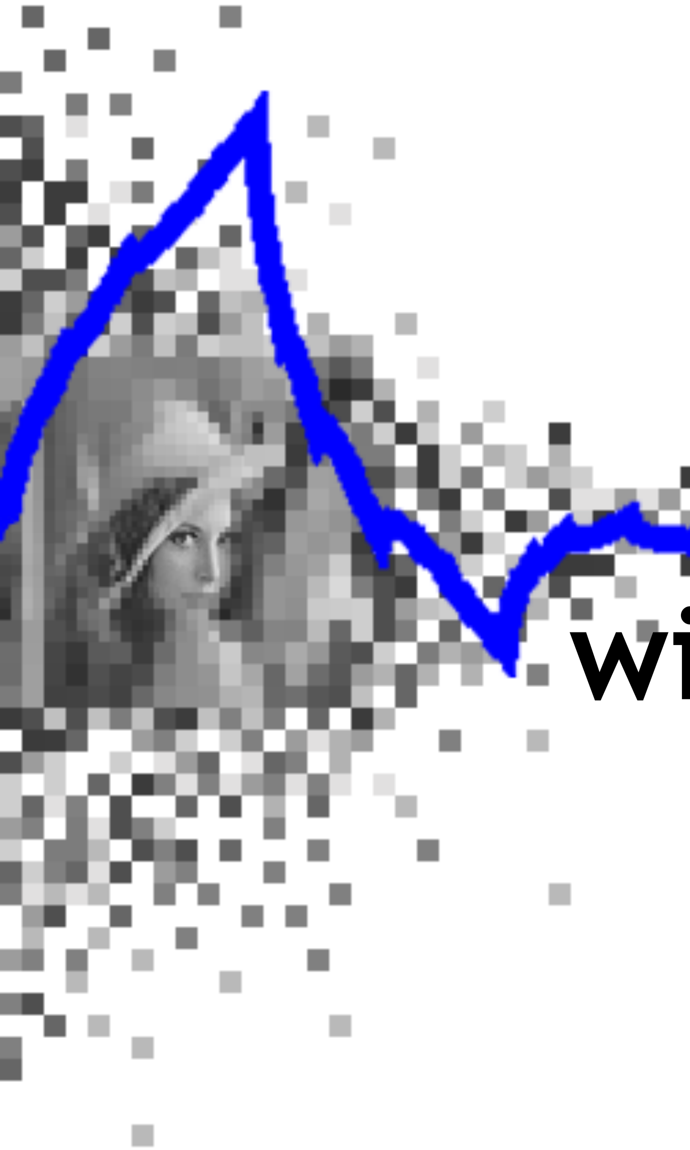

Signal and Image Processing with Orthogonal Decompositions

Gabriel Peyré

http://www.ceremade.dauphine.fr/~peyre/numerical-tour/

CNRS
CEREMADE
DAUPHINE UNIVERSITÉ PARIS

Signals, Images and More

Analog signal: $f_0 \in \mathrm{L}^2([0,1])$.

Analog image: $f_0 \in \mathrm{L}^2([0,1]^2)$.

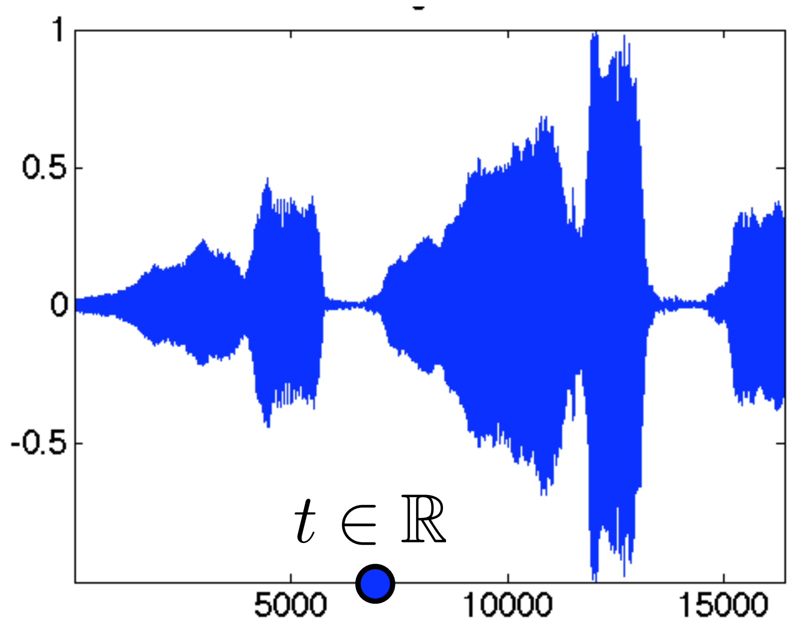

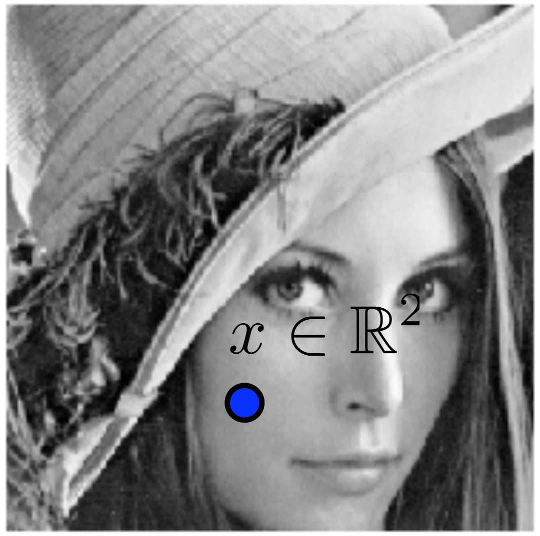

Signals, Images and More

Analog signal: $f_0 \in \mathrm{L}^2([0,1])$.

