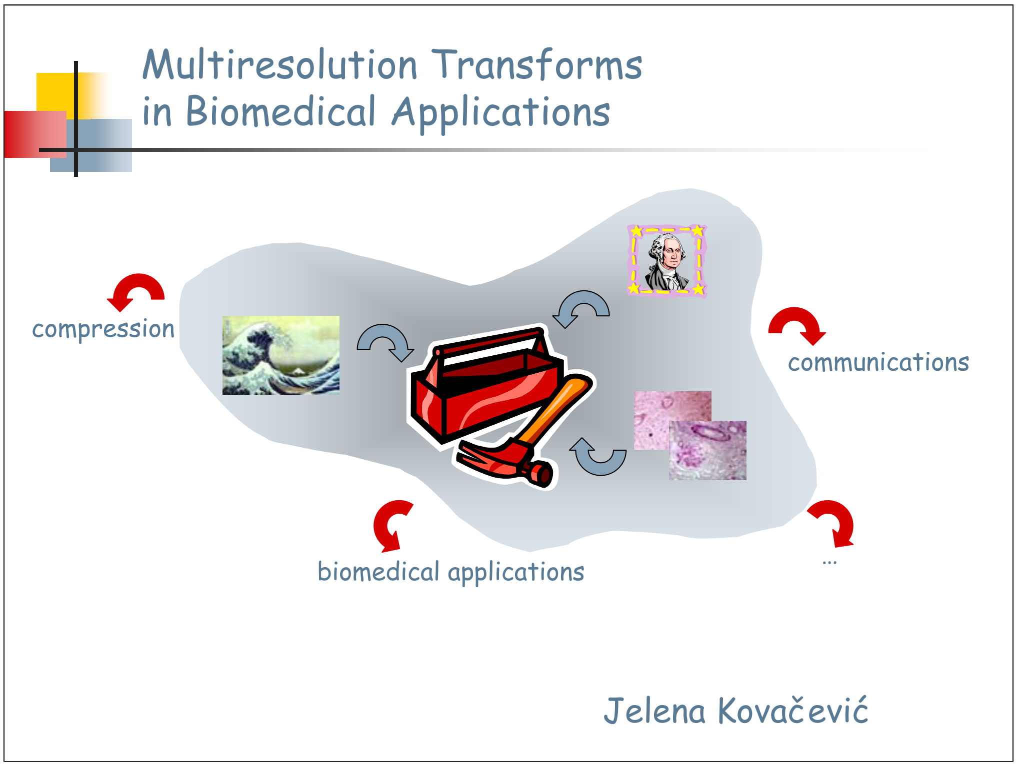

Almost all the images and signals were obtained in Matlab, using the Wavelet Toolbox. When this is not the case, I will specify it.

In instructions on how to obtain them, first do

wavemenu

in Matlab. This brings up the mail Wavelet Toolbox GUI. I used four different functions:

Wavelet 1-D	(W1D)
Wavelet Packet 1-D	(WP1D)
Wavelet 2-D	(W2D)
Wavelet Packet 2-D	(WP2D)

I will refer to them as specified in parentheses.

MR: Why?

- Localization
- Adaptivity
- Computational efficiency
- Wide use in medical imaging
- Subsumes many well known transforms

Comments on the slide

Some of the most important and useful characteristics of multiresolution are listed above. We now look at each individually.

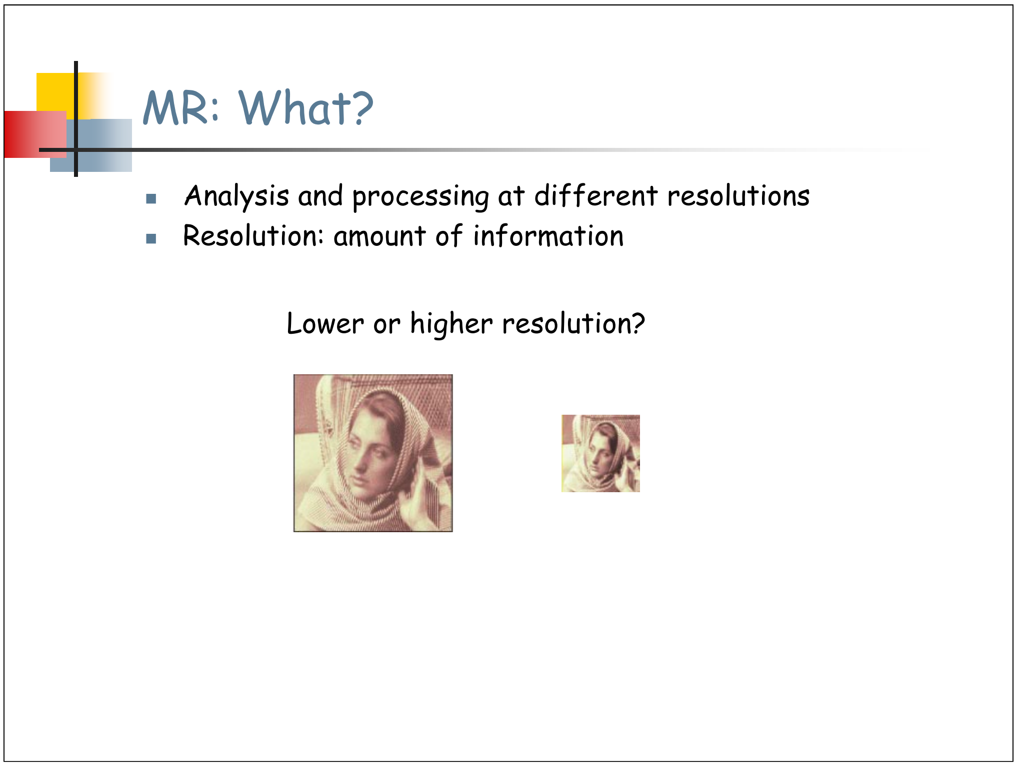

•**Image on the left:** image_original.jpg

How obtained: W2D

File → Example Analysis → At level 2, with Haar → Woman

Choose the original image

•**Image on the right:** image_lr_1.jpg

How obtained: W2D

File → Example Analysis → At level 2, with Haar → Woman

Choose 1 level

Choose the UL (upper left) image in the Image Selection box

Comments on the slide

This slide illustrates the concept of resolution as the amount of information in a signal. If one looks at the two images and they look as they were “equally sharp”, the one on the left is of higher resolution just by the virtue of the fact that it is larger. In fact, the image on the right was obtained as the LL subband in 1-level DWT.

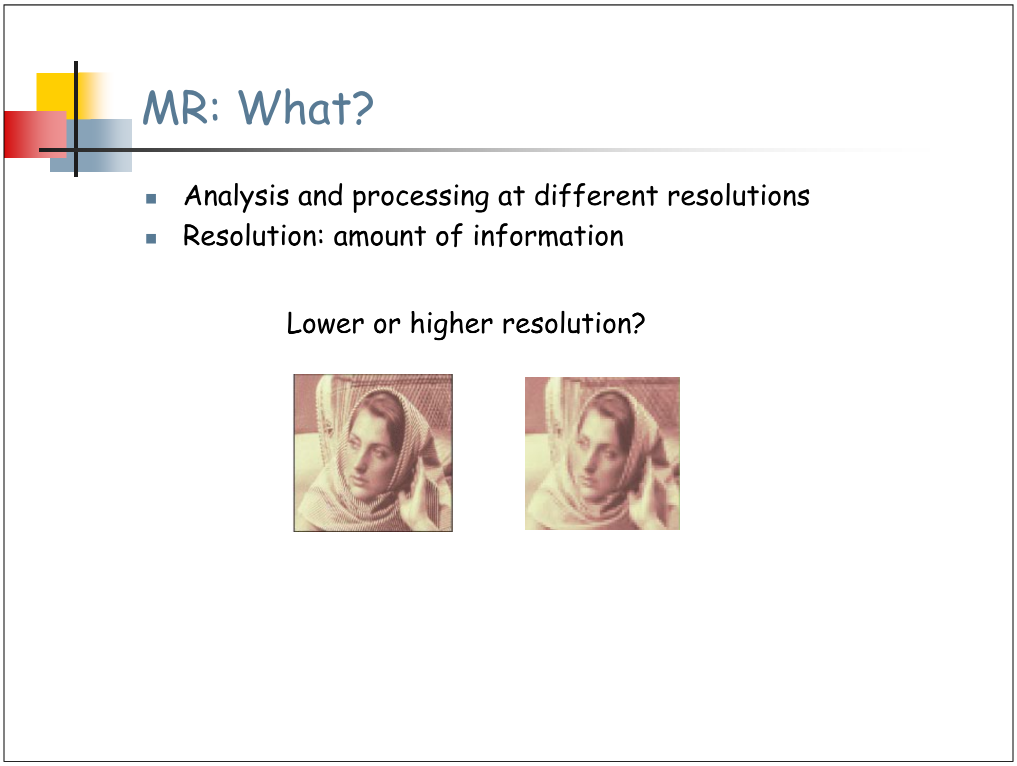

•Image on the left: image_original.jpg

How obtained: W2D

File → Example Analysis → At level 2, with Haar → Woman

Choose the original image

•Image on the right: image_lr_reconstructed.jpg

How obtained: W2D

File → Example Analysis → At level 2, with Haar → Woman

Select the UL (upper left) image in the Image Selection box

Click on Reconstruct in Operations on Selected Image (right panel)

Choose the image Recons. Approx. coef. of level 2

Comments on the slide

This slide continues the story from the previous one. Here, the one can say that the left is of higher resolution than the right since it is clear that the right one is missing some information (sharpness of detail). This is true since the right image is obtained by interpolating the right image from the previous slide.

