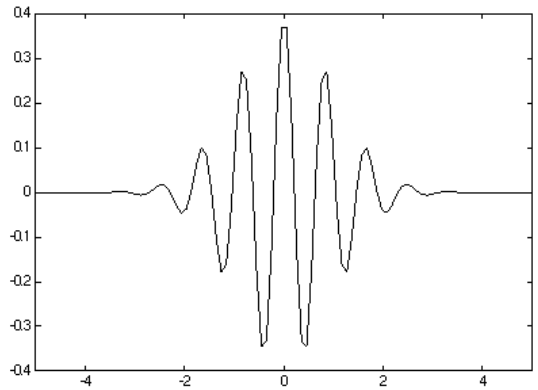

- Daubechies wavelet

  are obtained by iteration;

  no closed representation exists.

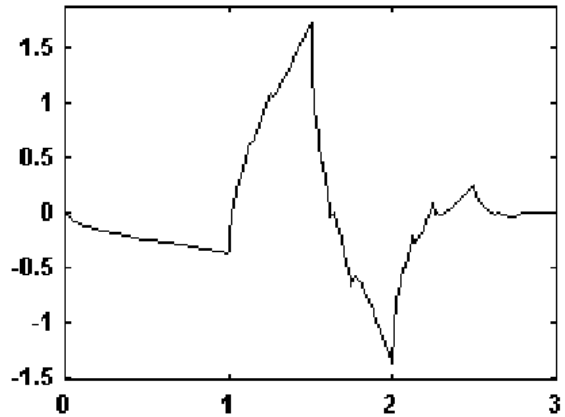

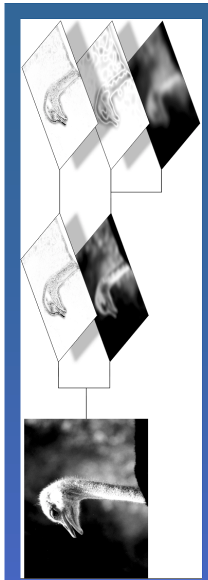

Which Wavelet?

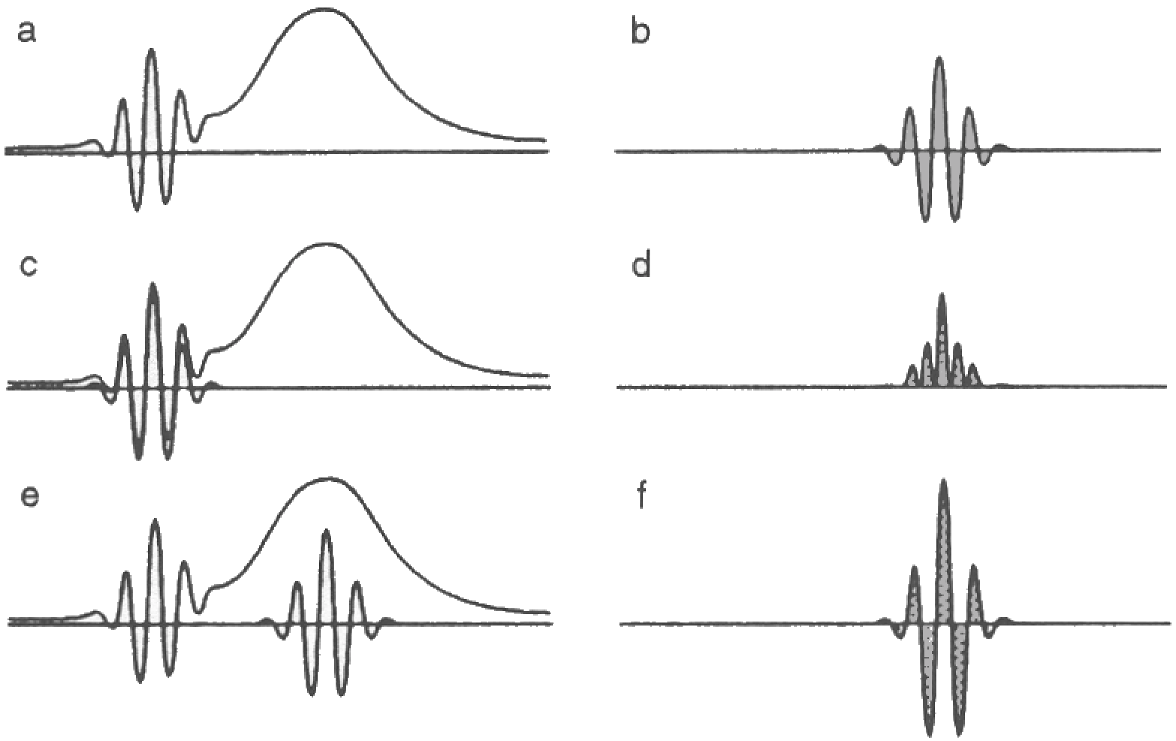

a) Original signal

b) Wavelet

c) Wavelet analyzes the signal at a position where both shapes are similar

d) The integral is large, indicating large similarity

e) Wavelet analyzes the signal at a position where both shapes differ largely

f) The integral almost vanishes, indicating no similarity

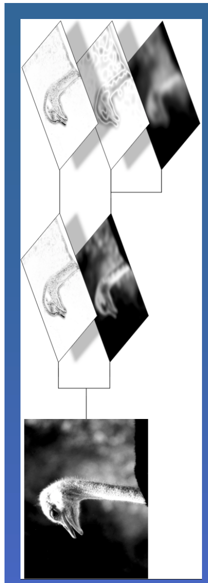

Integral Wavelet Transform

Definition. The integral wavelet transform of a function $f \in L_2(\mathbb{R})$ with regard to the admissible wavelet ψ is given by

$$f \longmapsto \tilde{f}_\psi(a,b) := \frac{1}{\sqrt{a}} \int_{\mathbb{R}} f(t)\psi^* \left(\frac{t-b}{a}\right) dt = \int_{\mathbb{R}} f(t)\psi^*_{a,b}(t)dt$$

where ψ^* is the complex conjugate of ψ.

$a > 0$ is called the *dilation factor* and b is the *translation* parameter, thus $\psi_{a,b}$ denotes a dilated and translated wavelet.

Remarks:

1. The wavelet transform is linear.
2. A one-dim. signal is transformed into a two-dim. space.

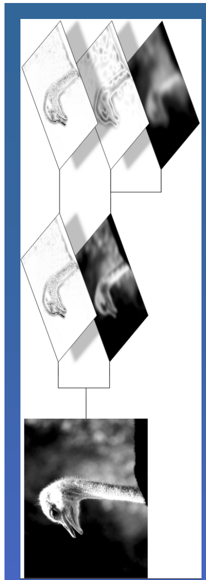

Wavelet Basis (I)

A wavelet transform decomposes a signal f into coefficients for a corresponding wavelet ψ. Since all wavelets ‚live' in $L_2(\mathbb{R})$, we would like to know whether *every* function $f \in L_2(\mathbb{R})$ can be approximated with arbitrary precision. This is the case: The set of wavelets

$$\Psi = \{\psi \in L_2(\mathbb{R}) \, : \, \psi \text{ is admissible}\}$$

is a *dense* subset of $L_2(\mathbb{R})$. That is, every function in $L_2(\mathbb{R})$ can be approximated by wavelets, and the approximation error gets arbitrarily small.

Moreover, we can restrict the ‚pool of wavelet base functions' to dilated and translated versions of one *mother wavelet* ψ.

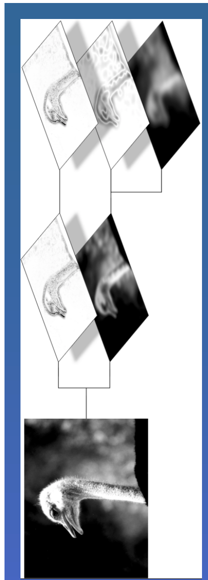

Wavelet Basis (II)

Finally, the parameter $a > 0$ which steers the dilation of the wavelet ψ can be restricted further:

The *dyadic* wavelet transform of f,

$$\tilde{f}_\psi(2^j, b) = \frac{1}{\sqrt{2^j}} \int_{\mathbb{R}} f(t)\psi^*\left(\frac{t-b}{2^j}\right) dt$$

defines a complete and stable representation of f if the frequency axis is completely covered by dilated dyadic wavelets.

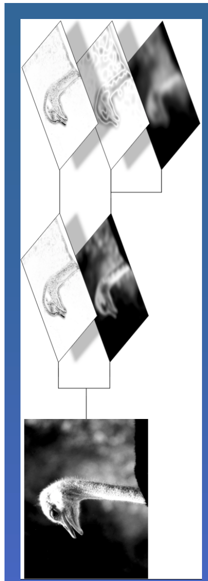

1.3 Multiscale Analysis

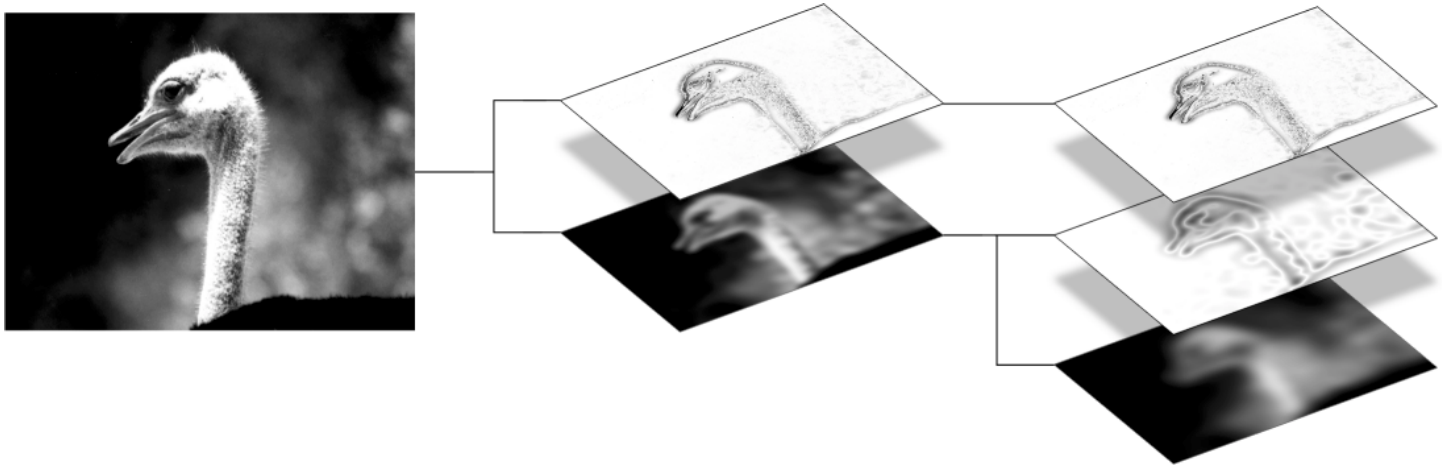

Successive decomposition of a signal into a series of *approximations* and *details*.

- Approximation: contains the low frequencies,
- Detail: ‚collects‘ the remaining high frequencies.