Analog image: $f_0 \in \mathrm{L}^2([0,1]^2)$.

More complex data sets: $f_0 : [0,1]^d \rightarrow [0,1]^s$.

- videos: $d = 3, s = 1$.

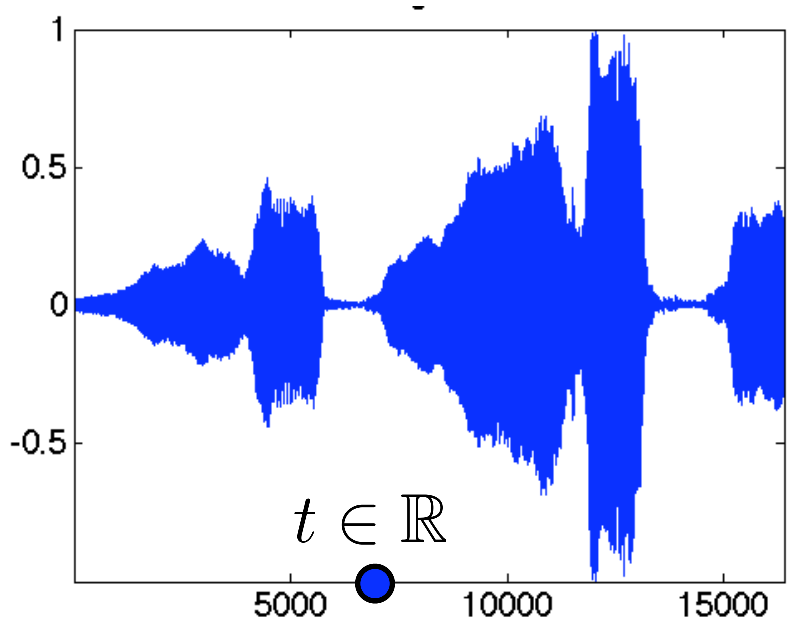

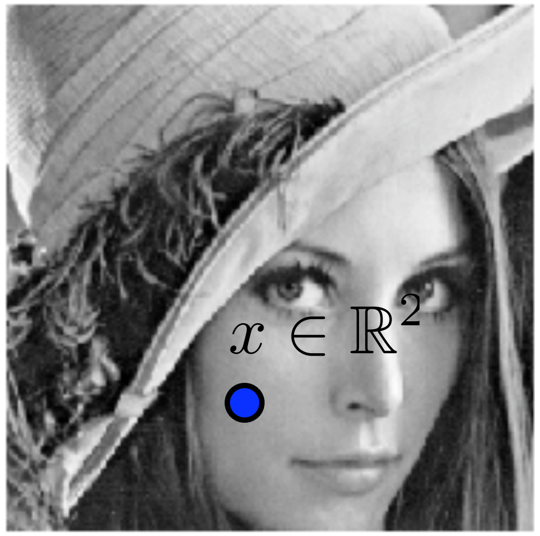

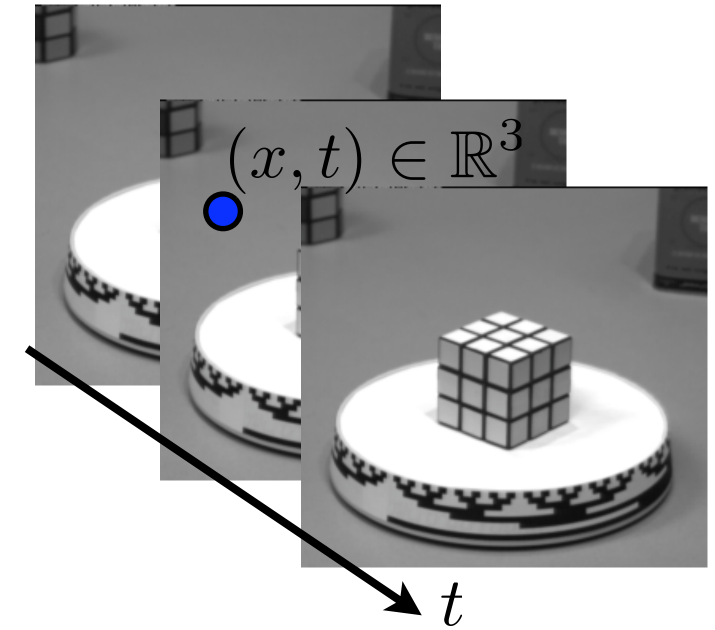

Signals, Images and More

Analog signal: $f_0 \in \mathrm{L}^2([0,1])$.

Analog image: $f_0 \in \mathrm{L}^2([0,1]^2)$.

More complex data sets: $f_0 : [0,1]^d \rightarrow [0,1]^s$.

- videos: $d = 3, s = 1$.
- color images: $d = 2, s = 3$.