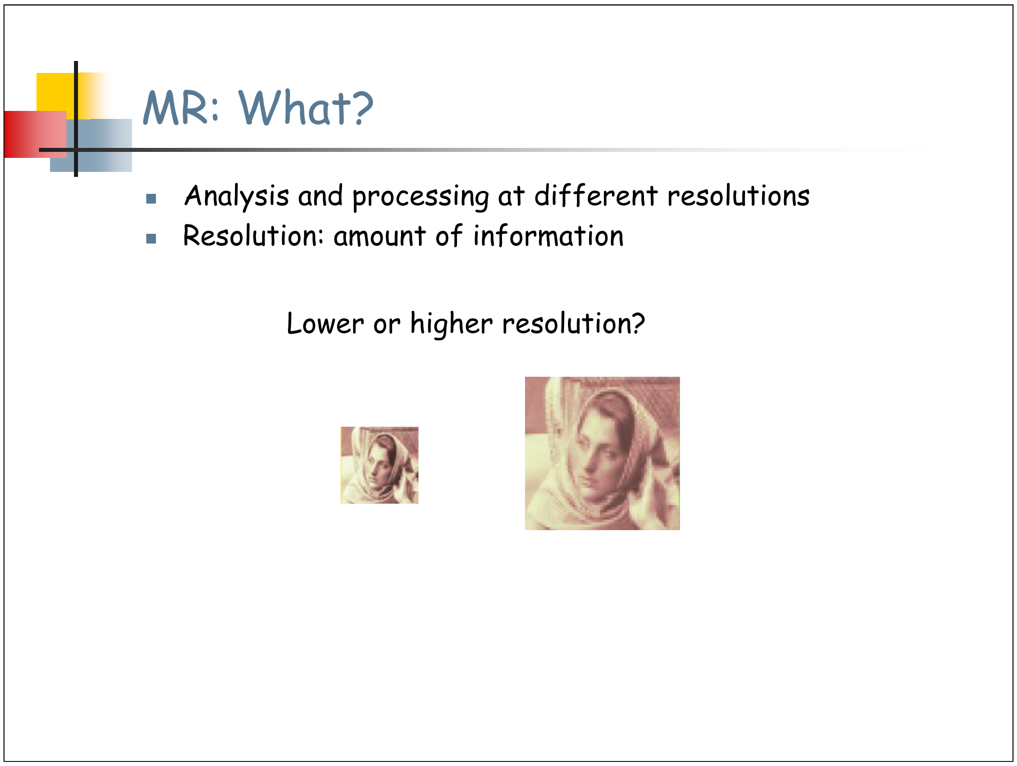

•**Image on the left:** image_lr_1.jpg

How obtained: W2D

File → Example Analysis → At level 2, with Haar → Woman

Choose 1 level

Choose the UL (upper left) image in the Image Selection box

•**Image on the right:** image_lr_reconstructed.jpg

How obtained: W2D

File → Example Analysis → At level 2, with Haar → Woman

Select the UL (upper left) image in the Image Selection box

Click on Reconstruct in Operations on Selected Image (right panel)

Choose the image Recons. Approx. coef. of level 2

Comments on the slide

Same story as before. Now it is not as clear what one should pick and just by looking at them it is hard to be sure which one is of higher resolution and which one is lower. In this particular case, they are both of the same resolution as the right one is the interpolated version of the left one, and thus, no new information has been added.

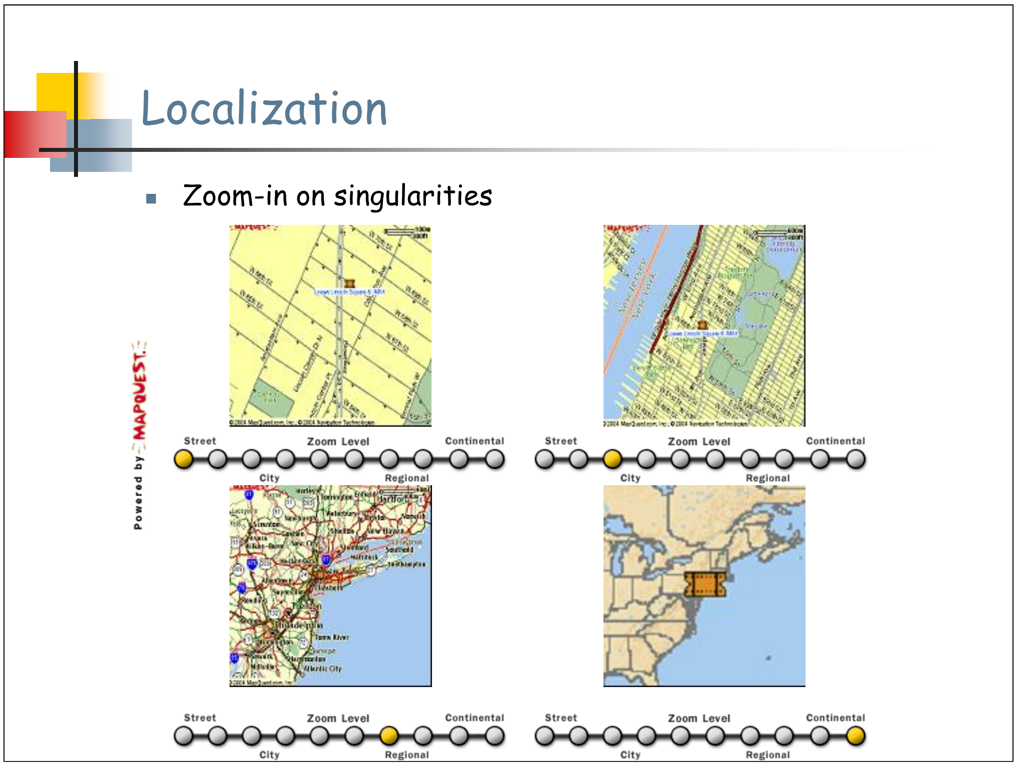

MapQuest: Search for Loews Theatre on 68th & Broadway, NYC

Comments on the slide

Wavelets have the ability to zoom in on singularities. An example is MapQuest where if you search for something, you can look on which continent it is located, followed by a region, followed by a city and finally, followed by a street view. You refine your search by using a zooming tool.

Localization

- Zoom-in on singularities

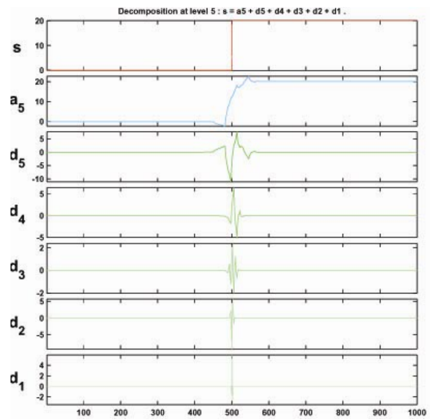

•**Graph:** decomposition_step_db2_5.jpg

How obtained: W1D

File → Example Analysis → Basic Signals → w db2 at level 5 → Step signal

Comments on the slide

Here is an example of a signal—step signal, and how the WT zooms in on the singularity—step. You can see that at the finest level, the detail coefficients d1 are all 0 except in the vicinity of the step. This analysis is performed with the db2 filters which are of length 4. As the level gets higher (coarser), the support of the filter operating on the signal gets longer and longer, and thus more and more coefficients in the vicinity of the step catch it, making the result fatter.

Localization

- Zoom-in on singularities

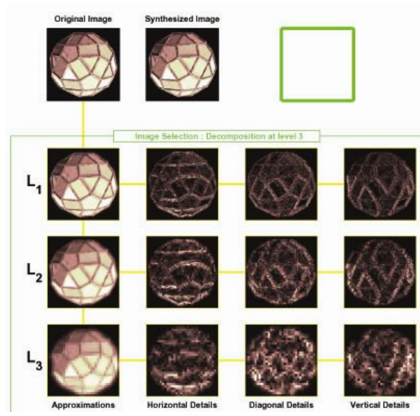

•**Image:** decomposition_ball_bi5.5_3.jpg

How obtained: W1D

File → Example Analysis → At level 3, with bior5.5 → Facets

Choose View Mode: Tree

Comments on the slide

This slide shows the same as the previous except on images. Again, the singularities (edges) are best seen at the finest (lowest) level while they become blurred as the level becomes higher.

Localization

- Characterize singularities

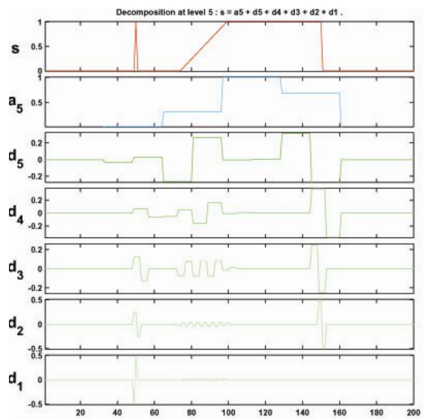

•**Graph:** singularities.jpg

How obtained: W1D

File → Load → Signal → singularities.mat

Choose Wavelet: Haar, Level: 5

Click Analyze

Comments on the slide

This slides demonstrates the ability of the wavelet transform not only to detect singularities but to characterize them as well. The graph has three different types of singularities and they all behave differently across scales. By measuring the amplitude and behavior of the singularity across scales, one can determine its location and type.

Denoising

- Separate signal from noise

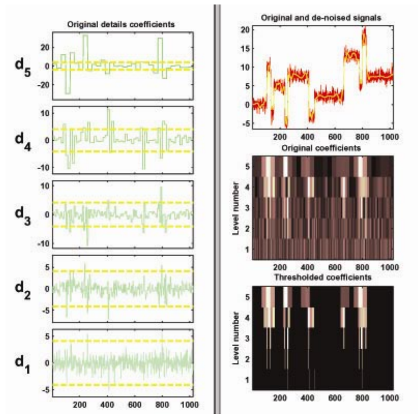

•**Graph:** decomposition_blocks_denoise.jpg

How obtained: W1D

File → Example Analysis

→ Noisy Signals – Constant Noise Variance

→ with sym8 at level 5 → Noisy blocks

Click De-noise

Click De-noise again in the new window that pops up

Comments on the slide

The power of multiresolution shows itself also in its ability to separate signal from noise. If something about the noise is known apriori, one can look at the behavior of coefficients across scales and decide which belong to noise and which to the signal. Those deemed to belong to noise are simply removed and the signal is reconstructed without them. The yellow lines on the left side of the graph denote the threshold below which the coefficients are eliminated. The original signal is given in red on the right side and the denoised one in yellow.