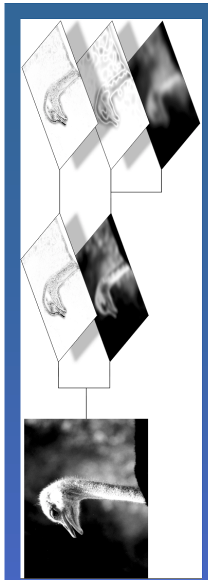

Projection onto Subspaces

In multiscale analysis, a signal $f \in L_2(\mathbb{R})$ is projected onto a subspace V_k of $L_2(\mathbb{R})$. The projection separates out the detail of the signal and only maintains the approximation on level k.

Iteration: $$V_k = V_{k+1} \oplus W_{k+1}, \qquad k = 0, 1, \ldots$$

Dyadic approach: $k = 2^j$

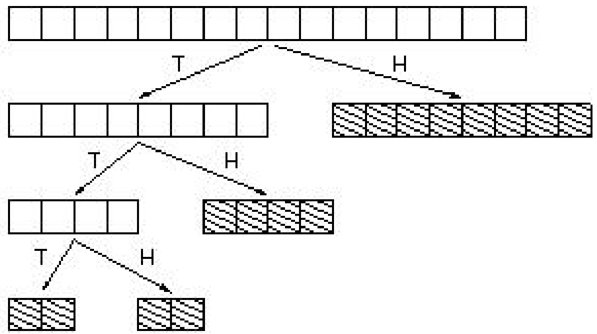

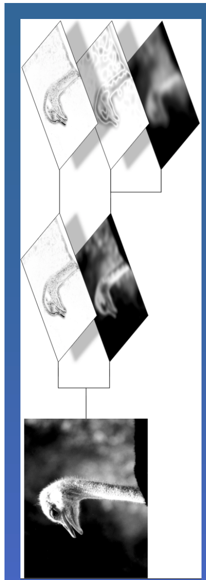

Approximation

Theorem. Let $\{V_{2^j}\}_{j\in\mathbb{Z}}$ be a series of closed nested subspaces:

$$\{0\} \subset \ldots \subset V_{2^{j+1}} \subset V_{2^j} \subset V_{2^{j-1}} \subset \ldots \subset L_2(\mathbb{R})$$

Then there exists a single function $\varphi \in L_2(\mathbb{R})$ such that

$$\left\{\frac{1}{\sqrt{2^j}}\varphi\left(\frac{t-k2^j}{2^j}\right)\right\}_{j,k\in\mathbb{Z}}$$

is an orthonormal base of V_{2^j}.

φ is called *scaling function*. Its explicit form is written as recursive difference:

$$\varphi(t) = \sqrt{2}\sum_{k\in\mathbb{Z}} h_0[k]\varphi(2t-k)$$

where h_0 is called the *filter mask*.

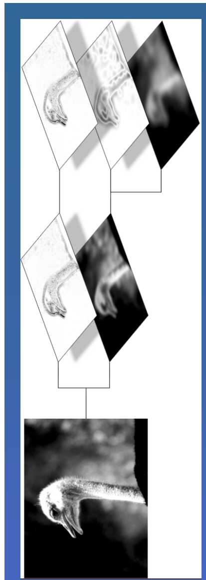

Example 1

- Let φ be the *indicator function* on [0,1). On the scale twice as fine, φ would need two representatives, i.e.,

$$\varphi(t) = \varphi(2t) + \varphi(2t-1)$$

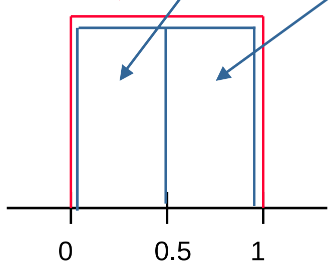

Here, the filter coefficients are:

$h_0[0] = h_0[1] = 1$ and $h_0[k] = 0$ else.

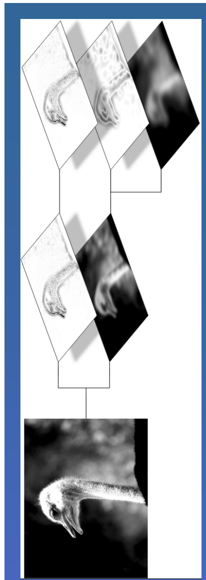

Example 2

- Let φ be the *hat function* on [0,2). On the scale twice as fine, φ would need three representatives, i.e.,

$$\varphi(t) = \frac{1}{2}\varphi(2t) + \varphi(2t-1) + \frac{1}{2}\varphi(2t-2)$$

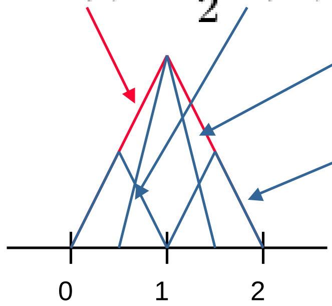

Here, the filter coefficients are:

$h_0[0] = h_0[2] = \frac{1}{2}, h_0[1] = 1$ and $h_0[k] = 0$ else.

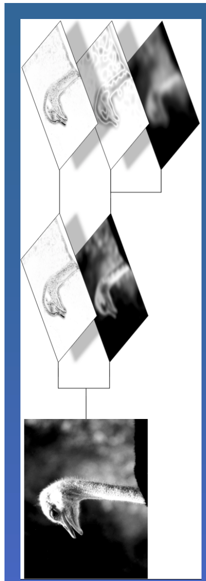

Detail

Theorem. Let $\{W_{2^j}\}_{j\in\mathbb{Z}}$ be a multiscale analysis of $L_2(\mathbb{R})$. Then there exists a single function $\psi \in L_2(\mathbb{R})$ such that

$$\left\{\frac{1}{\sqrt{2^j}}\psi\left(\frac{t-k2^j}{2^j}\right)\right\}_{j,k\in\mathbb{Z}}$$

is an orthonormal base of W_{2^j}.

ψ is called *orthogonal wavelet*. Its explicit form is written as recursive difference:

$$\psi(t) = \sqrt{2}\sum_{k\in\mathbb{Z}} h_1[k]\varphi(2t-k)$$

where h_1 is called the *filter mask*.

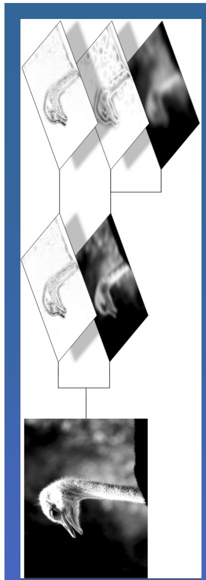

Summary: Spaces

Signal	Space
Original signal $f(t)$	$L_2(\mathbb{R})$
Signal is the sum of all its details	$L_2(\mathbb{R}) = \sum_{j\in\mathbb{Z}} W_{2^j}$
Detail in level 2^j	W_{2^j}
Approximation in Level 2^j	V_{2^j}
Relation between the approximation levels	$V_{2^j} = V_{2^{j+1}} \oplus W_{2^{j+1}}$
Decomposition of the signal	$L_2(\mathbb{R}) = V_{2^J} \oplus \sum_{j<J} W_{2^j}$

Relations between signals and spaces in multiscale analysis.

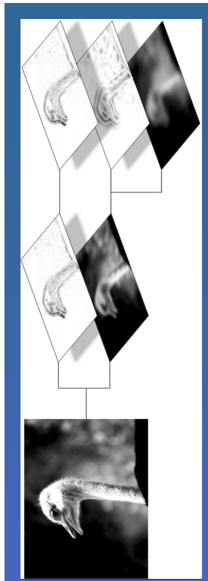

Summary: Subband Coding

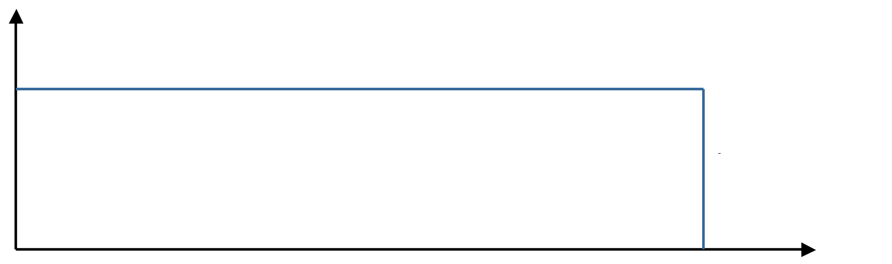

The original signal encompasses a certain frequency range.

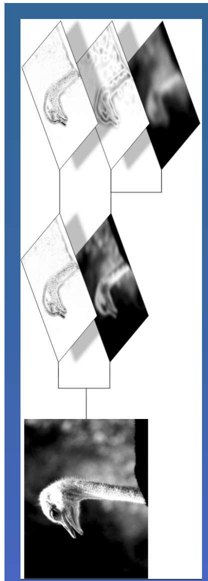

Summary: Subband Coding

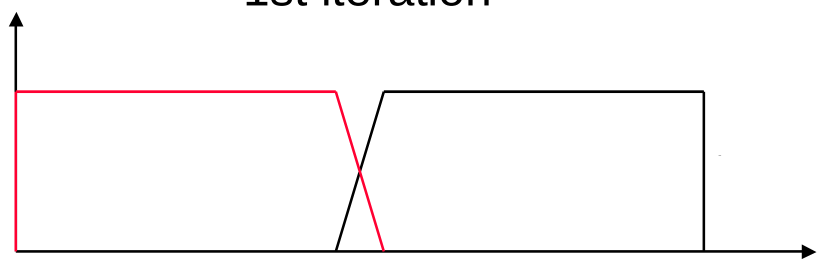

Subband coding:

In each iteration, half the resolution is ‚separated out' as details. The remaining approximation is then further subdivided. In each iteration, the scaling function determines the remaining approximation that sub-summarizes all the yet unconsidered parts.

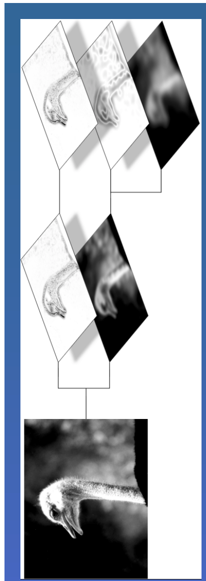

Summary: Subband Coding

2nd iteration

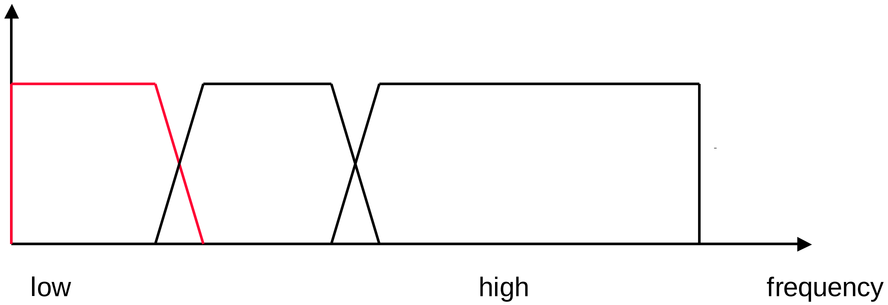

Subband coding:

In each iteration, half the resolution is ‚separated out‘ as details. The remaining approximation is then further subdivided. In each iteration, the scaling function determines the remaining approximation that sub-summarizes all the yet unconsidered parts.