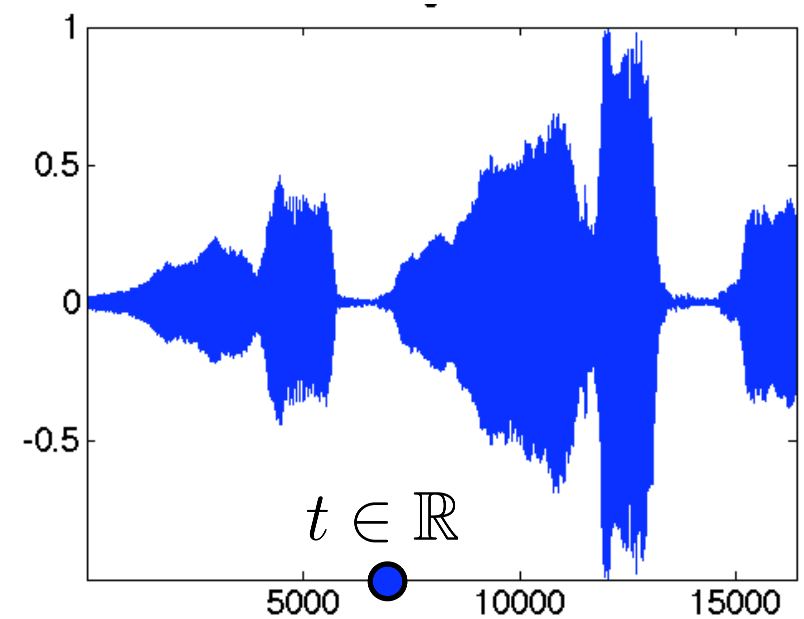

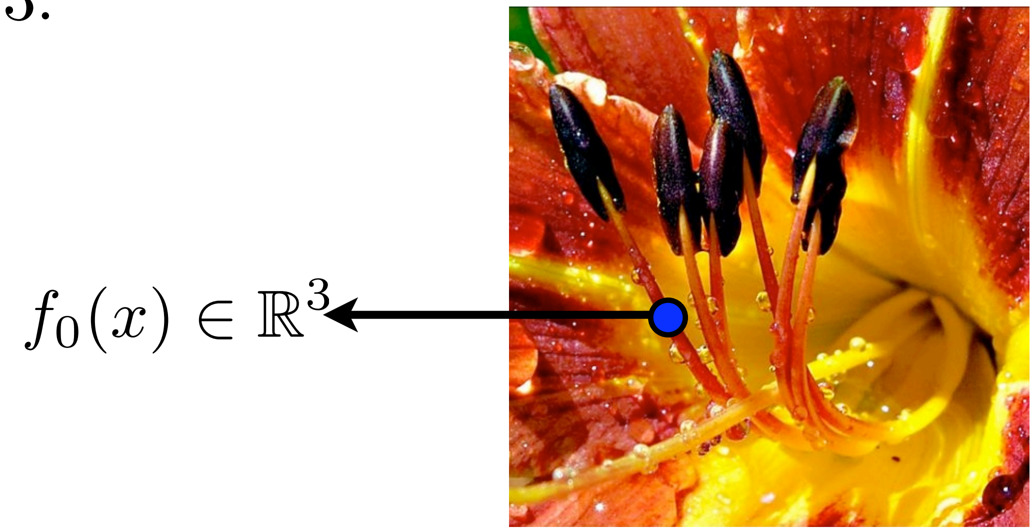

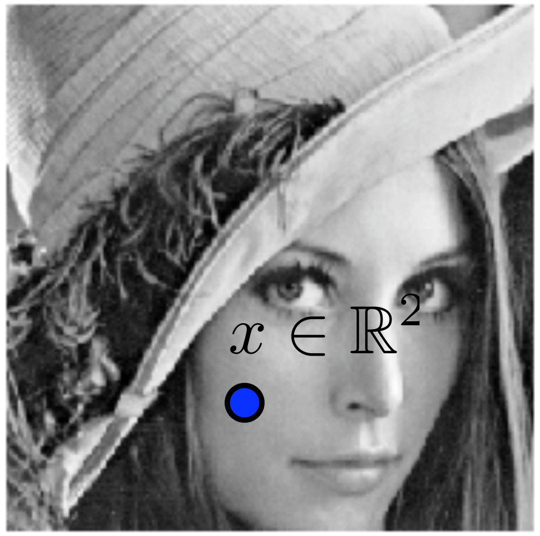

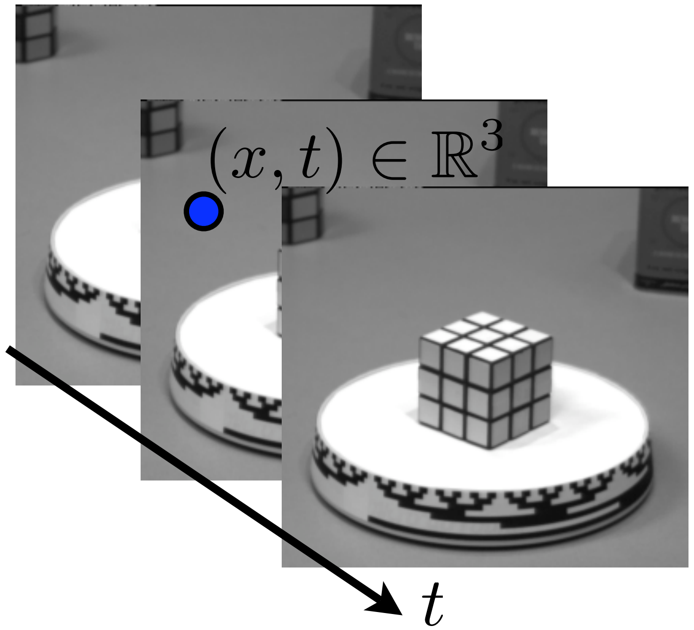

Signals, Images and More

Analog signal: $f_0 \in \mathrm{L}^2([0,1])$.

Analog image: $f_0 \in \mathrm{L}^2([0,1]^2)$.