Adaptivity: Holy Grail of Signal Analysis/Processing

- Understand the "blob"-like structure of the energy distribution in the time-frequency space
- Design a representation reflecting that

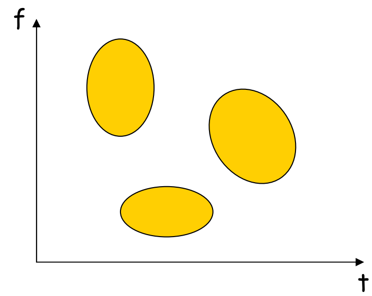

Comments on the slide

The adaptivity is a big advantage of MR techniques. The Fourier transform does analysis with a fixed window both in time and frequency and thus localizing features in either time or frequency is possible only within a fixed window.

Adaptivity

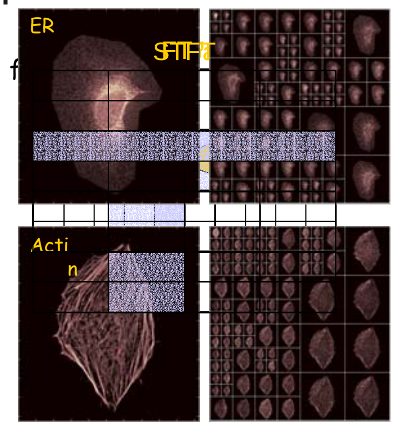

- "Holy Grail" of Signal Analysis/Processing
  - Understand the "blob"-like structure of the energy distribution in the time-frequency space
  - Design a representation reflecting that

The adaptivity is a big advantage of MR techniques. The Fourier transform does analysis with a fixed window both in time and frequency and thus localizing features in either time or frequency is possible only within a fixed window.

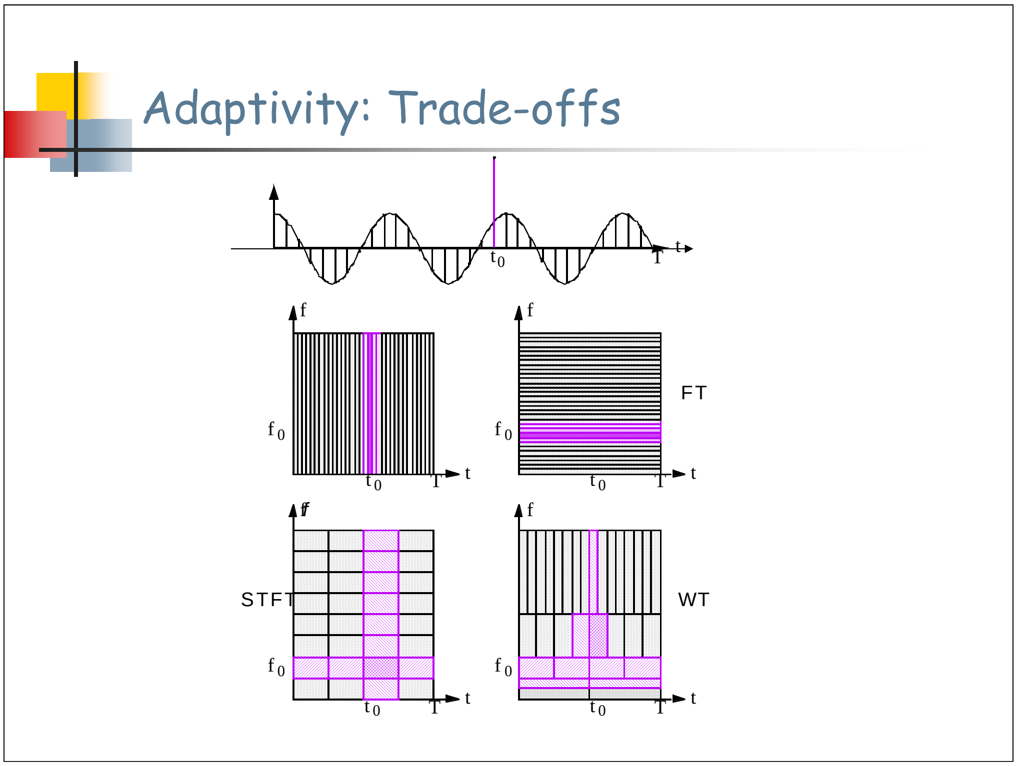

Comments on the slide

To illustrate that concept, consider a sampled sinusoid of fairly low frequency with an impulse superimposed on it (top graph). The upper left graph shows what would happen if we analyzed this signal with a basis consisting of shifted impulses---each impulse would catch what happens at exactly that point in time. This analysis is perfect for isolating time discontinuities but does nothing for isolating the frequency of the sinusoid---it is spread over all frequencies. At the other end of the spectrum is the Fourier transform (upper right graph) where, obviously, the sinusoid is isolated perfectly while the time discontinuity cannot be resolved as it has been spread over all time. The bottom left tries to alleviate that problem by windowing the sinusoids in the Fourier transform (short-time Fourier transform). Indeed, some time localization has been bought at the expense of frequency localization. The problem is that the once fixed, the window determines how finely we can isolate locally fast events both in time and frequency. Finally, the bottom right shows the wavelet transform and the trade-off it offers; at high frequencies, the time discontinuity is caught while at low frequencies the frequency discontinuity shows up (the sinusoid). Note a catch here: if the sinusoid had been of high frequency, then it would not have been isolated well. This is solved by using arbitrary trees, those which adapt themselves to the signal at hand.

Adaptivity: Trade-offs

high frequency lower frequency

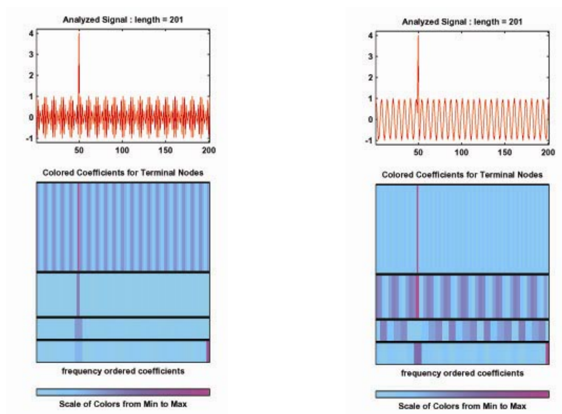

•**Graph on the left:** sin_0.9pi_dirac.jpg

How obtained: WP1D

File → Load → Signal → sin_0.9pi_dirac.mat

Choose Wavelet: Haar, Level: 3

Click Analyze

Click Wavelet Tree

•**Graph on the right:** sin_0.3pi_dirac.jpg

How obtained: WP1D

File → Load → Signal → sin_0.3pi_dirac.mat

Choose Wavelet: Haar, Level: 3

Click Analyze

Click Wavelet Tree

Comments on the slide

Here we have a similar example where on the left side we have a high frequency with an impulse and on the right a lower frequency and the impulse. Below are the time-frequency tilings obtained by using the DWT. It is obvious that on the right the frequency is identified to fall somewhere in the upper half of the spectrum while on the right it falls within the middle quarter or so.

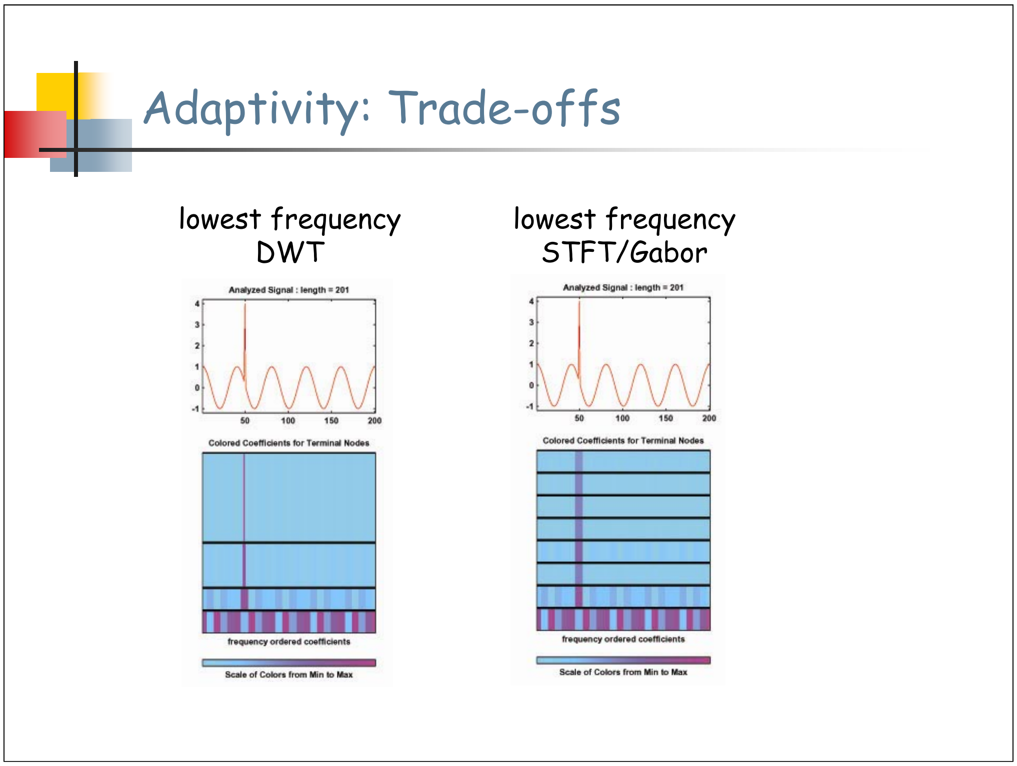

•**Graph on the left:** sin_0.05pi_dirac.jpg

How obtained: WP1D

File → Load → Signal → sin_0.05pi_dirac.mat

Choose Wavelet: Haar, Level: 3

Click Analyze

Click Wavelet Tree

•**Graph on the right:** sin_0.05pi_dirac_stft.jpg

How obtained: WP1D

File → Load → Signal → sin_0.05pi_dirac.mat

Choose Wavelet: Haar, Level: 3

Click Analyze

Comments on the slide

Here we have a very low frequency and the impulse. On the left is analysis with the DWT and on the right with the Fourier. Both analyses isolate the frequency equally well; however, the DWT is more precise in isolating the impulse. Note how the precision with which the FT isolates the pulse is not as fine (the impulse if fatter) as that of the DWT at high frequencies.