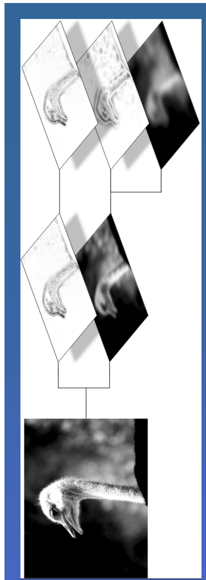

Summary : Subband Coding

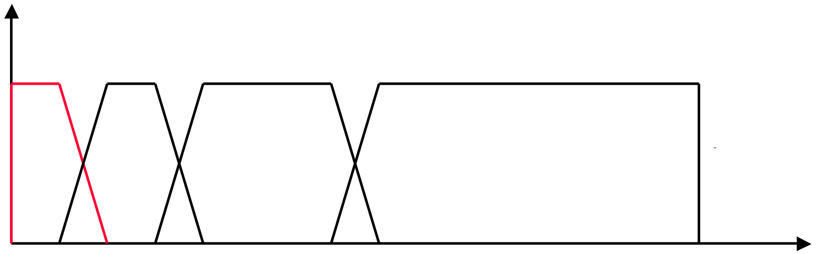

Subband coding:

In each iteration, half the resolution is ‚separated out‘ as details. The remaining approximation is then further subdivided. In each iteration, the scaling function determines the remaining approximation that sub-summarizes all the yet unconsidered parts.

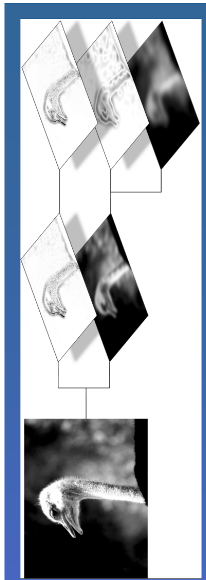

Summary: Tiling

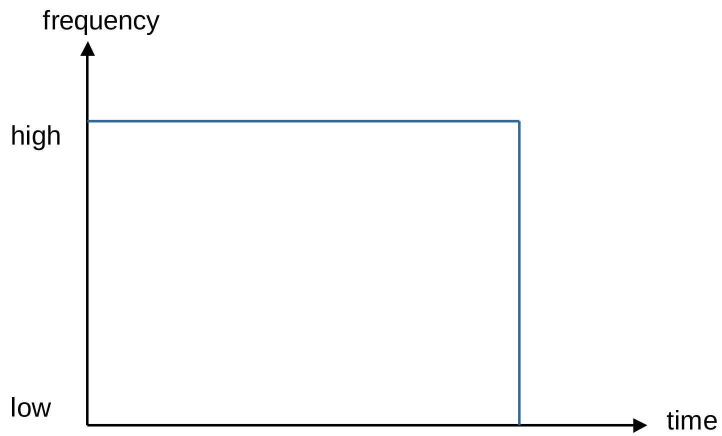

Tiling the time-scale domain for the dyadic wavelet transform.

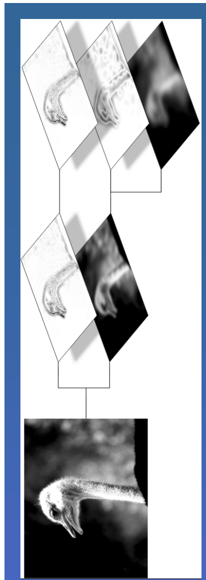

Summary: Tiling

1st iteration

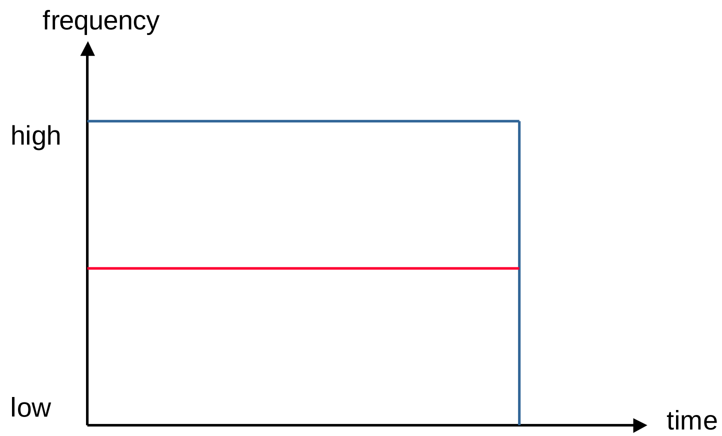

Tiling the time-scale domain for the dyadic wavelet transform.

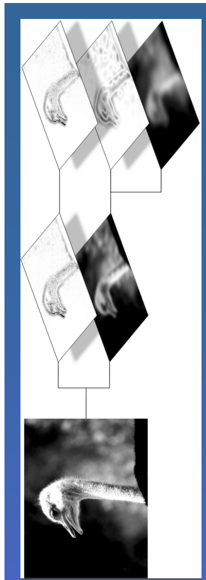

Summary: Tiling

2nd iteration

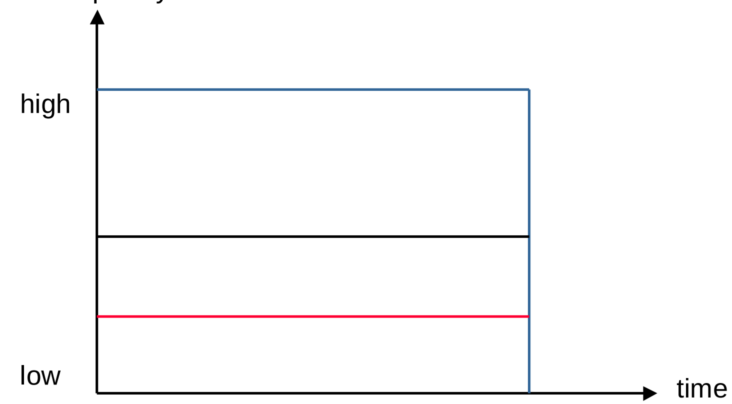

Tiling the time-scale domain for the dyadic wavelet transform.

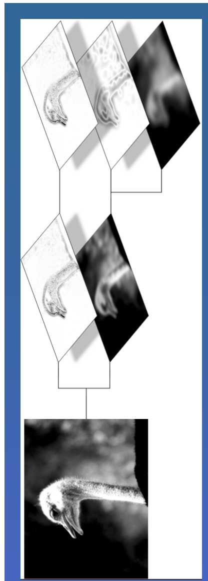

Summary: Tiling

3rd iteration

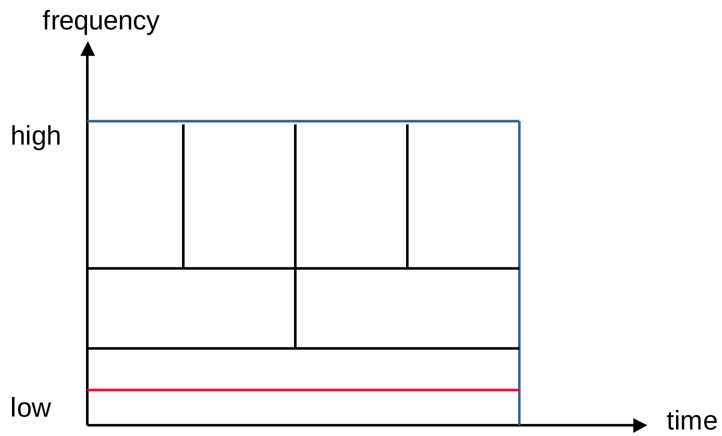

Tiling the time-scale domain for the dyadic wavelet transform.

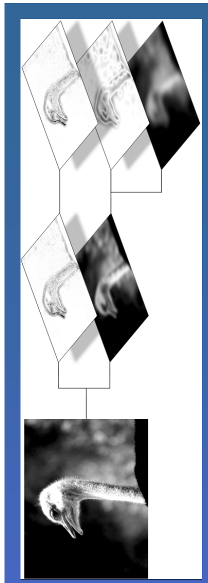

1.4 Transformation Based on the Haar Wavelet

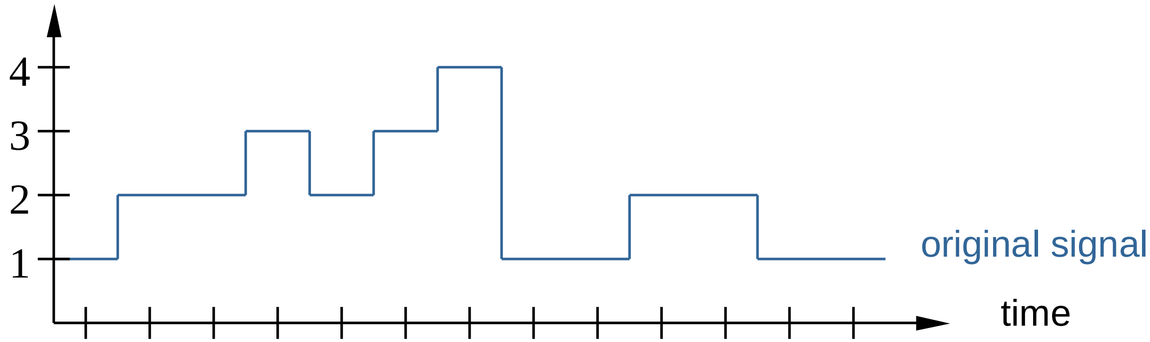

graphical representation

1	2	2	3	2	3	4	1	1	2	2	1	1

coefficients in time

A signal shall be approximated with fewer coefficients. An easy approach is to take the average of each two neighboring coefficients as approximation. The remaining error then is the difference of the ‚true‘ values towards these approximations.

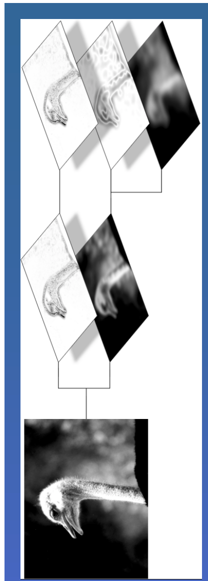

Haar Transform (I)

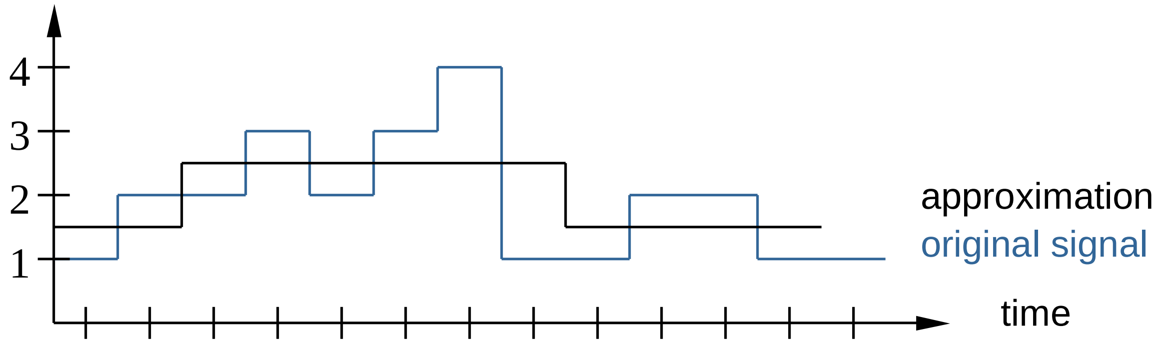

1	2	2	3	2	3	4	1	1	2	2	1	1	original signal
1.5		2.5		2.5		2.5		1.5		1.5			approximation