More complex data sets: $f_0 : [0,1]^d \to [0,1]^s$.

- videos: $d = 3, s = 1$.
- color images: $d = 2, s = 3$.
- multi-spectral images: $d = 2, s \gg 3$.

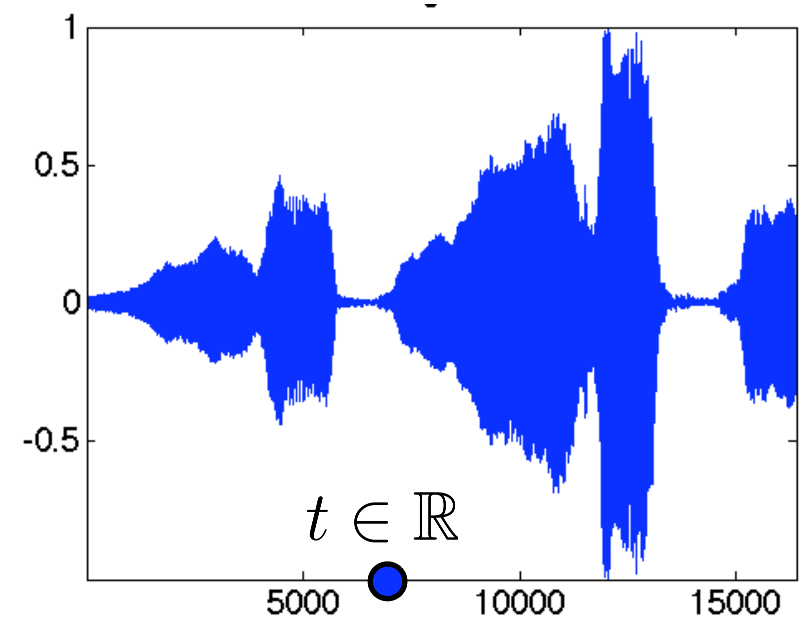

$f_0(x) \in \mathbb{R}^3$

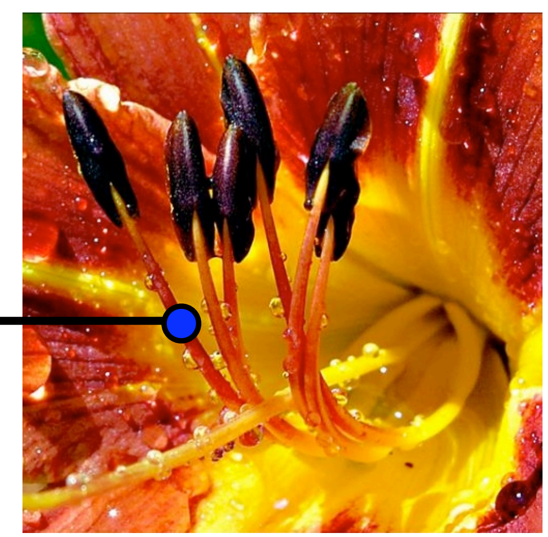

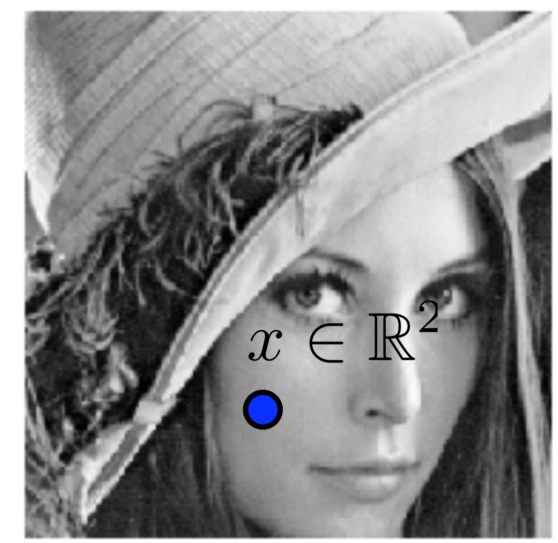

$f_0(x) \in \mathbb{R}^s, s = 32$

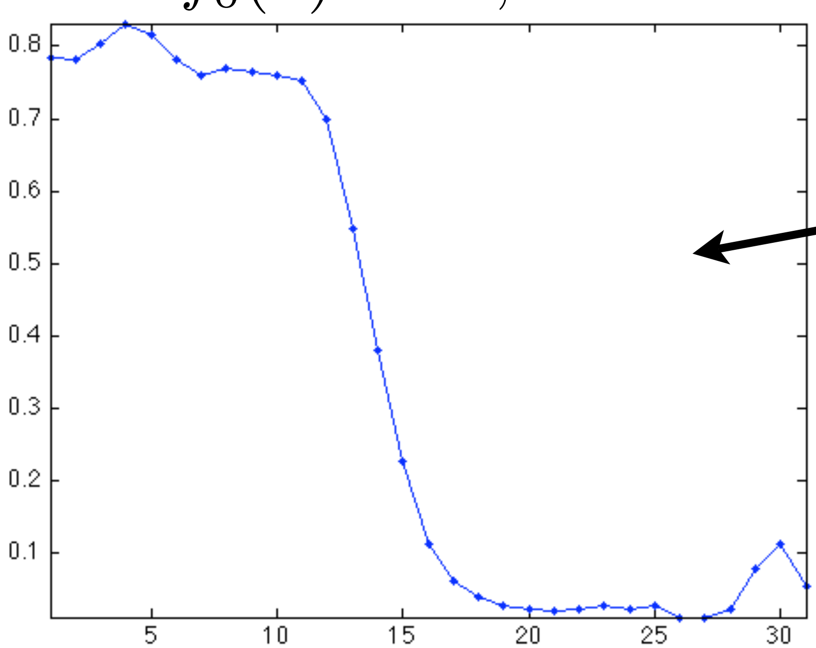

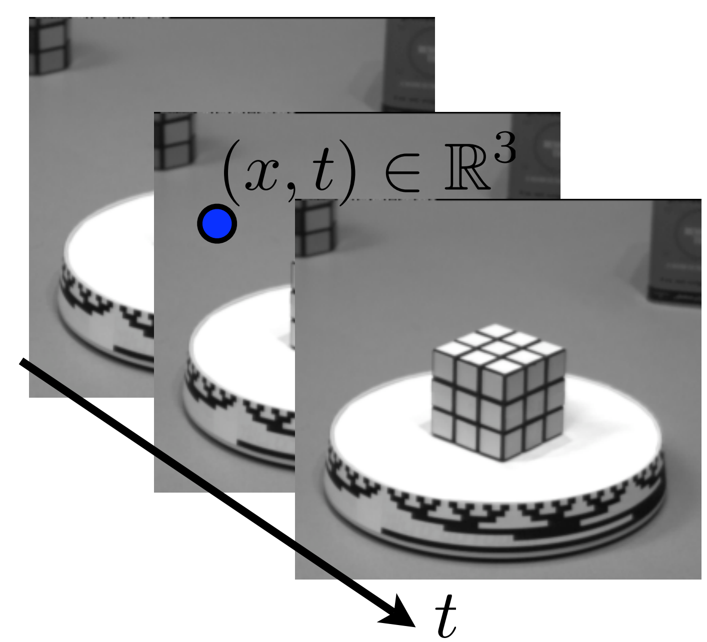

Overview

- 1D Orthogonal Decompositions
- Discretization
- 2D Bases
- Linear and Non-linear Approximations
- Decompositions for Processing

Orthogonal Decompositions

Analog signal $f_0 \in \mathrm{L}^2([0,1]^d)$.

$$\langle f_0,\, g_0 \rangle = \int_{[0,1]^d} f_0(x) \bar{g}_0(x) \mathrm{d}x$$