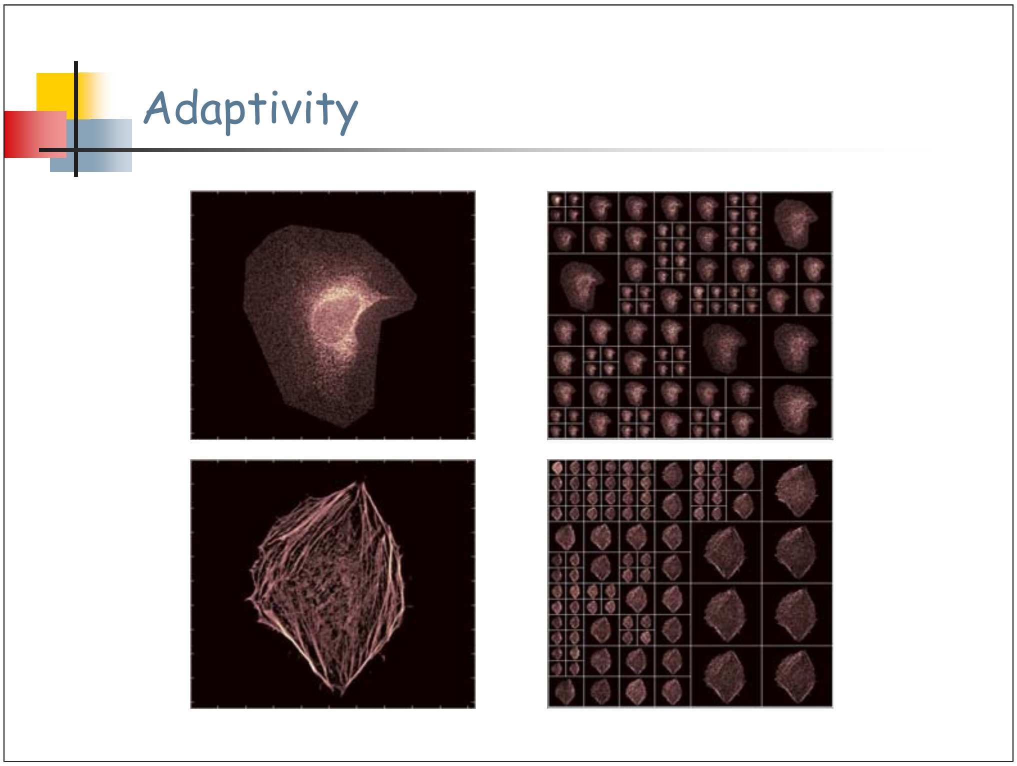

These were produced by Tad Merryman.

Comments on the slide

This slide illustrates the power of wavelet packets on an image. These images depict two different proteins in the cell (both are fluorescence microscopy images) and on the right is the best wavelet packet tree found by using the energy in the subband as the criterion.

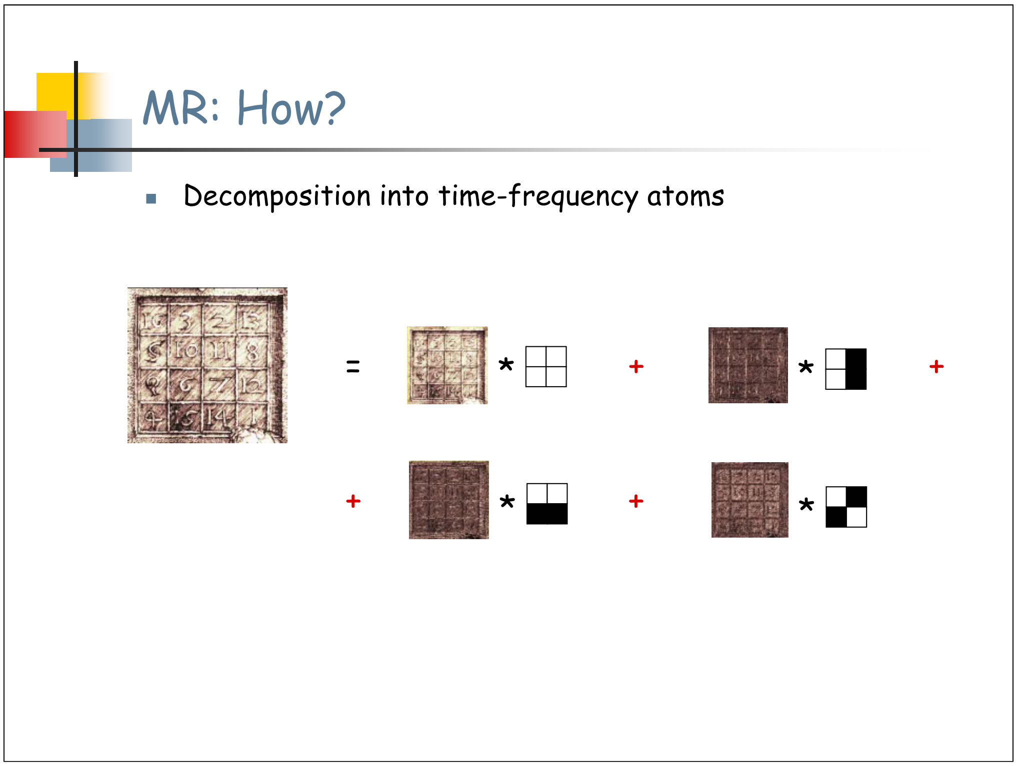

•**Image on the left:** detail_original.jpg

How obtained: W2D

File → Example Analysis → At level 3, with sym4 → Detail Durer

Choose Wavelet: Haar, Level: 1

Choose Original Image

•**Small images on the right:** detail_LL.jpg, detail_LH.jpg, detail_HL.jpg, detail_HH.jpg

How obtained: W2D

File → Example Analysis → At level 3, with sym4 → Detail Durer

Choose Wavelet: Haar, Level: 1

Choose the four images in the Image Selection box (UL=LL, UR=LH, LL=H, LR=HH)

Comments on the slide

This slide demonstrates how we do decomposition into time-frequency atoms. In this case, these are Haar, applied separately in both directions (thus you obtain the little 2x2 squares where black denotes a -1 and white denotes a 1). The image is obtained as a linear combination of these time-frequency atoms, where the coefficients in the expansion are the four images given above.

MR: How?

- Decomposition into time-frequency atoms

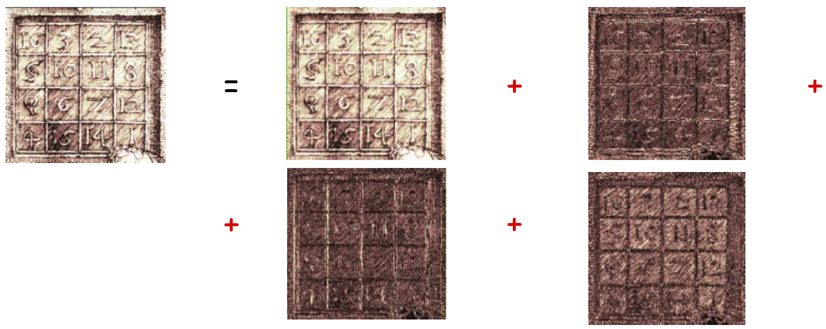

•**Image on the left:** detail_original.jpg

How obtained: W2D

File → Example Analysis → At level 3, with sym4 → Detail Durer

Choose Wavelet: Haar, Level: 1

Choose Original Image

•**Images on the right:** detail_LL_rec.jpg, detail_LH_rec.jpg, detail_HL_rec.jpg, detail_HH_rec.jpg

How obtained: W2D

File → Example Analysis → At level 3, with sym4 → Detail Durer

Choose Wavelet: Haar, Level: 1

Select the UL/UR/LL/LR image in the Image Selection box

Click on Reconstruct in Operations on Selected Image

Choose the image Recons. Approx. coef./detail of level 1

Comments on the slide

When the products are computed, one can look at the decomposition also as the sum of 4 subimages. The first one keeps high frequencies (lowpassed both in horizontal and vertical directions, LL), the second one has been highpassed horizontally and lowpassed vertically (HL), the third one is the reverse LH and the fourth one has been highpassed in both directions HH.