Filter for the computation of the average:

1/2	1/2

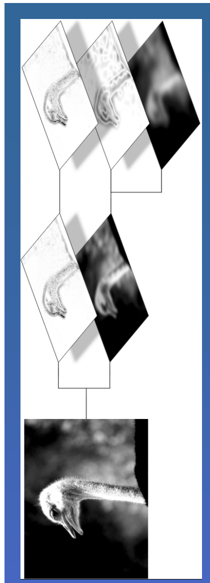

Haar Transform (II)

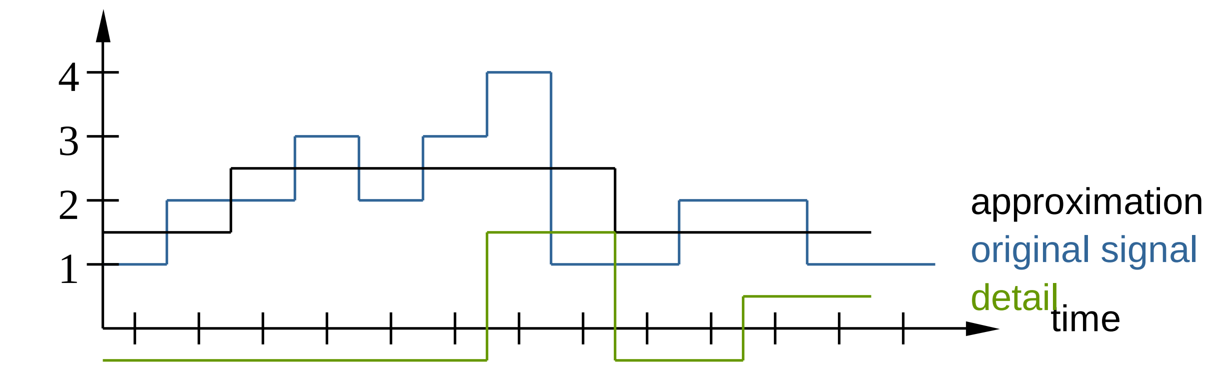

1	2	2	3	2	3	4	1	1	2	2	1	1	original signal
1.5		2.5		2.5		2.5		1.5		1.5			approximation
-0.5		-0.5		-0.5		1.5		-0.5		0.5			detail

Filter for the computation of the detail:

1/2	-1/2

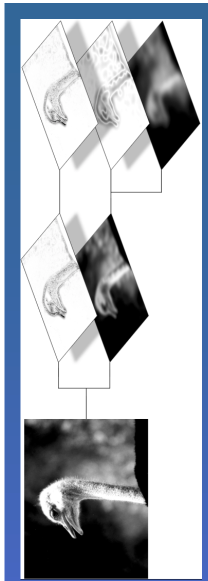

Haar Transform (III)

Synthesis: the information of the original signal can be recovered with *synthesis filters*.

1	1	synthesis filter for 1st signal entry
1	-1	synthesis filter for 2nd signal entry

Thus:

- 1.5*1+(-0.5)*1 = 1 (synthesis of 1. entry)
- 1.5*1+(-0.5)*(-1) = 2 (synthesis of 2. entry)

- 2.5*1+(-0.5)*1 = 2 (synthesis of 1. entry)
- 2.5*1+(-0.5)*(-1) = 3 (synthesis of 2. entry)

- ......

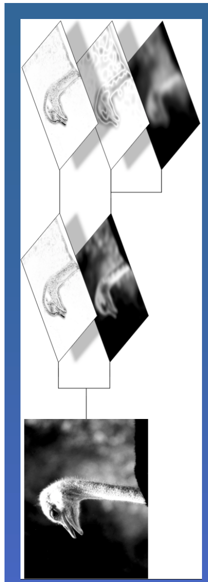

Haar Transform (IV)

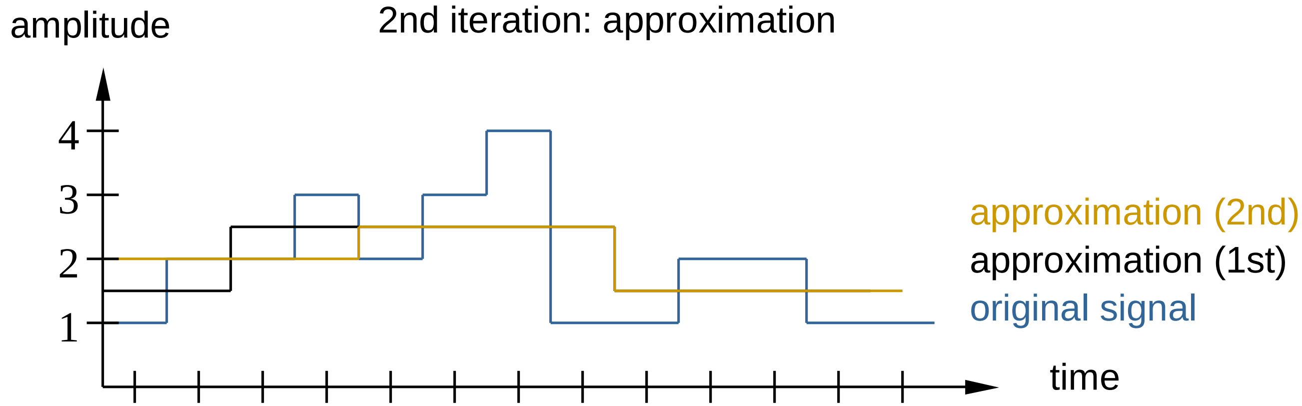

original signal	1	2	2	3	2	3	4	1	1	2	2	1	1
approximation (1st)	1.5		2.5		2.5		2.5		1.5		1.5		
approximation (2nd)	2				2.5				1.5				

Iteration on the approximation.

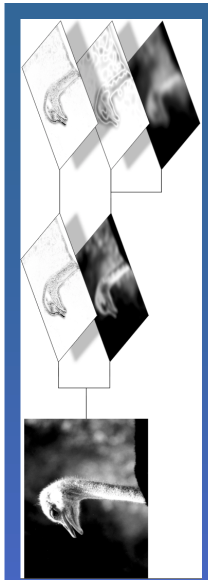

Haar Transform (V)

In total, we have used four filters for analysis and synthesis of a signal:

- approximation:
- detail:
- synthesis 1:
- synthesis 2:

1/2	1/2
1/2	-1/2
1	1
1	-1

In literature, the Haar filter is sometimes referred to as:

$1/\sqrt{2}$	$1/\sqrt{2}$
$1/\sqrt{2}$	$-1/\sqrt{2}$
$1/\sqrt{2}$	$1/\sqrt{2}$
$1/\sqrt{2}$	$-1/\sqrt{2}$

here, the factor $1/\sqrt{2}$ has been shifted from the analysis to the synthesis.

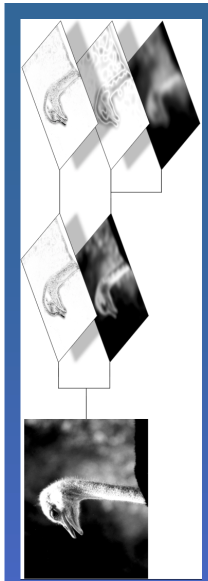

Convolution-based Transform

General filters are longer than the two entries of the Haar filter. The approach, however, to use two analysis and two synthesis filters, holds in general.

Even the longer filters are shifted by 2 signal coeffients.

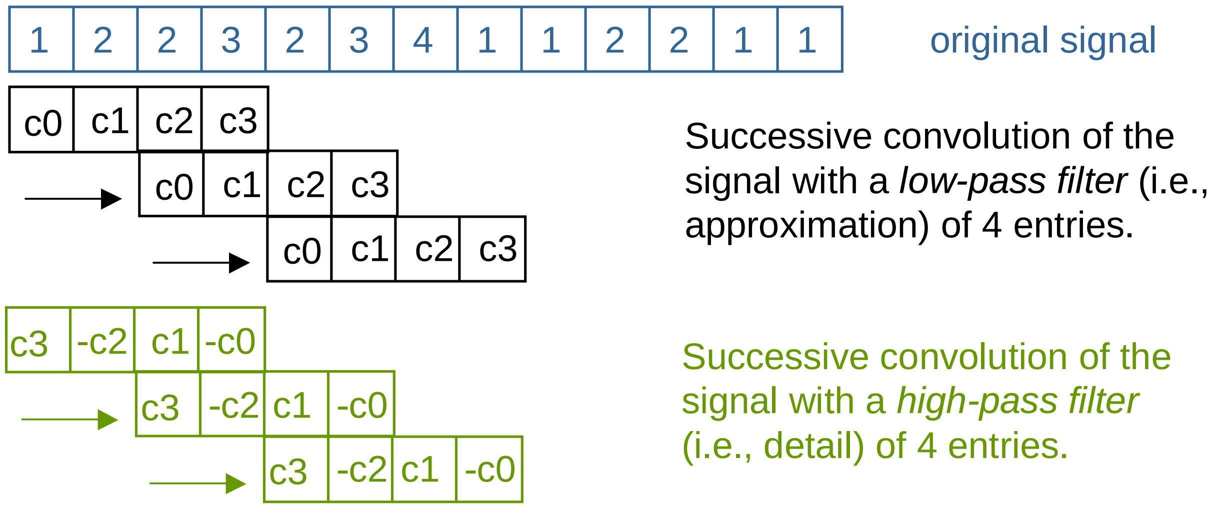

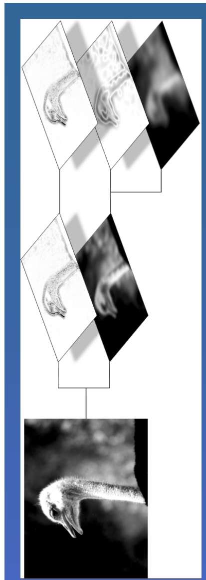

2.1 Wavelets in Multiple Dimensions

Until now: multiscale analysis in one dimension

$$V_{2^{-1}} = V_{2^J} \oplus W_{2^J} \oplus \ldots \oplus W_{2^1} \oplus W_{2^0}$$

Application of the wavelet transform on still images and video requires an approximation into multiple dimensions.

- *Separable* approach: successive application of a one-dimensional filter into one dimension and afterwards into a second dimension is mathematically identical to a two-dimensional transform from the outset.
- *Non-separable* approach: the *real* idea of multiple dimensions. Current research of groups around Kovacevic, Vetterli, and Tay.

Here: separable approach.