Orthogonal basis $\{\psi_m\}_{m \in \mathbb{Z}}$ of $\mathrm{L}^2([0,1]^d)$.

Decomposition: $$f_0 = \sum_{m \in \mathbb{Z}} \langle f_0,\, \psi_m \rangle \psi_m$$

Energy conservation:

$$\|f_0\|^2 = \int_{[0,1]^d} |f_0(x)|^2 \mathrm{d}x = \sum_{m \in \mathbb{Z}} |\langle f_0,\, \psi_m \rangle|^2$$

Orthogonal Decompositions

Analog signal $f_0 \in \mathrm{L}^2([0,1]^d)$.

$$\langle f_0,\, g_0 \rangle = \int_{[0,1]^d} f_0(x) \bar{g}_0(x) \mathrm{d}x$$

Orthogonal basis $\{\psi_m\}_{m \in \mathbb{Z}}$ of $\mathrm{L}^2([0,1]^d)$.

Decomposition: $$f_0 = \sum_{m \in \mathbb{Z}} \langle f_0,\, \psi_m \rangle \psi_m$$

Energy conservation:

$$\|f_0\|^2 = \int_{[0,1]^d} |f_0(x)|^2 \mathrm{d}x = \sum_{m \in \mathbb{Z}} |\langle f_0,\, \psi_m \rangle|^2$$