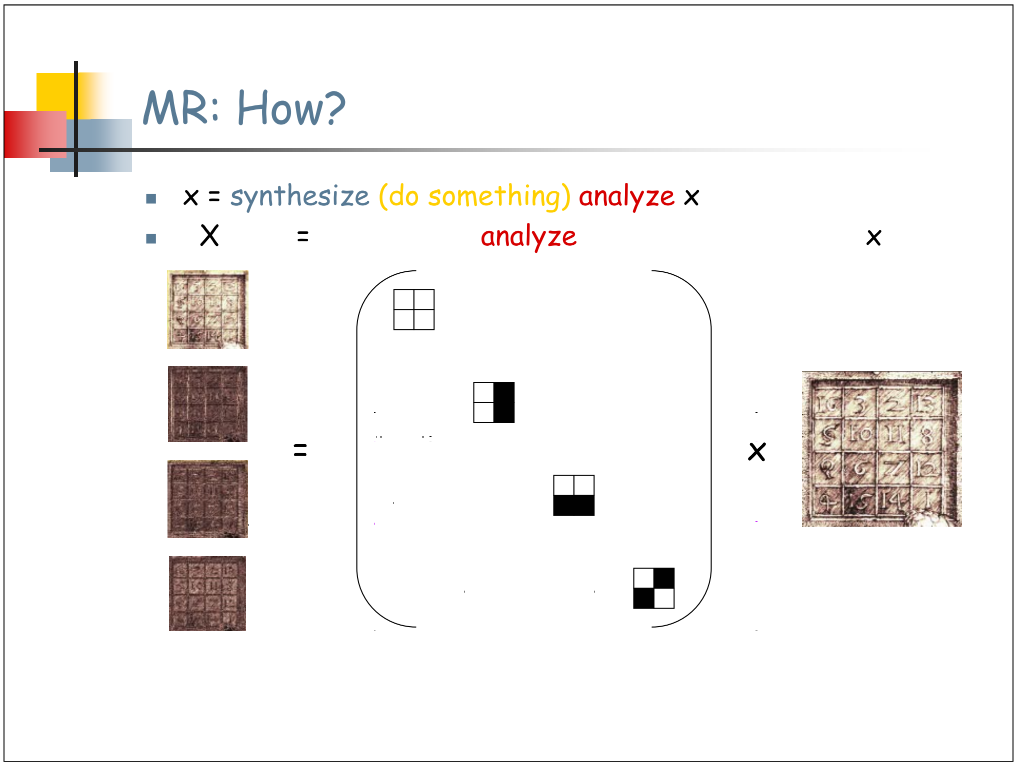

Same images as on Slide 16.

Comments on the slide

Mathematically, the analysis process is given as a vector/matrix operation, where we have a matrix operating on the input signal to produce the four middle signals (called subbands, channels, etc.)

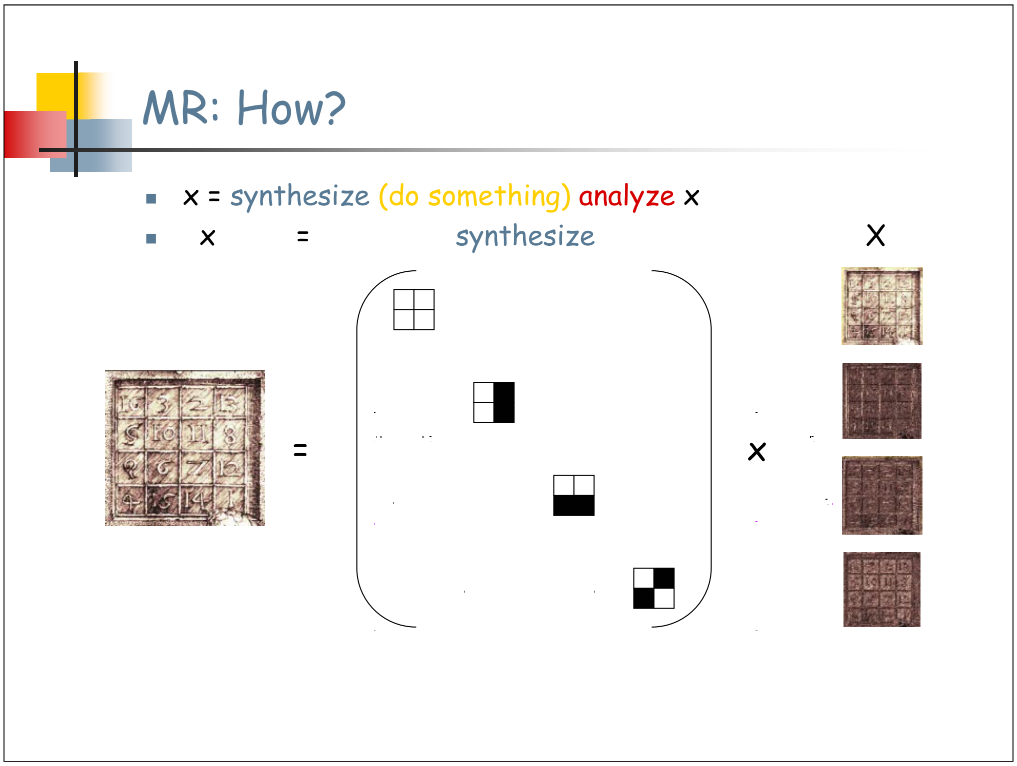

Same images as on Slide 16.

Comments on the slide

On the synthesis side, we do the opposite. In the Haar case, the filters/basis functions used in analysis/synthesis are the same as the basis is orthonormal.

Aside: How X is Represented

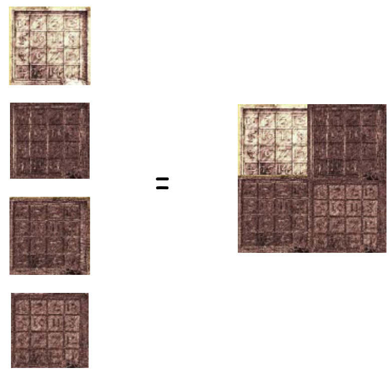

Same images as on Slide 16.

Comments on the slide

Digression: In image processing, and in Matlab, one often represents the four subimages in the same area as the image itself, with the LL subband in the upper left-hand corner.

MR: How?

- x = synthesize (do something) analyze x
- Structure of analyze and synthesize given by TF

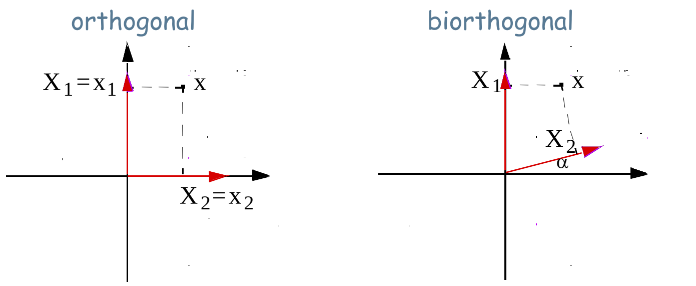

Comments on the slide

We have seen how one can analyze and synthesize. The way we choose these operators are given to us by time-frequency constraints. This slide illustrates two possible nonredundant bases, orthogonal and biorthogonal.

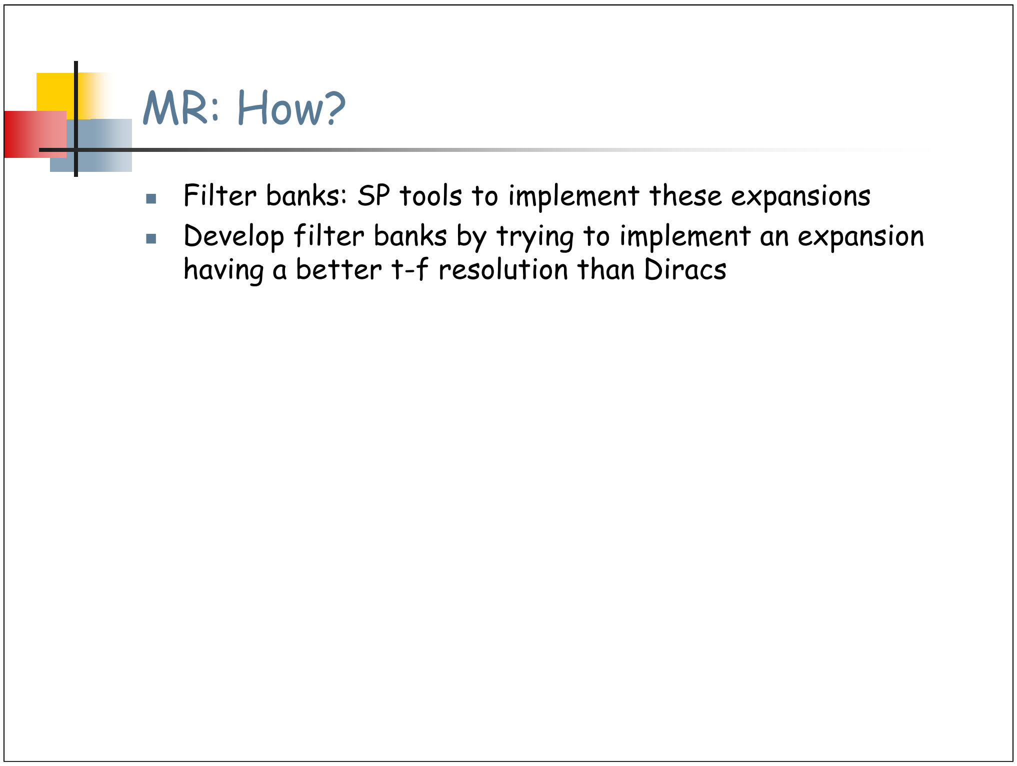

Comments on the slide

We now answer how to implement these multiresolution expansions.

Slightly Better TF than Diracs

- Split the frequency in half
  means we can downsample by 2
  to reconstruct upsample by 2
  filter to remove unwanted images and add
- Basic building block: Two-channel filter bank

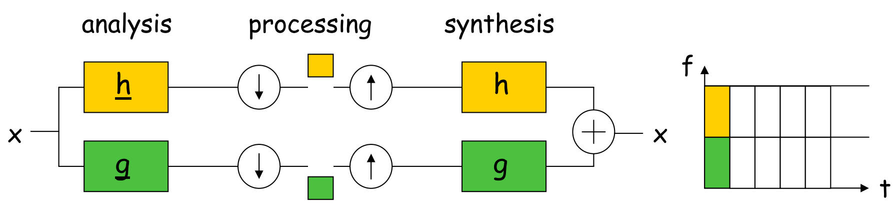

Comments on the slide

The motivation is to obtain a time-frequency analysis slightly better than that of Diracs (impulses). This means we want to refine the frequency resolution. A natural way is to try to divide the spectrum in two parts. As the Nyquist frequency has now been lowered, we can sample, and to reconstruct we have to upsample and filter again. This gives rise to a basic building block --- the two-channel filter bank. As you can see, the tf tiling is indeed better in frequency with a slightly worse one in time.