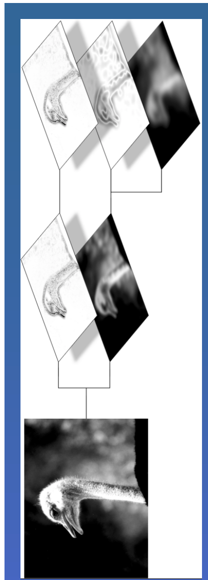

Separability

The separable wavelet transform on still images is defined via the tensor product, i.e., $V_0^{(2)} = V_0 \times V_0$.
This two-dimensional space decomposes into

$$\begin{aligned} V_0^{(2)} &= V_0 \times V_0 \\ &= (V_1 \oplus W_1) \times (V_1 \oplus W_1) \\ &= V_1 \times V_1 \oplus V_1 \times W_1 \oplus W_1 \times V_1 \oplus W_1 \times W_1 \\ &=: (\square). \end{aligned}$$

idea

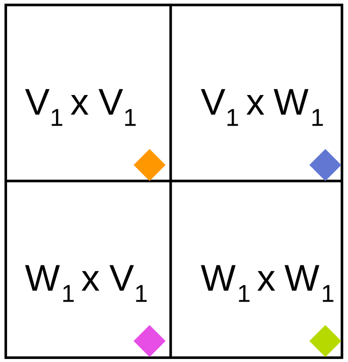

visualization

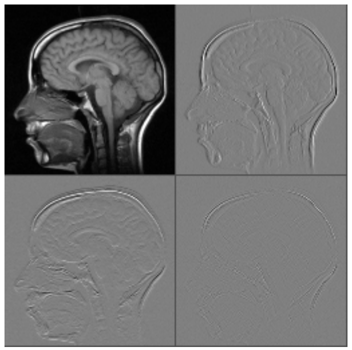

Frequency location in Fourier domain

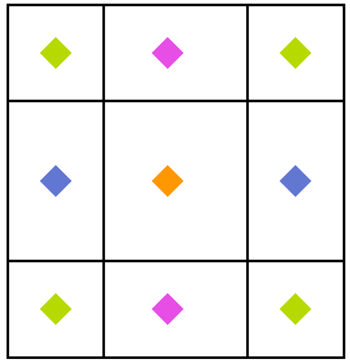

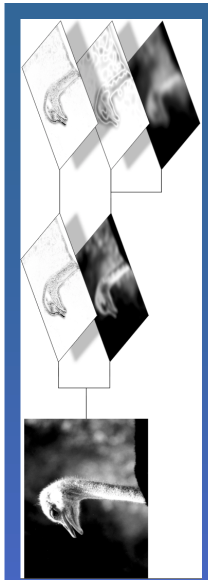

Standard Decomposition

In the following iteration steps, the standard decomposition iterates on *all* approximation spaces:

$$
\begin{aligned}
(\square) &= (V_2 \oplus W_2) \times (V_2 \oplus W_2) \oplus (V_2 \oplus W_2) \times W_1 \\
&\quad \oplus W_1 \times (V_2 \oplus W_2) \oplus W_1 \times W_1 \\
&= V_2 \times V_2 \oplus V_2 \times W_2 \oplus W_2 \times V_2 \oplus W_2 \times W_2 \\
&\quad \oplus V_2 \times W_1 \oplus W_2 \times W_1 \oplus W_1 \times V_2 \\
&\quad \oplus W_1 \times W_2 \oplus W_1 \times W_1
\end{aligned}
$$

4 iterations

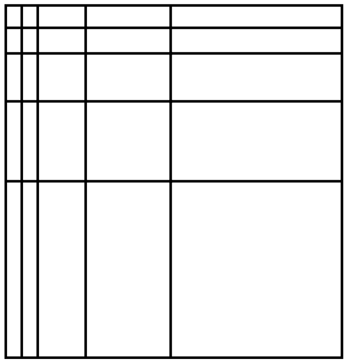

4 iterations - work in progress

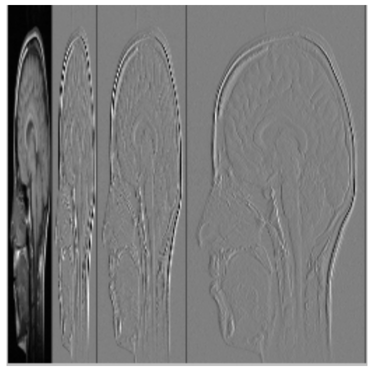

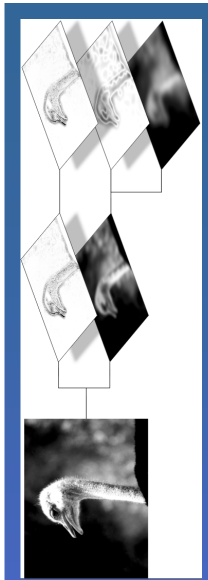

Non-standard Decomposition

In the following iteration steps, the standard decomposition only iterates the *purely* low-pass filtered approximations:

$$\begin{aligned}(\square) &= (V_2 \oplus W_2) \times (V_2 \oplus W_2) \oplus V_1 \times W_1 \\ &\quad \oplus W_1 \times V_1 \oplus W_1 \times W_1 \\ &= V_2 \times V_2 \oplus V_2 \times W_2 \oplus W_2 \times V_2 \oplus W_2 \times W_2 \\ &\quad \oplus V_1 \times W_1 \oplus W_1 \times V_1 \oplus W_1 \times W_1,\end{aligned}$$

4 iterations

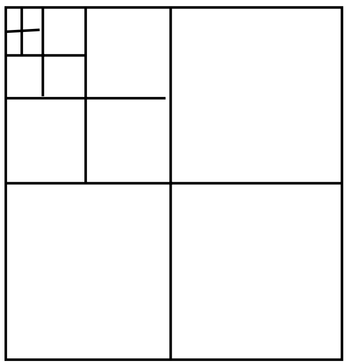

4 iterations

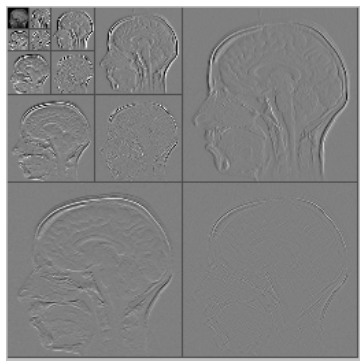

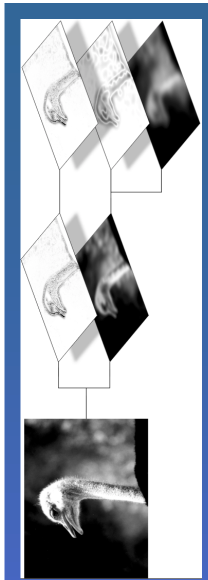

Comparison

- Standard decomposition:
  - more fine-grained: it realizes a better localization of a signal's energy in the approximation.
- Non-standard decomposition:
  - less complex,
  - mostly used in image coding applications.

Frequency location in Fourier domain

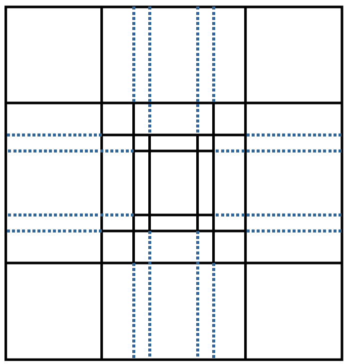

Non-standard, 3 iterations

Standard, 3 iterations

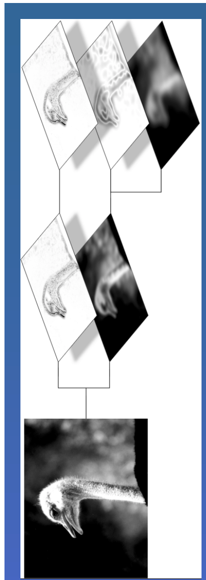

2.2 Signal Boundary

- A digital filter is applied to a signal by convolution. In order to result in a mathematically correct, reversible wavelet transform, *each* signal coefficient must enter into filter_length/2 calculations of convolution.
- Thus, each filter longer than Haar (i.e., 2 entries), requires a boundary extension.
- Boundary treatment more important the shorter the signal under consideration.
- Common policies:
  - circular convolution
  - padding

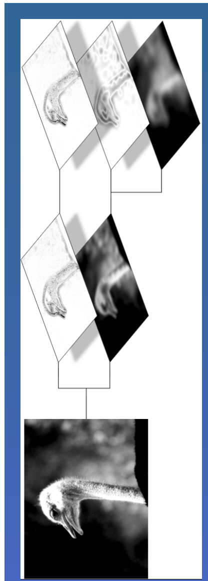

Circular Convolution

- Idea: ‘wrap’ the signal around

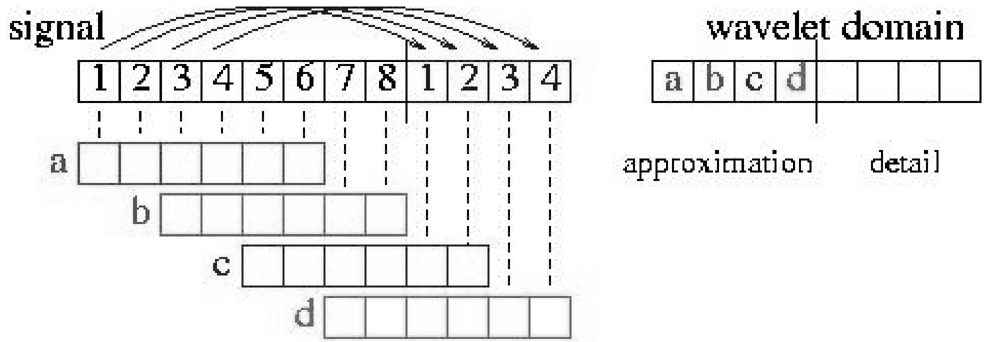

- Circular convolution is the only boundary policy that maintains the number of coefficients, thus simplifying storage handling.

- However, the time-information contained in the time-scale domain ‘blurs’.

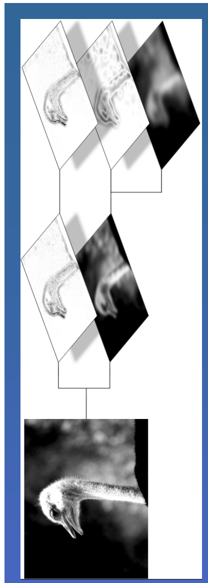

Padding Policies

- Idea: ‘pad’ the boundary with additional coefficients

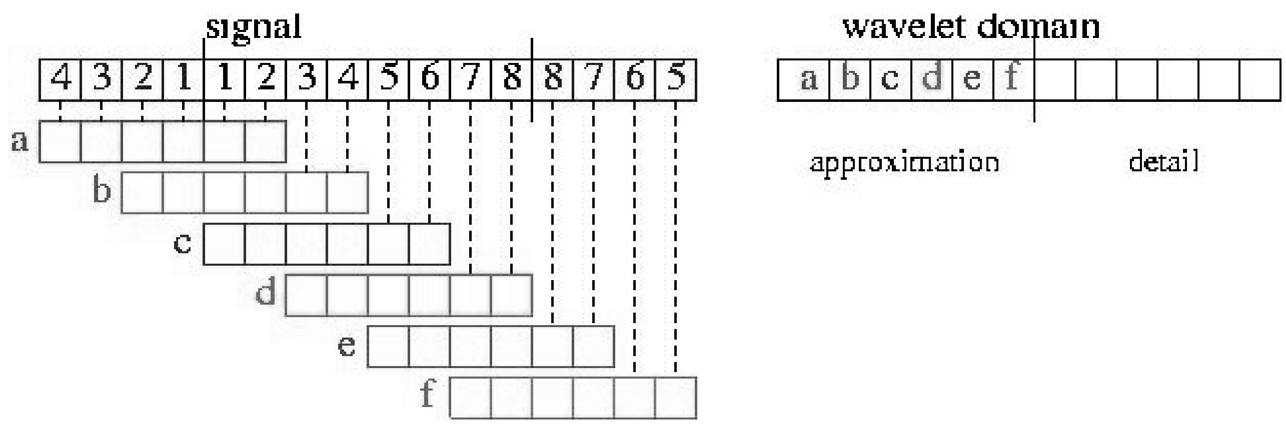

- Various padding policies: zero padding, constant padding, mirror padding, spline padding, ...
- Padding policies expand the transformed domain!
- Time-information of the time-scale domain is maintained.