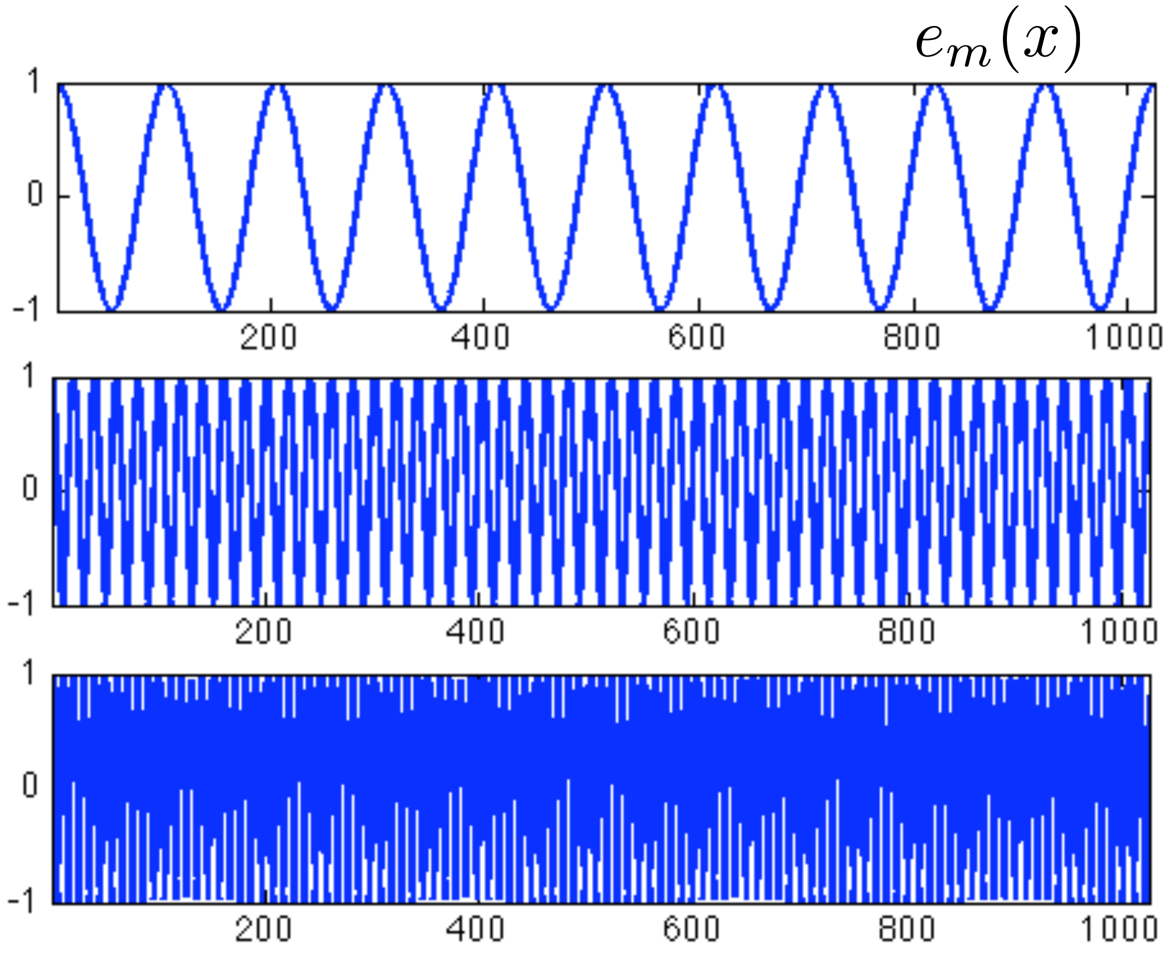

Fourier: $\psi_m(x) = e_m(x) = e^{2i\pi m x}$

$$= \cos(2\pi m x) + i \sin(2\pi m x)$$

Frequency m.

Orthogonal Decompositions

Analog signal $f_0 \in \mathrm{L}^2([0,1]^d)$.

$$\langle f_0,\, g_0 \rangle = \int_{[0,1]^d} f_0(x)\bar{g}_0(x)\mathrm{d}x$$

Orthogonal basis $\{\psi_m\}_{m\in\mathbb{Z}}$ of $\mathrm{L}^2([0,1]^d)$.

Decomposition: $$f_0 = \sum_{m\in\mathbb{Z}} \langle f_0,\, \psi_m \rangle \psi_m$$

Energy conservation:

$$\|f_0\|^2 = \int_{[0,1]^d} |f_0(x)|^2 \mathrm{d}x = \sum_{m\in\mathbb{Z}} |\langle f_0,\, \psi_m\rangle|^2$$