Slightly Better TF than Diracs

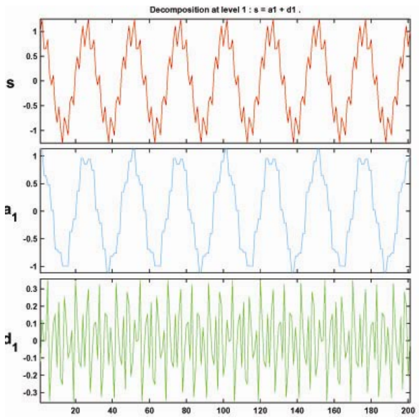

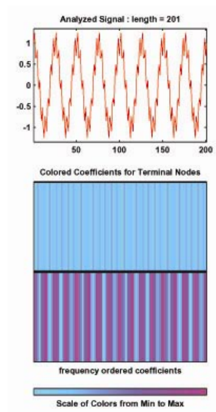

•**Graph on the left:** twosines.jpg

How obtained: W1D

File → Load → Signal → twosines.mat

Choose Wavelet: Haar, Level: 1

Click Analyze

•**Graph on the right:** twosines_WP_1_STFT.jpg

How obtained: WP1D

File → Load → Signal → twosines.mat

Choose Wavelet: Haar, Level: 1

Click Analyze

Comments on the slide

An example of how this works. This graph has a sum of two sines, one large of low frequency, and one tiny of much higher frequency, superimposed on top of the large one. If you look at the second and third graph on the left side, you can see that the coarse channel gets the low-frequency sinusoid out and the other channel gets the tiny sinusoid almost out. Also in the tiling on the right side, it is obvious that the low-frequency sinusoid has been well separated. The high-frequency one is not that obvious since its amplitude is really small.

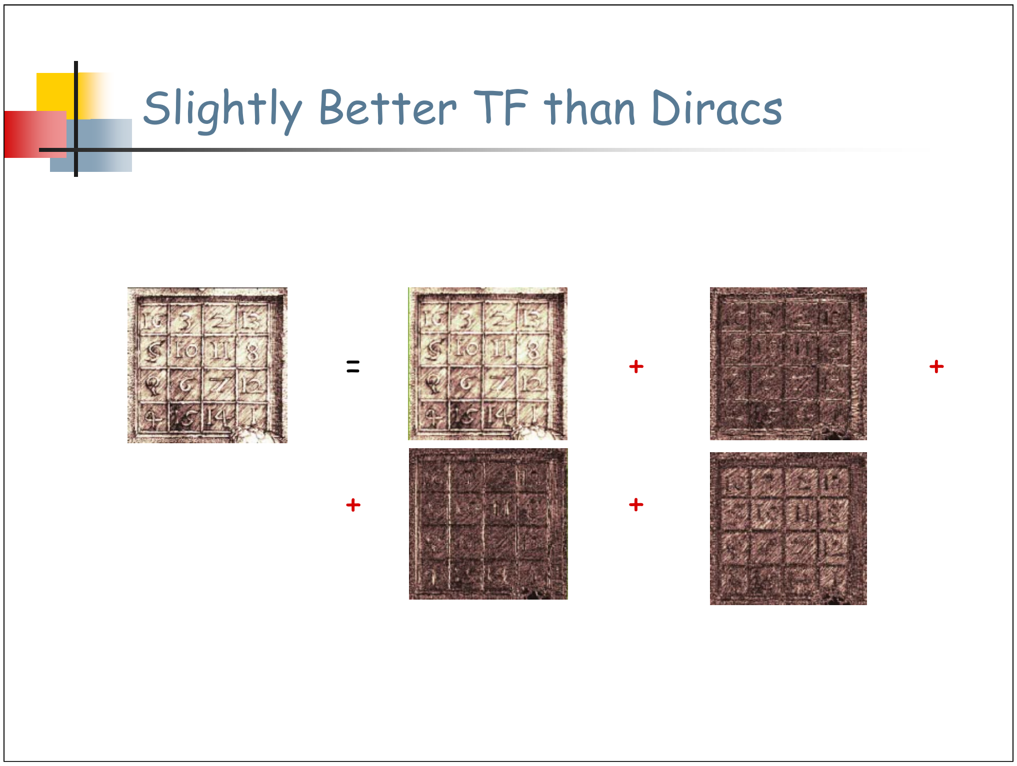

Same images as on Slide 17.

Comments on the slide

In 2D, we typically use the filter banks first across the rows and then across the columns (although this is not really necessary). The result are four subbands as we have seen before.

What Else Can We Do?

- Grow the full tree – STFT/Gabor

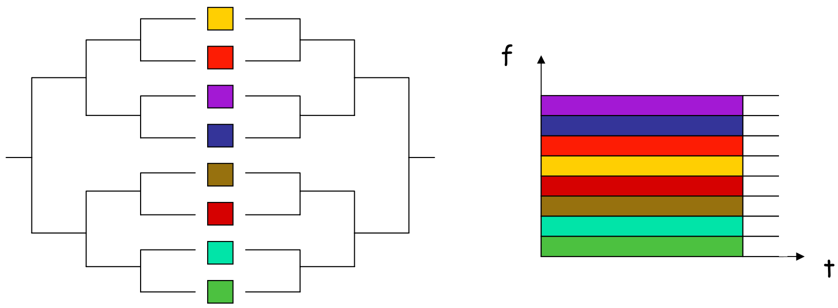

Comments on the slide

What else can we do? We can split into a full tree, splitting every subband further, resulting in a STFT splitting.

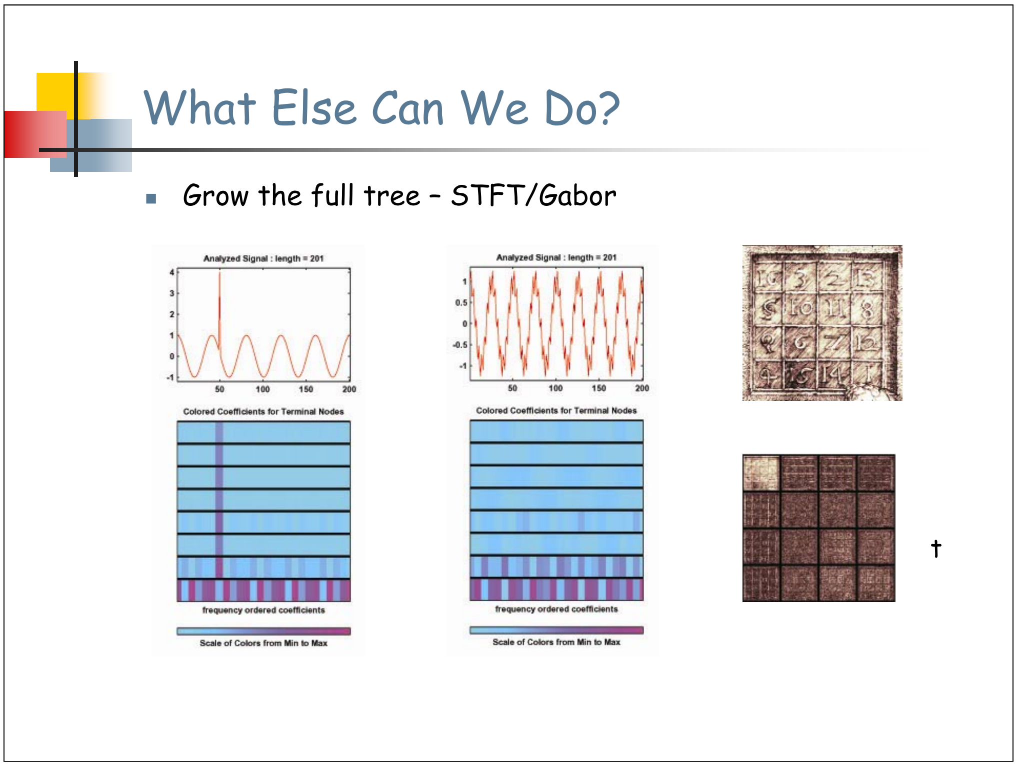

•Graph on the right: sin_0.05pi_dirac_stft.jpg

How obtained: WP1D

File → Load → Signal → sin_0.05pi_dirac.mat

Choose Wavelet: Haar, Level: 3

Click Analyze

•Graph in the middle: twosines_WP_3_STFT.jpg

How obtained: WP1D

File → Load → Signal → twosines.mat

Choose Wavelet: Haar, Level: 3

Click Analyze

•Images on the right: detail_WP_2_STFT.jpg

How obtained: WP2D

File → Example Analysis → db1-depth:1-ent:shannon → detail

Choose Wavelet: Haar, Level: 2

Click Analyze

Comments on the slide

This slides shows the STFT for the low frequency plus an impulse, two sinusoids and a 2D signal. One can see how the STFT isolates the sinusoids well but not the impulse.

What Else Can We Do?

- Iterate only on the lowpass channel – DWT

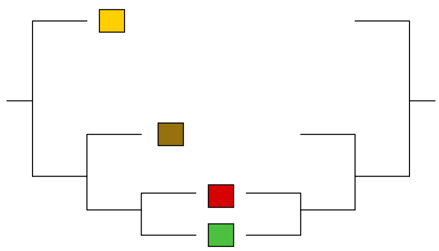

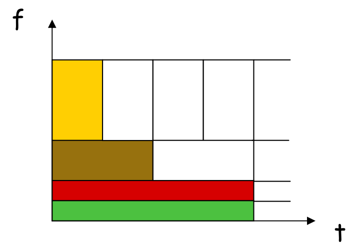

Comments on the slide

Another option is to iterate the splitting only on the lowpass channel, leading to the DWT.

What Else Can We Do?