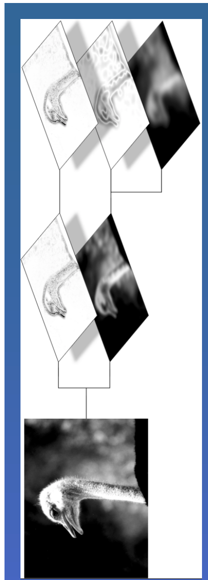

2.3 Painting the Time-scale Domain

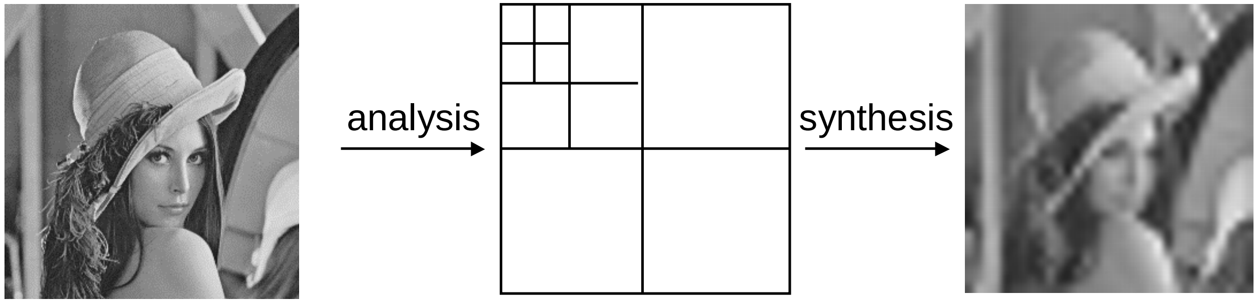

- So far, we have discussed wavelet *analysis*, i.e., the decomposition of a signal into its coefficients in the time-scale domain.
- Now: visualization of the time-scale domain.
- Wavelet-transformed coefficients are *not* pixel values. Consideration of
  - normalization and
  - range

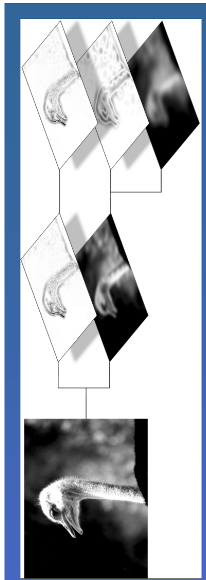

Normalization

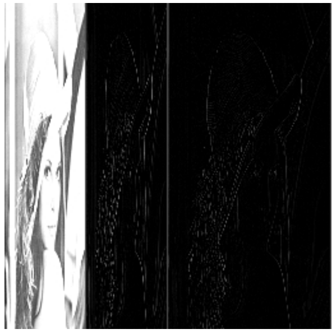

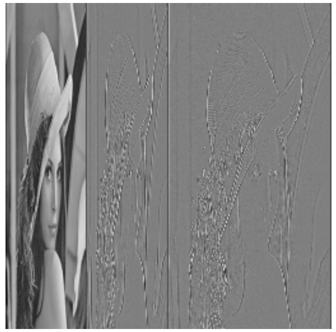

Two possible realizations of ‚painting the time-scale domain‘:

- No normalization:
  - details vary about zero, but this means black in image coding,
  - approximation is lifted by factor of $\sqrt{2} > 1$ (for Daubechies filters). Thus, the luminance is lifted by $\sqrt{2}$ in each iteration.
- Normalization:
  - lift the details by 128, i.e., by a medium gray color,
  - divide the approximations through $\sqrt{2}$ before painting.

Thus, all the images of the time-scale domain in this tutorial are ‘cheated’ since they are edited before visualization.

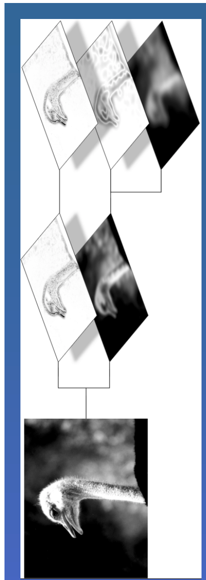

Growing Spatial Range with Padding

We have seen that boundary padding policies result in an enlarged time-scale domain.

Iteration level	Size of ‚upper left corner' Haar	Daub-20
1	128 x 128	147 x 147
2	64 x 64	93 x 93
3	31 x 32	66 x 66
...	...	...
8	1 x 1	39 x 39

All coefficients in the time-scale domain with zero padding

All coefficients in the time-scale domain with mirror padding

...and ‚what we would prefer for painting'.

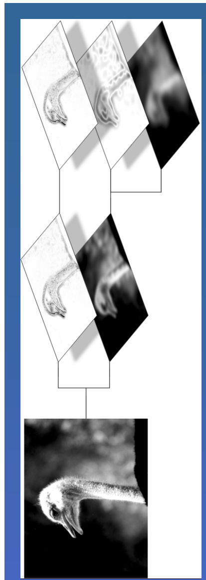

2.4 Lifting

- A different technique to construct biorthogonal wavelets and multiresolution has been introduced by Sweldens: *lifting scheme* or *second generation wavelets*.
- Advantages:
  - amount of floating point operations can reduced by a factor of 2,
  - allows fully *in-place-calculation*,
  - is *not* defined via the Fourier transform, thus is easier to understand.

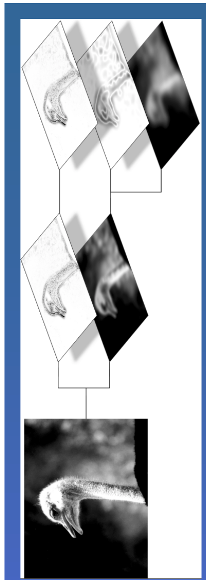

Example (I)

Look at the Haar transform from a different perspective.

- Signal a_0 with sampling distance 1 shall again be decorrelated.
- By subsampling the even samples of the original signal, one obtains a new sequence of *approximations*:

  $$a_{1,k} := a_{0,2k} \quad \text{for} \quad k \in \mathbb{Z}$$

- A trivial way to capture the lost information is to say the *detail* is simply contained in the odd samples:

  $$d_{1,k} = a_{0,2k+1}$$

- A more elaborate way is to recover the original samples from the subsampled coefficients $a_{1,k}$. Then, the odd samples indicate to what extend the signal 'fails to be linear':

  $$d_{1,k} := a_{0,2k+1} - \frac{1}{2}(a_{0,2k} + a_{0,2k+2})$$

  The expected value of these details is small.
- In order to preserve the average value of all coefficients at each level, i.e., $2\sum_k a_{j+1,k} = \sum_k a_{j,k}$ the approximations are *lifted* again:

  $$a_{1,k} = a_{0,2k} + \frac{1}{4}(d_{1,k-1} + d_{1,k})$$

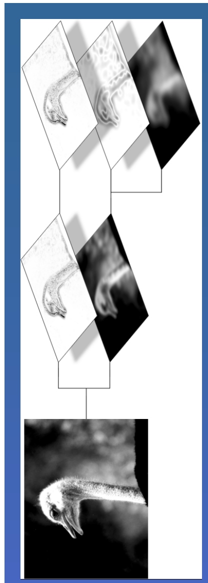

Example (II)

The wavelet transform on each level now consists of two stages:

$$d_{1,k} := a_{0,2k+1} - \frac{1}{2}(a_{0,2k} + a_{0,2k+2})$$

$$a_{1,k} = a_{0,2k} + \frac{1}{4}(d_{1,k-1} + d_{1,k})$$

This is demonstrated in the following scheme:

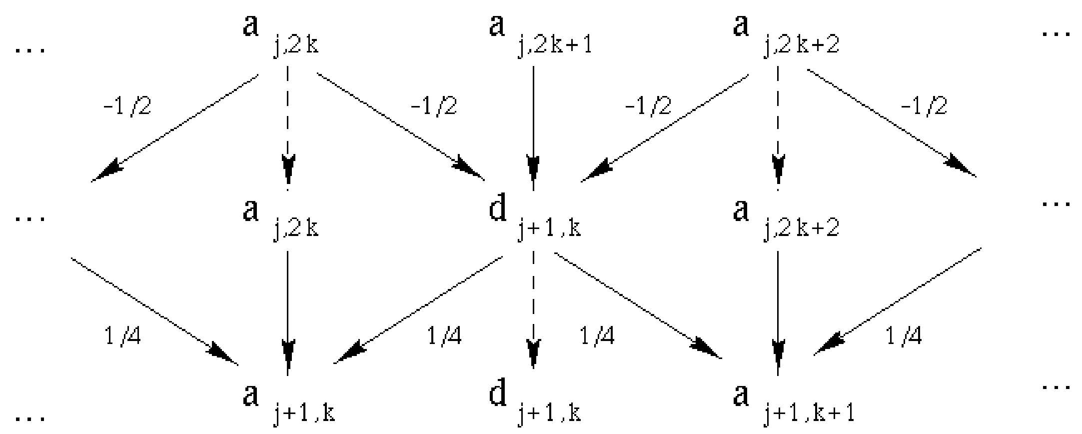

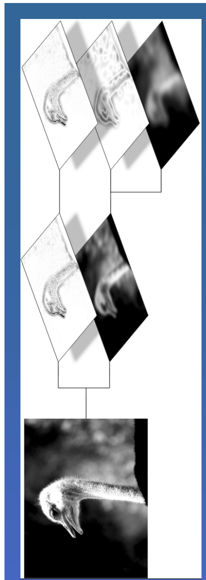

Example (III)

The synthesis is simply the reverse of the two equations above.

The filters induces by this lifting example are:

- high-pass filter (details): $[-\frac{1}{2}\ 1\ -\frac{1}{2}]$
- low-pass filter (approx.): $[-\frac{1}{8}\ \frac{2}{8}\ \frac{6}{8}\ \frac{2}{8}\ -\frac{1}{8}]$

This is the default *reversible wavelet transform* Daub-5/3 suggested in JPEG2000. An *irreversible wavelet transform* is defined as well, denoted Daub-9/7.

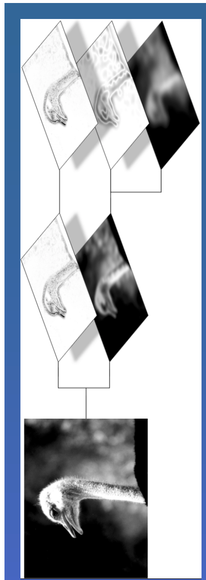

Example (IV)

Filter coefficients of the two default wavelet filter banks of JPEG2000.