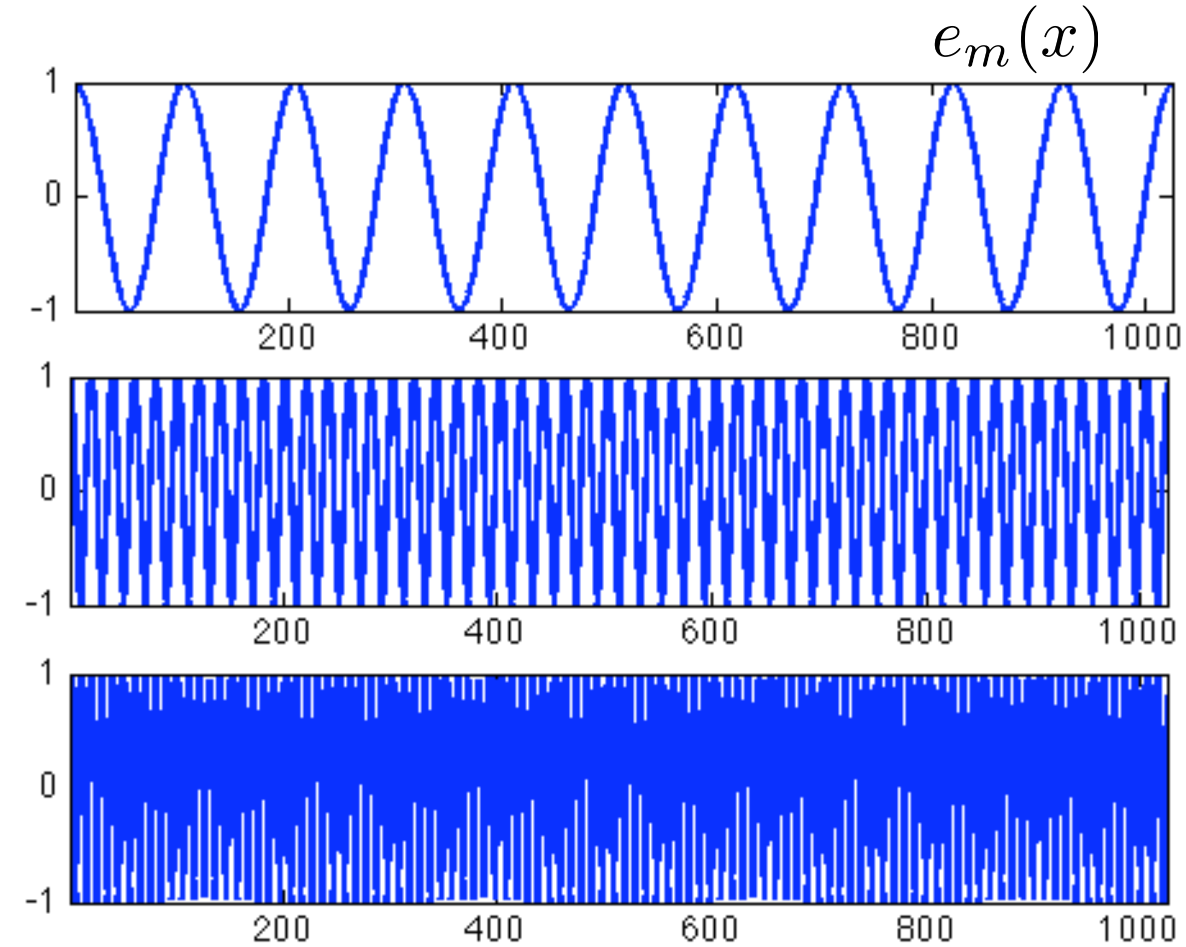

Fourier: $\psi_m(x) = e_m(x) = e^{2i\pi m x}$

$$= \cos(2\pi m x) + i \sin(2\pi m x)$$

Frequency m.

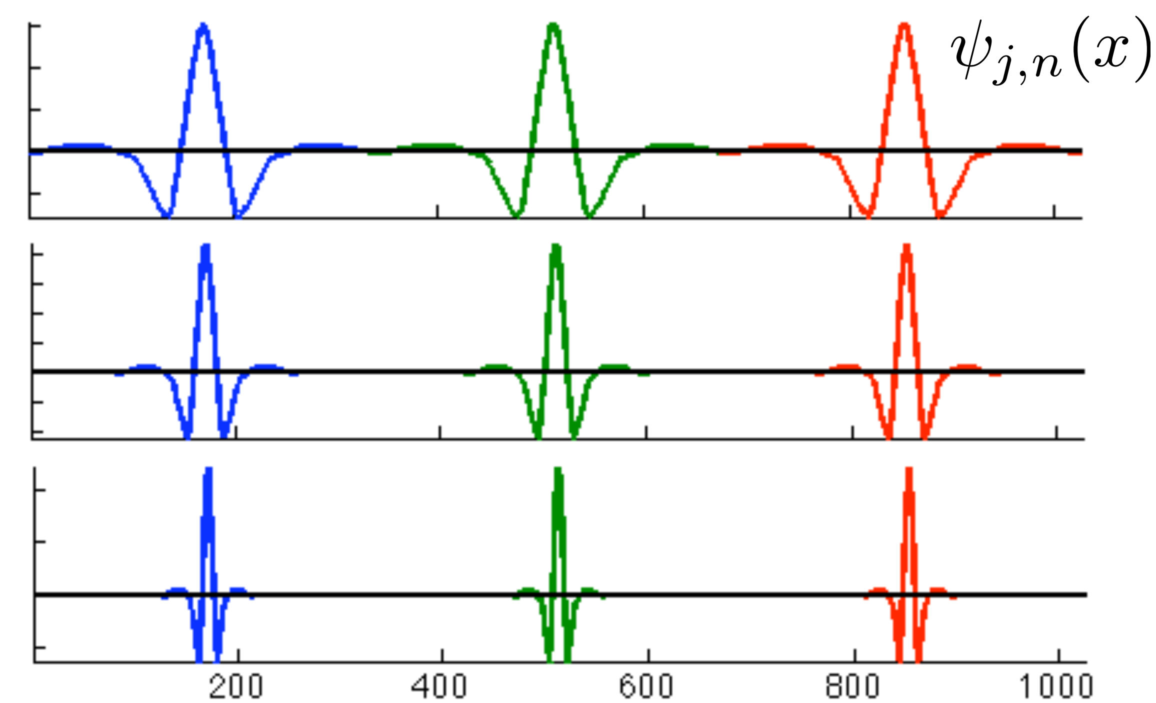

Wavelets: $\psi_m(x) = 2^{j/2}\psi(2^{-j}x - n)$, $m = (j,n)$

Position n.

Scale 2^j.

Examples of Wavelets

Fourier transform: the one and only.