- Iterate only on the lowpass channel – DWT

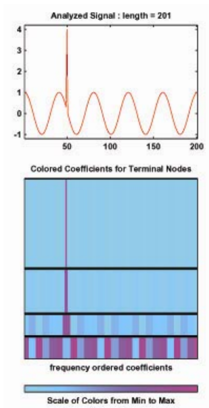

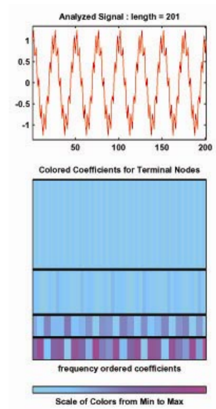

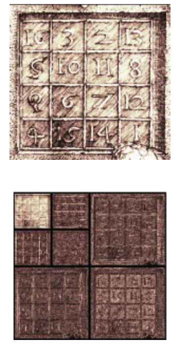

•Graph on the right: sin_0.05pi_dirac.jpg

How obtained: WP1D

File → Load → Signal → sin_0.05pi_dirac.mat

Choose Wavelet: Haar, Level: 3

Click Analyze

•Graph in the middle: twosines_WP_3_DWT.jpg

How obtained: WP1D

File → Load → Signal → twosines.mat

Choose Wavelet: Haar, Level: 3

Click Analyze

Click Wavelet Tree

•Images on the right: detail_WP_2_DWT.jpg

How obtained: WP2D

File → Example Analysis → db1-depth:1-ent:shannon → detail

Choose Wavelet: Haar, Level: 2

Click Analyze

Click Wavelet Tree

Comments on the slide

Here is an example of the DWT applied on the low frequency plus an impulse, two sinusoids and a 2D signal. Again, one can see the trade-off offered by the DWT, it isolates well both the low-frequency sinusoid and the impulse. However, it does not isolate well the high-frequency sinusoid in the second graph.

What Else Can We Do?

- First grow the full tree

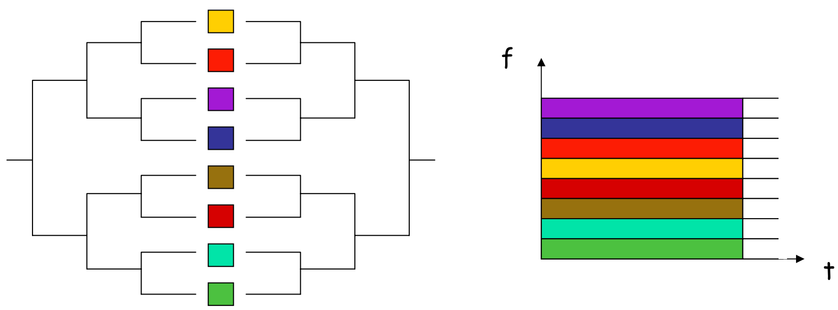

Comments on the slide

Finally, to adapt ourselves to the signal, we first grow the full tree.

What Else Can We Do?

- Then prune the tree – wavelet packets (WP)

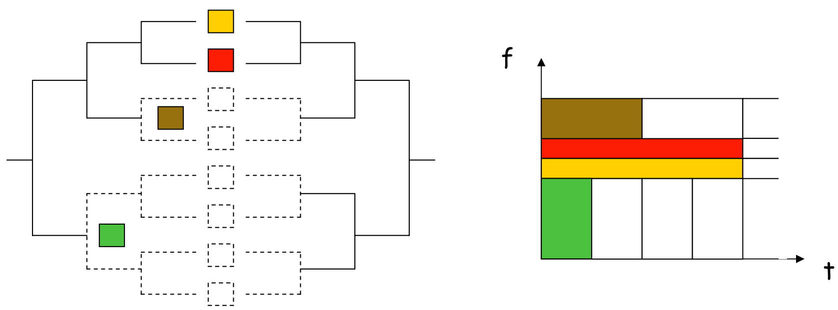

Comments on the slide

We then prune the tree based on some cost function. For example, we might ask ourselves, is it going to cost me more or less to code these two children or the parent? If it costs me less to code the children, I keep the children, if not, I keep the parent. This yields a time-frequency tiling adapted to the signal at hand.

Wavelet Packets

- First stage: full decomposition

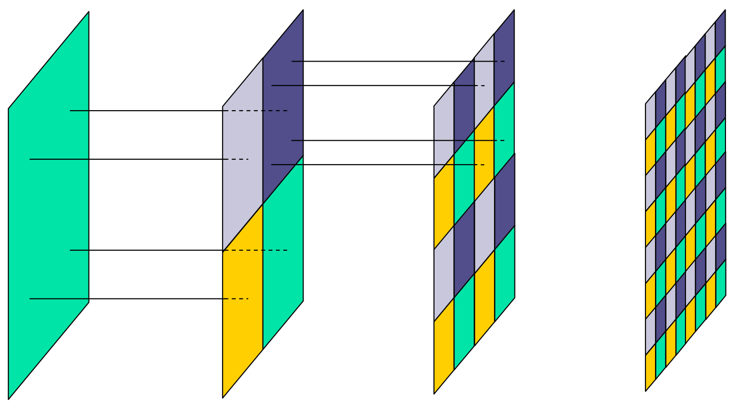

Wavelet Packets

- Second stage: pruning

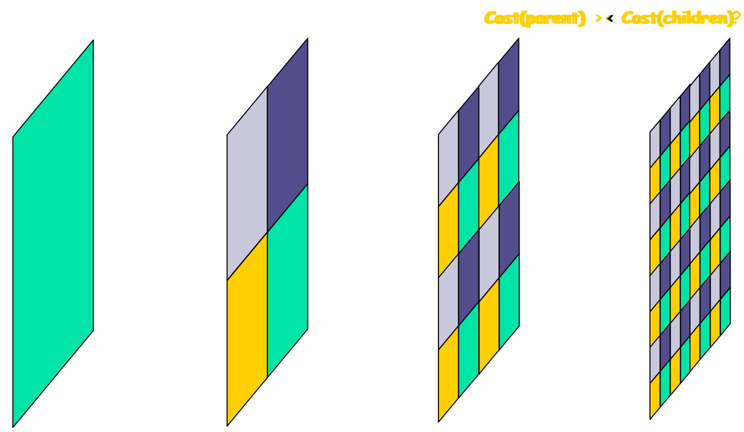

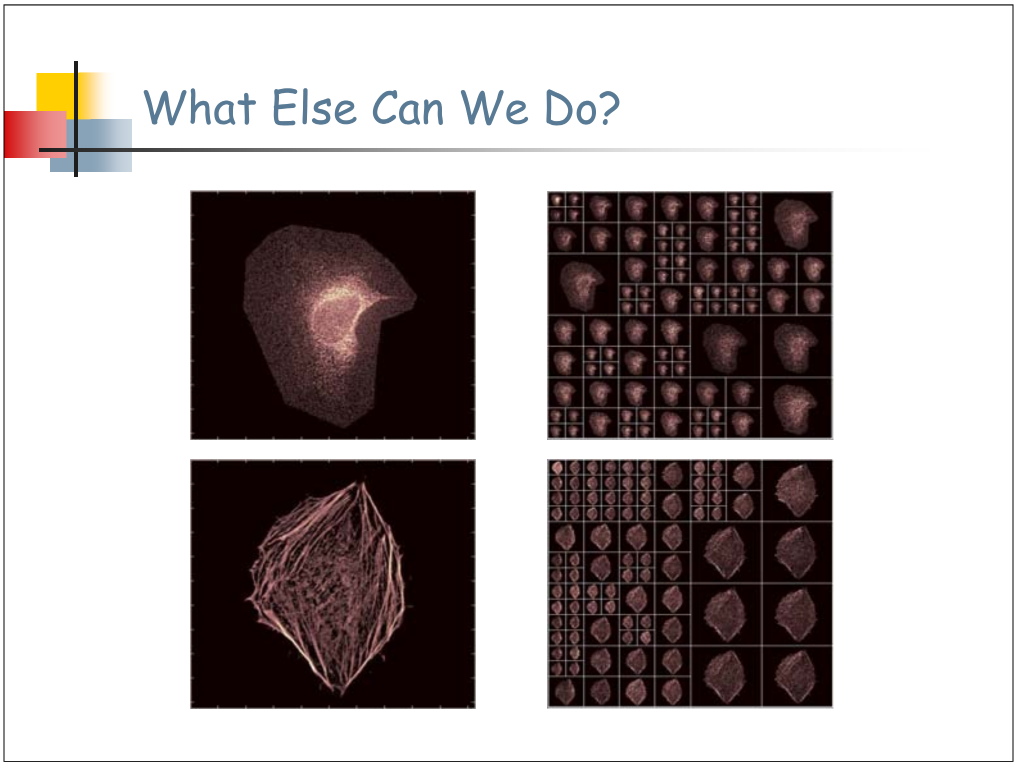

Same images as on Slide 16.

Comments on the slide

Here is an example with the two proteins and WP trees for each one.

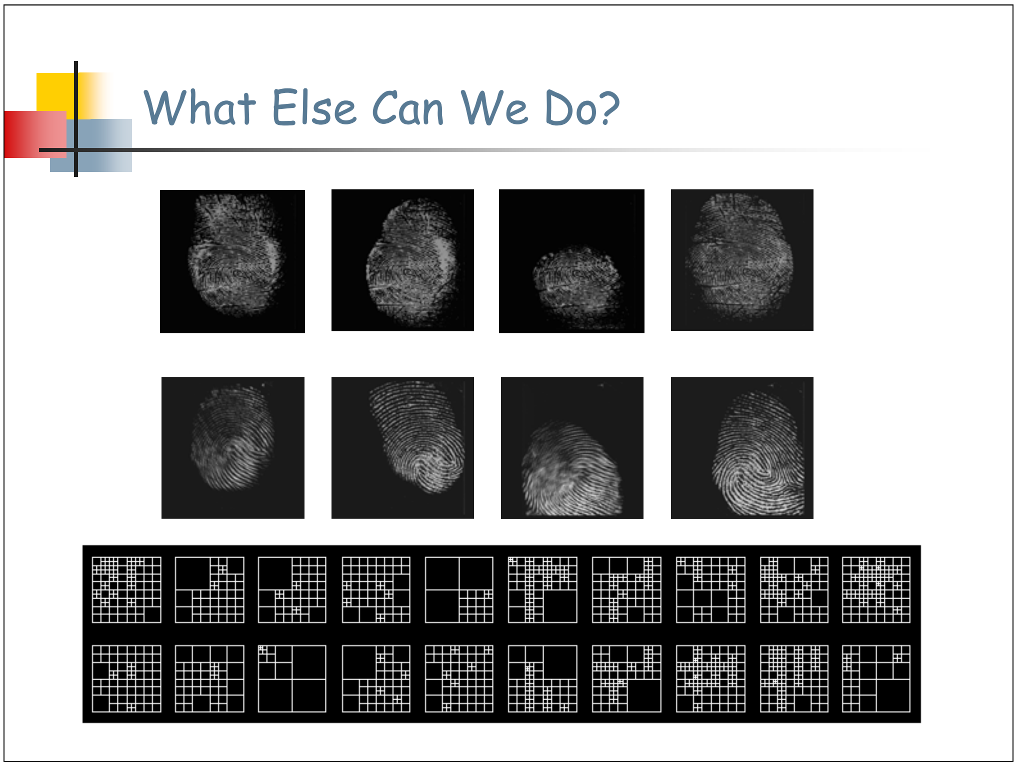

These images were produced by Pablo Henning Yeomans and Jason Thornton.