Daub–9/7 Analysis and Synthesis Filter Coefficients

i	Analysis Filter low–pass	Analysis Filter high–pass	Synthesis Filter low–pass	Synthesis Filter high–pass
0	0.6029490182363579	1.115087052456994	1.115087052456994	0.6029490182363579
±1	0.2668641184428723	-0.5912717631142470	0.5912717631142470	-0.2668641184428723
±2	-0.07822326652898785	-0.05754352622849957	-0.05754352622849957	-0.07822326652898785
±3	-0.01686411844287495	0.09127176311424948	-0.09127176311424948	0.01686411844287495
±4	0.02674875741080976			0.02674875741080976

Daub–5/3 Analysis and Synthesis Filter Coefficients

i	Analysis Filter low–pass	Analysis Filter high–pass	Synthesis Filter low–pass	Synthesis Filter high–pass
0	6/8	1	1	6/8
±1	2/8	-1/2	1/2	-2/8
±2	-1/8			-1/8

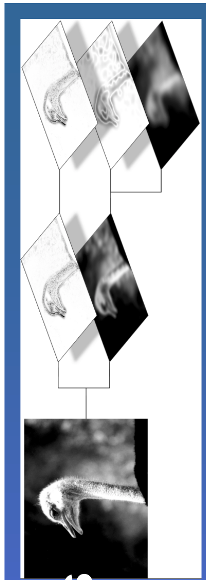

3.1 JPEG2000

- The JPEG2000 has been released on January 2, 2001
- Based on the wavelet transform

Part	Content
1	JPEG2000 Still Image Coding
2	Extensions
3	Motion-JPEG2000
4	Conformance
5	Reference Software
6	Compound Image File Format

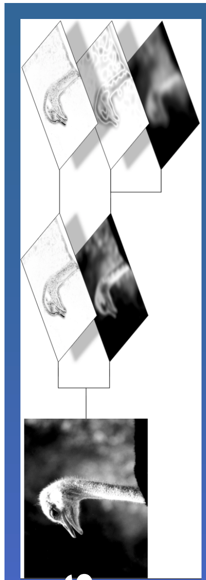

Design Goals

- Better performance at lower bitrates.
- Lossy and lossless compression.
- Progressive data transmission.
- Definition and coding of regions-of-interest.
- Random access.
- Robustness towards bit errors.
- Open architecture.
- Possibility of content description.
- Transparency.
- Watermarking.
- Support of images of arbitrary components.

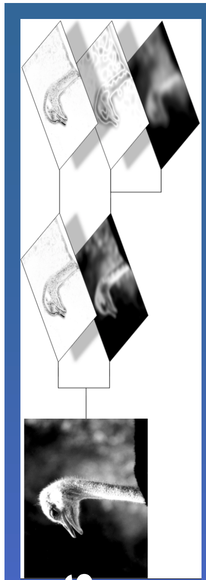

Architecture

- Image is decomposed into color components which are processed separately.
- Each color component is subject to a *tiling* process.
- Each tile is subject to the wavelet transform
  - standard reversible filter: Daub-5/3 (see Section 2.5)
  - standard irreversible filter: Daub-9/7 (see Section 2.5)
- The different scales are ordered such that they describe specific regions of the image. The resulting blocks are called *subbands*.
- Subbands are quantized and stored in *code blocks*.
- The bit layers of the code blocks are entropy encoded.
- Specific treatment of *regions-of-interest*.
- A file format allows the storage of the data stream.

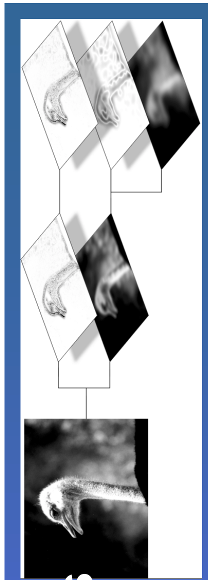

Performance

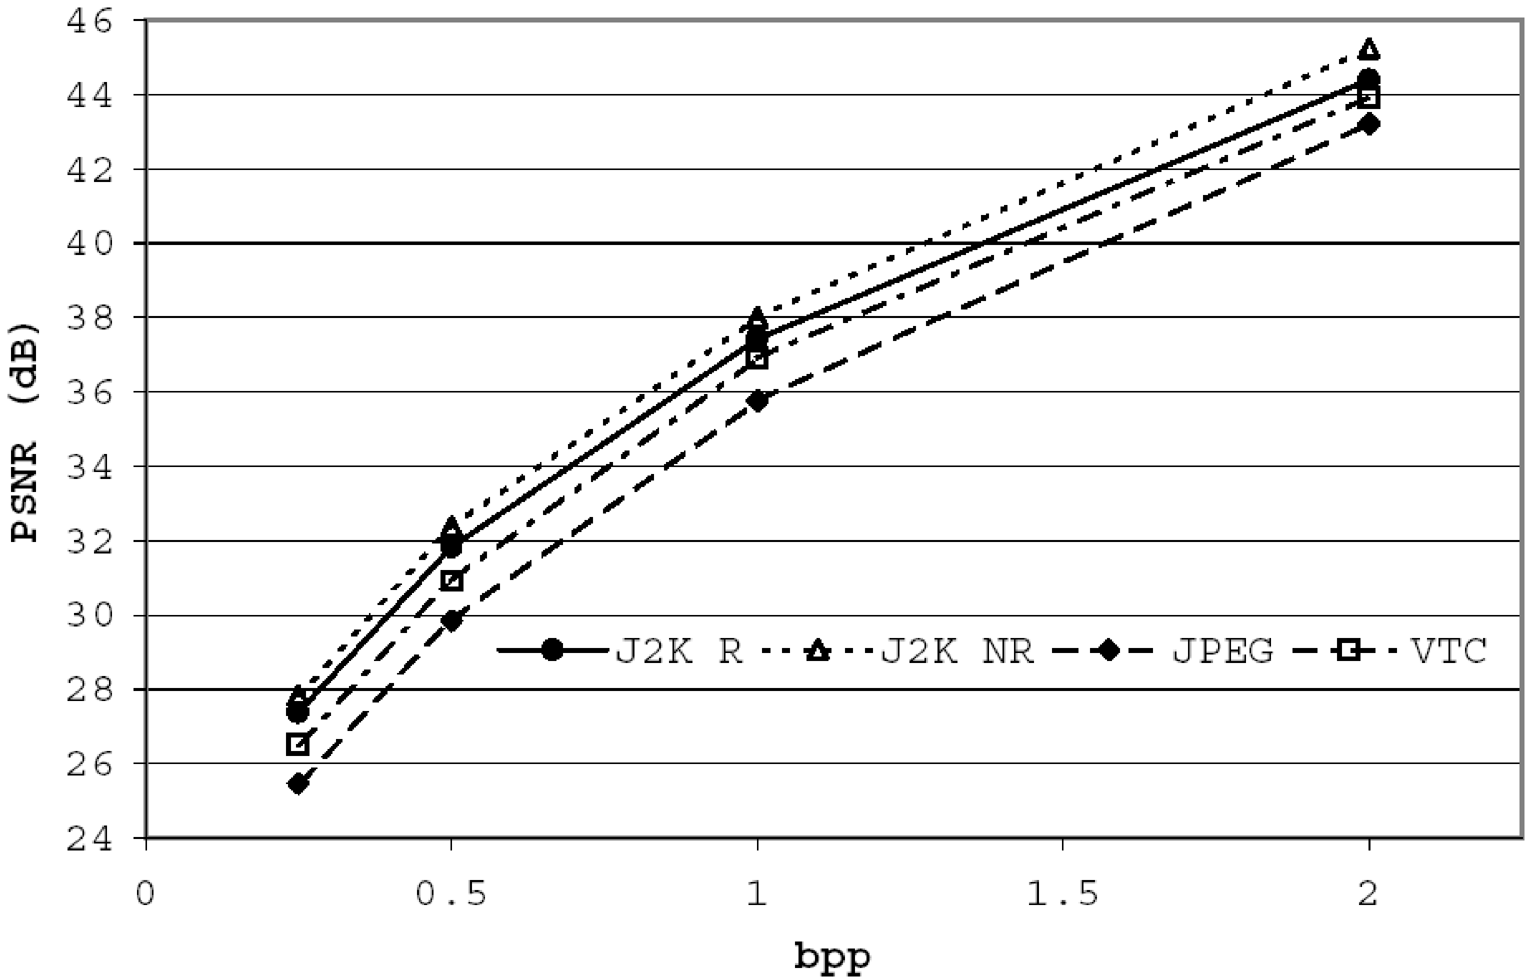

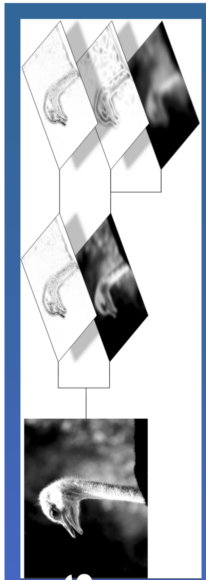

Coding Detail (I)

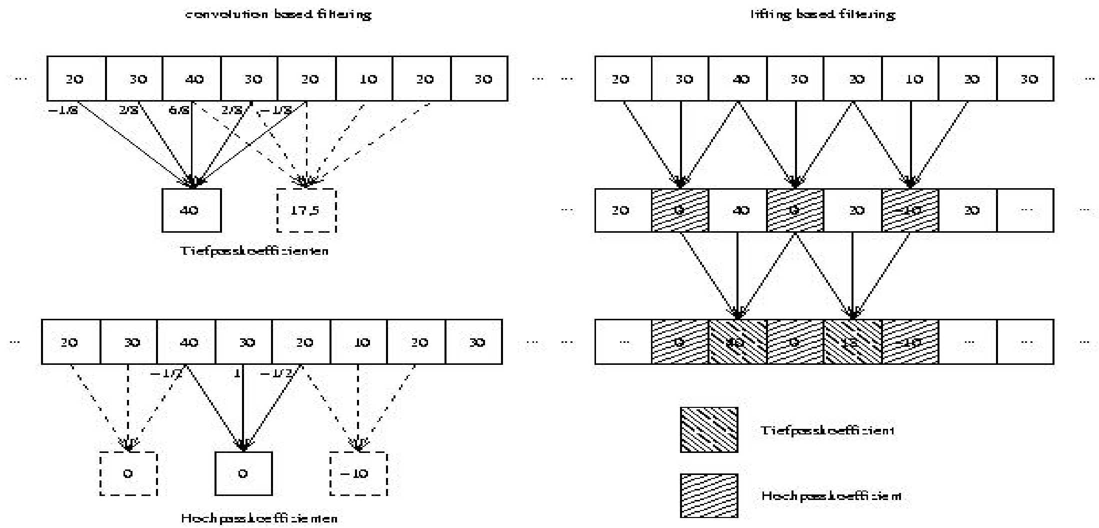

Standard wavelet transform with interleaved storage.

Coding Detail (II)

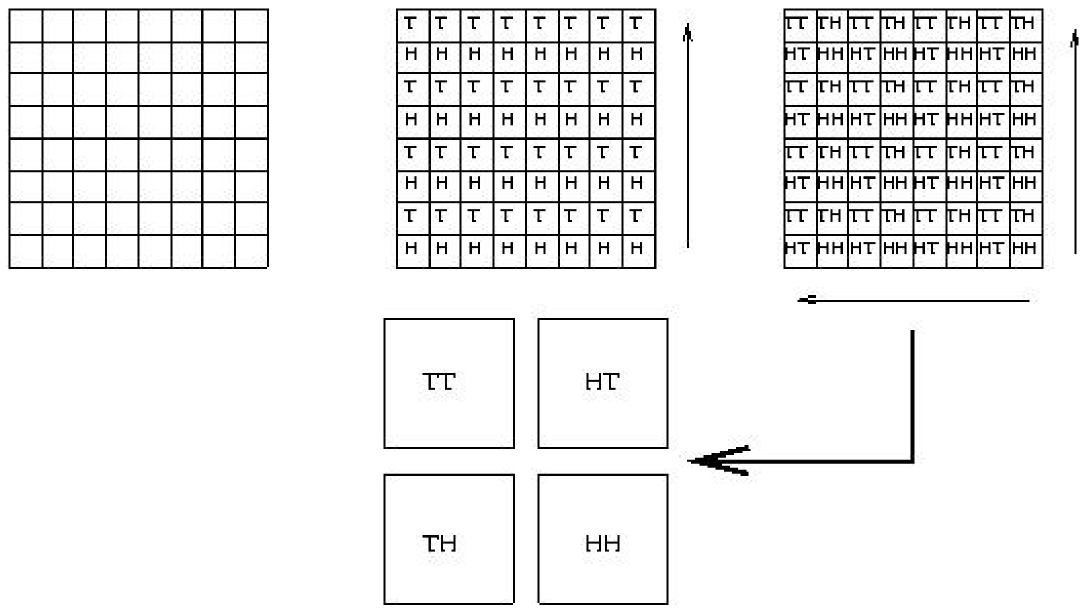

Interleaved storage in two dimensions.

Regions-of-interest in JPEG2000

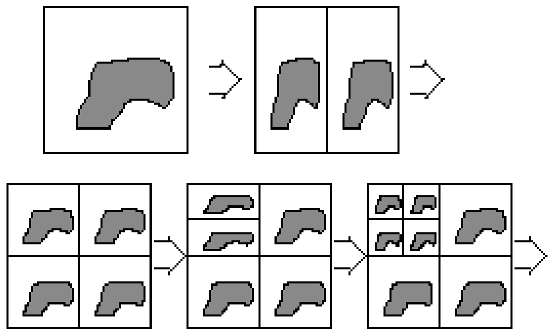

Due to the time (or: location) information that subsists in the time-scale domain, it is possible to track specific *regions-of-interest* (ROI) in their encoded representation.

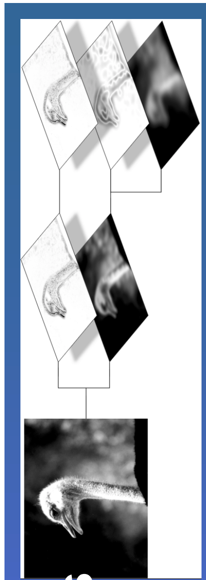

What is of interest?

- The investigation of ROI requires a pragmatic approach of the term ‘interest’.
  - regions of higher coding quality (RHQ),
  - regions of minor coding quality (RMQ).
- Classifications for segmentation
  - according to information content:

  - according to visual perception

Shape of ROI segments

- Shapes might be arbitrary like in the previous examples, or pre-defined:

- Trade-off between coding complexity and utility.

MAXSHIFT-method

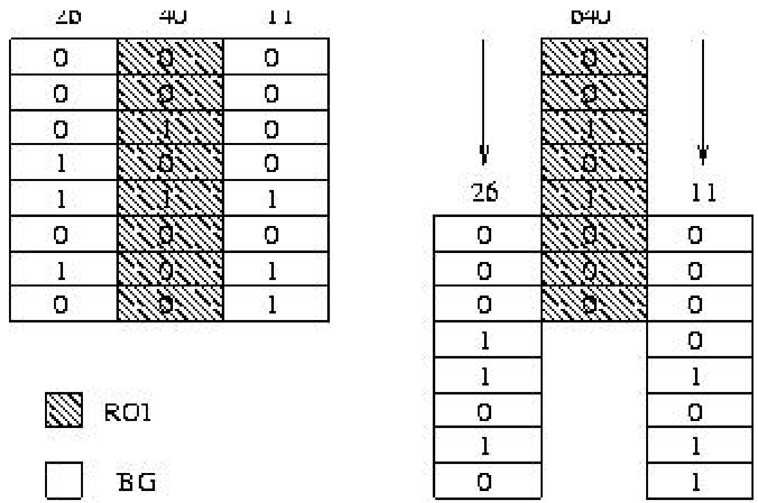

JPEG2000: MAXSHIFT-method defines the (arbitrary) shape of a ROI.

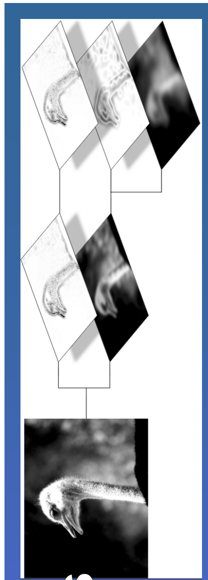

3.2 Hierarchical Video Coding

- A major drawback for the rapid deployment of streaming video in the internet is its heterogeneity.
- Solutions:
  - redundant coding or
  - hierarchical coding

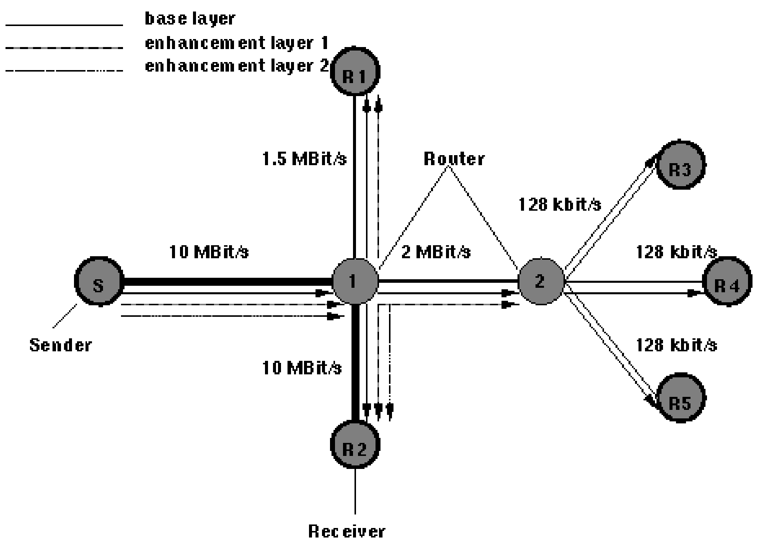

Layered data transmission in a heterogeneous network. The sender sends the base layer plus all enhancement layers. Each receiver decides how many layers he/she can receive.

Heuristic for Comparison

Pyramid encoding

Layered DCT frequencies

Bit layering

Layered wavelet transform frequencies

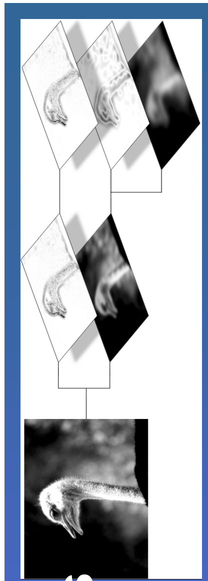

Layering Policies

Reasonable layering of wavelet-transformed data into base layer and enhancement layers can be carried out according to three policies:

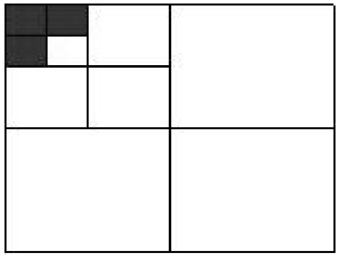

Policy 1: Blockwise. Layering and its respective synthesis work the other way round than analysis.

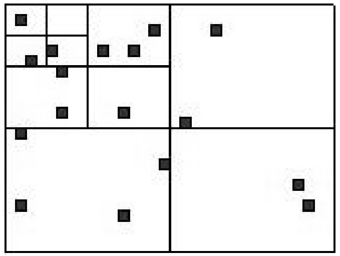

Policy 2: max Coefficients. The base layer should look for those coeffs with the highest (absolute) values, i.e., above a certain threshold. Subsequently smaller thresholds define the following layers.

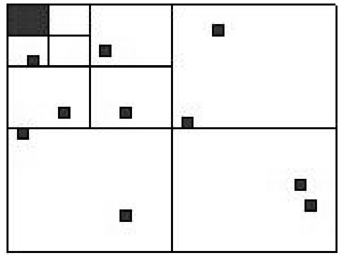

Policy 3: Mixture. First transmit the approximation and then subsequently fill the layers according to policy 2.

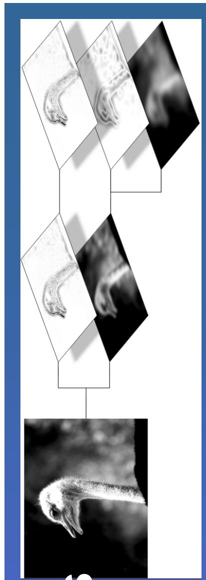

Different Perceptions

original frame

Synthesis with 6.25% of the information in the time-scale domain:

Blockwise layering.

Max. coefficients.

Mixture.

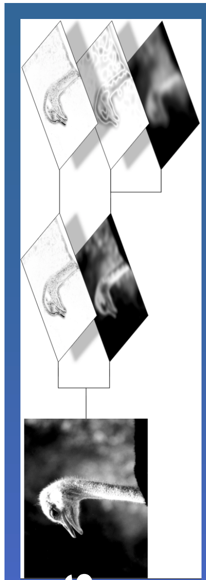

Open Issues

- Reliable estimation of the bitrate for different coding techniques, including
  - Huffman encoding,
  - entropy encoding,
- Thorough comparison based on the bitrate.
- Consideration of the network: adaptive coding according to the actual traffic.

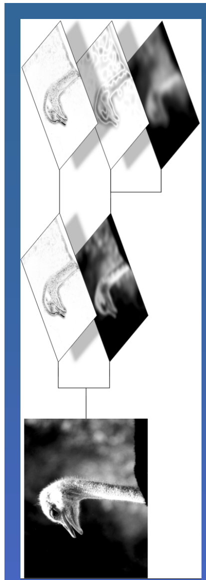

Conclusion

This tutorial on wavelets in theory and applications was subdivided into three major parts:

- Part I:
  Overview on the mathematical background of multiscale analysis and the wavelet transform.
- Part II:
  Discussion of implementation issues.
- Part III:
  Examples of wavelet applications in multimedia.

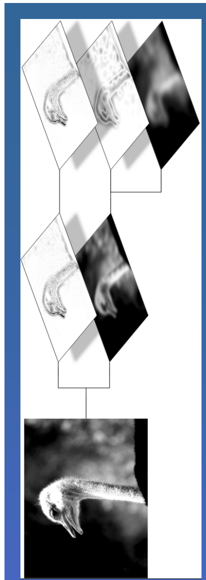

Information...

Claudia Schremmer
University of Mannheim, Germany
Department Praktische Informatik IV
http://www.informatik.uni-mannheim.de/informatik/pi4
schremmer@informatik.uni-mannheim.de