This was in the context of the identification/verification for biometrics.

Comments on the slide

This example is for fingerprint images (8 of them given on top). The WP trees for all 20 classes are given on the bottom illustrating how WP trees actually serve to distinguish between different classes.

What Else Can We Do?

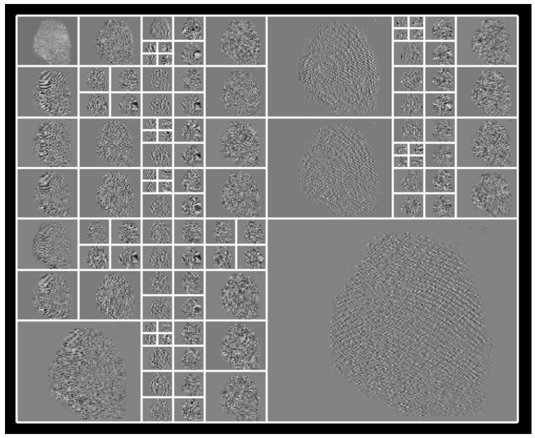

This image was produced by Pablo Henning Yeomans and Jason Thornton.

This was in the context of the identification/verification for biometrics.

Comments on the slide

This is a particular example of a fingerprint with its wavelet packet tree.

What If We Want to Go Redundant?

- Frames: Nonredundant decompositions
  - Robustness to noise
  - Robustness to losses
  - Freedom in design
  - Shift-invariance

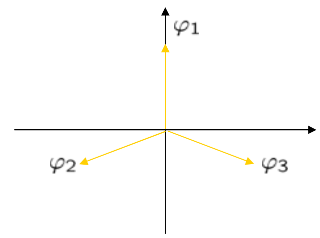

$$F = \begin{pmatrix} \varphi_{1,1} & \varphi_{1,2} \\ \varphi_{2,1} & \varphi_{2,2} \\ \varphi_{3,1} & \varphi_{3,2} \end{pmatrix}$$

Bases versus Frames?

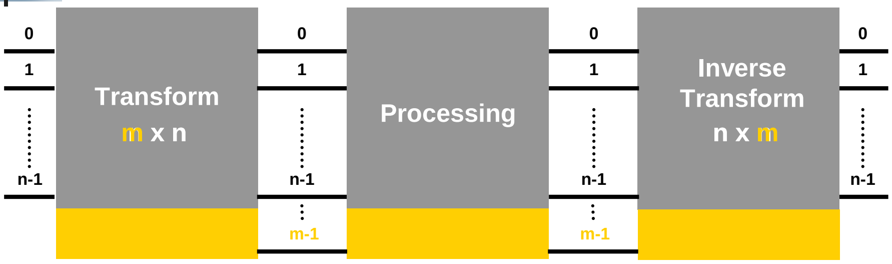

- Bases are nonredundant
  - Loss of one transform coefficient is irreplaceable
  - Sensitivity to noise is great
  - Space of possible solutions is restricted
- Solution: frames

Robustness to Noise

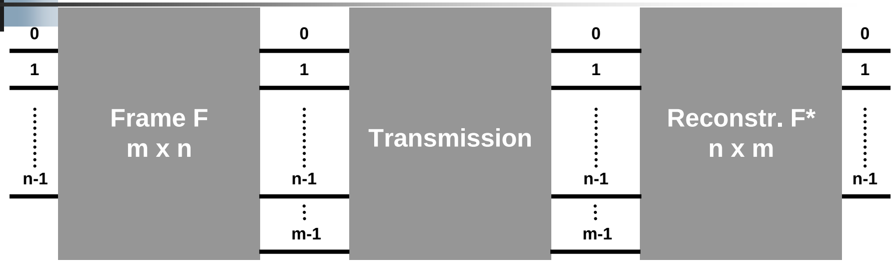

- Noise is spread over more components: easier to clean

$$MSE_{basis} = \sigma^2 = \frac{m}{n} MSE_{frame}$$

Robustness to Losses

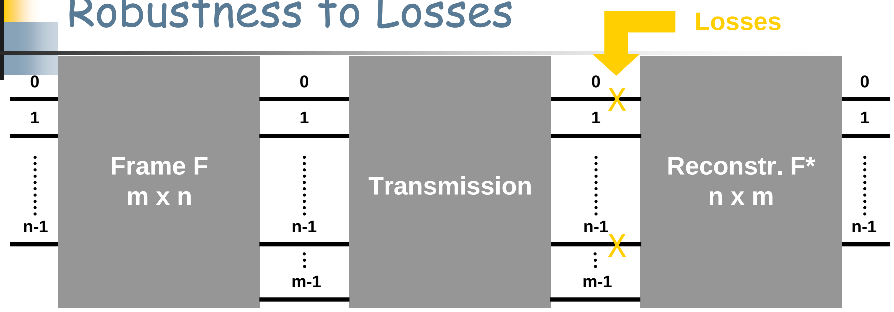

- Losses
  - Modeled as erasures
  - To reconstruct, inverse transform must exist
  - Mathematically: any (n x n) submatrix of the frame matrix must be full rank
    ➔ maximally robust to erasures (MR)

What are Frames?

- Generating system for R^n or C^n $\{\varphi_0, \ldots, \varphi_{m-1}\}$
- Usually represented by a matrix F

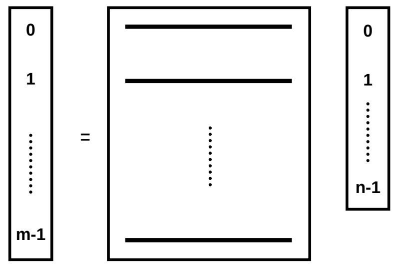

Frame coefficients y = F x

Frame Properties

Maximally robust (MR) | Tight (T) | Equal norm (EN)

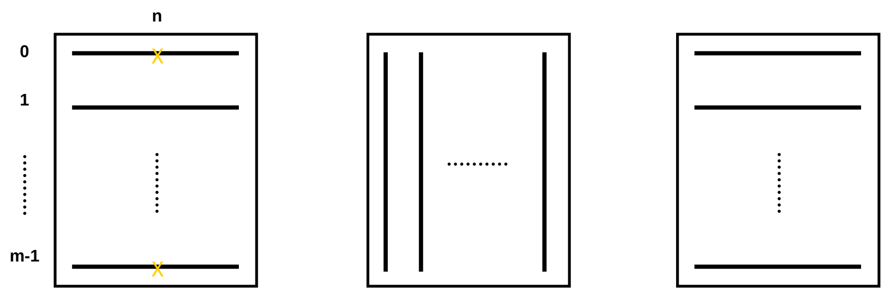

Any (n x n) submatrix is full rank | Columns are orthonormal | All rows have equal norm

Biomedical Imaging

- How to decide what to use?
  - Properties of signals
  - Properties of transforms
  - Task to perform

Properties of Bioimages

- Noise levels and types
- Lack of "ground truth"
- Large deviations
- Low definition and contrast
- Wide range of time- and frequency-localized features

Transform Properties

- Finite signal?
  - Boundary effects
- Representation type
  - Redundant or not?
- Orthogonal/tight vs. general bases/frames?
- Adaptivity and localization
- Properties of basis functions
  - Symmetry?
- Issues for images
  - Sampling lattice
  - Basis functions
    - Separable vs. nonseparable
    - Real vs. complex
    - Orientation selectivity
    - Rotational invariance?

Task to Perform

- Denoising
- Enhancement
- Feature extraction
- Classification
- Segmentation

Conclusions

- MR toolbox
  - Rich enough to address a whole set of problems

References on MR

- Light reading
  - "Wavelets: Seeing the Forest -- and the Trees", D. Mackenzie, Beyond Discovery, December 2001.
- Books
  - "A Wavelet Tour of Signal Processing", S. Mallat, Academic Press, 1999.
  - "Ten Lectures on Wavelets", I. Daubechies, SIAM, 1992.
  - "Wavelets and Subband Coding", M. Vetterli and J. Kovacevic, Prentice Hall, 1995.
  - "Wavelets and Filter Banks", G. Strang and T. Nguyen, Wells. Cambr. Press, 1996.
- Bioimaging
  - "A Review of Wavelets in Biomedical Applications", M. Unser and A. Aldroubi, Proc. IEEE, April 1996.
  - "Wavelets in Temporal and Spatial Processing of Biomedical Data", A. Laine, Annu. Rev. Biomed. Eng., 2000.
  - "Guest Editorial: Wavelets in Medical Imaging", M. Unser, A. Aldroubi and A. Laine, IEEE Trans. On Medical Imaging, March 2003.
  - "Wavelets in Bioinformatics and Computational Biology: State of the art and Perspectives", P. Lio, Bioinformatics Review, 2003.