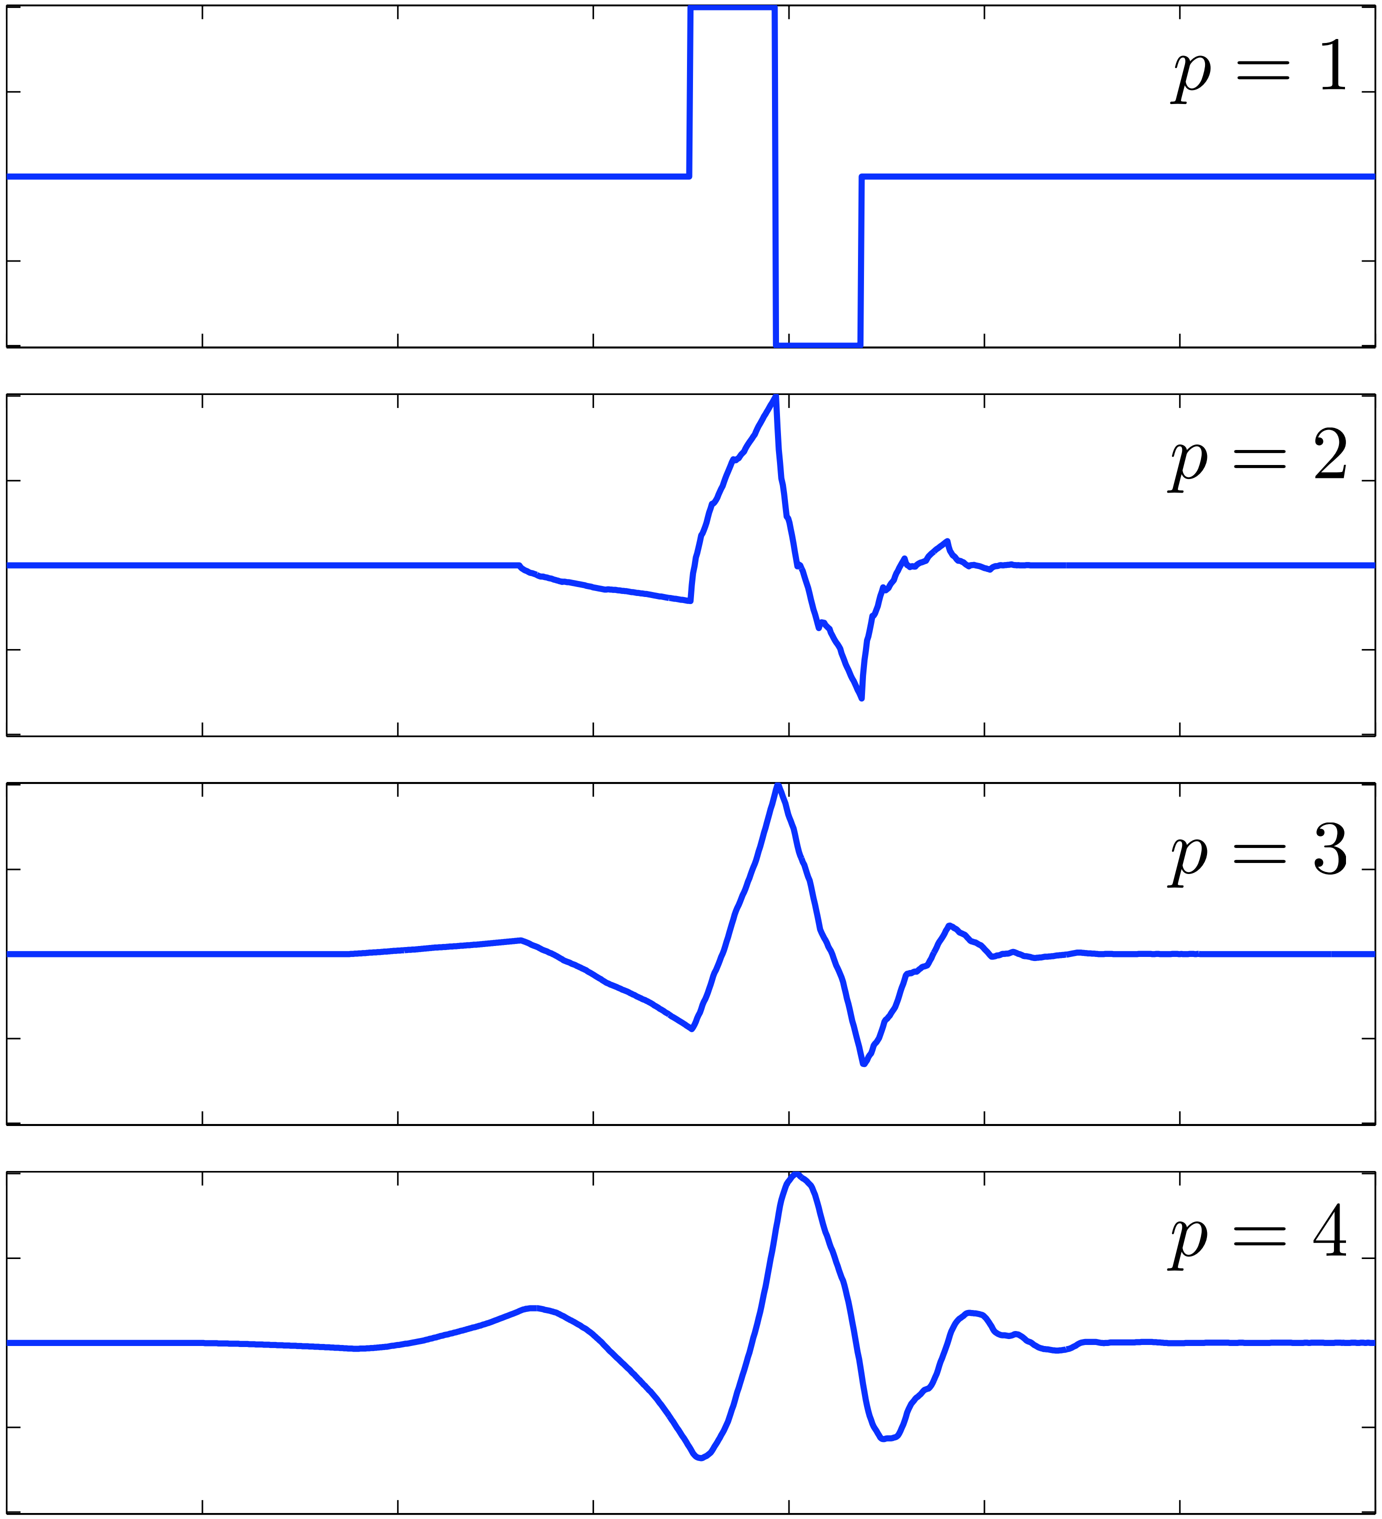

Wavelet transform:

$\rightarrow$ a large choice of "mother" wavelets.

- support size.
- number of oscilations.
- symmetry (non-orthogonal)
- smoothness.

Overview

- 1D Orthogonal Decompositions
- **Discretization**
- 2D Bases
- Linear and Non-linear Approximations
- Decompositions for Processing

Discretization

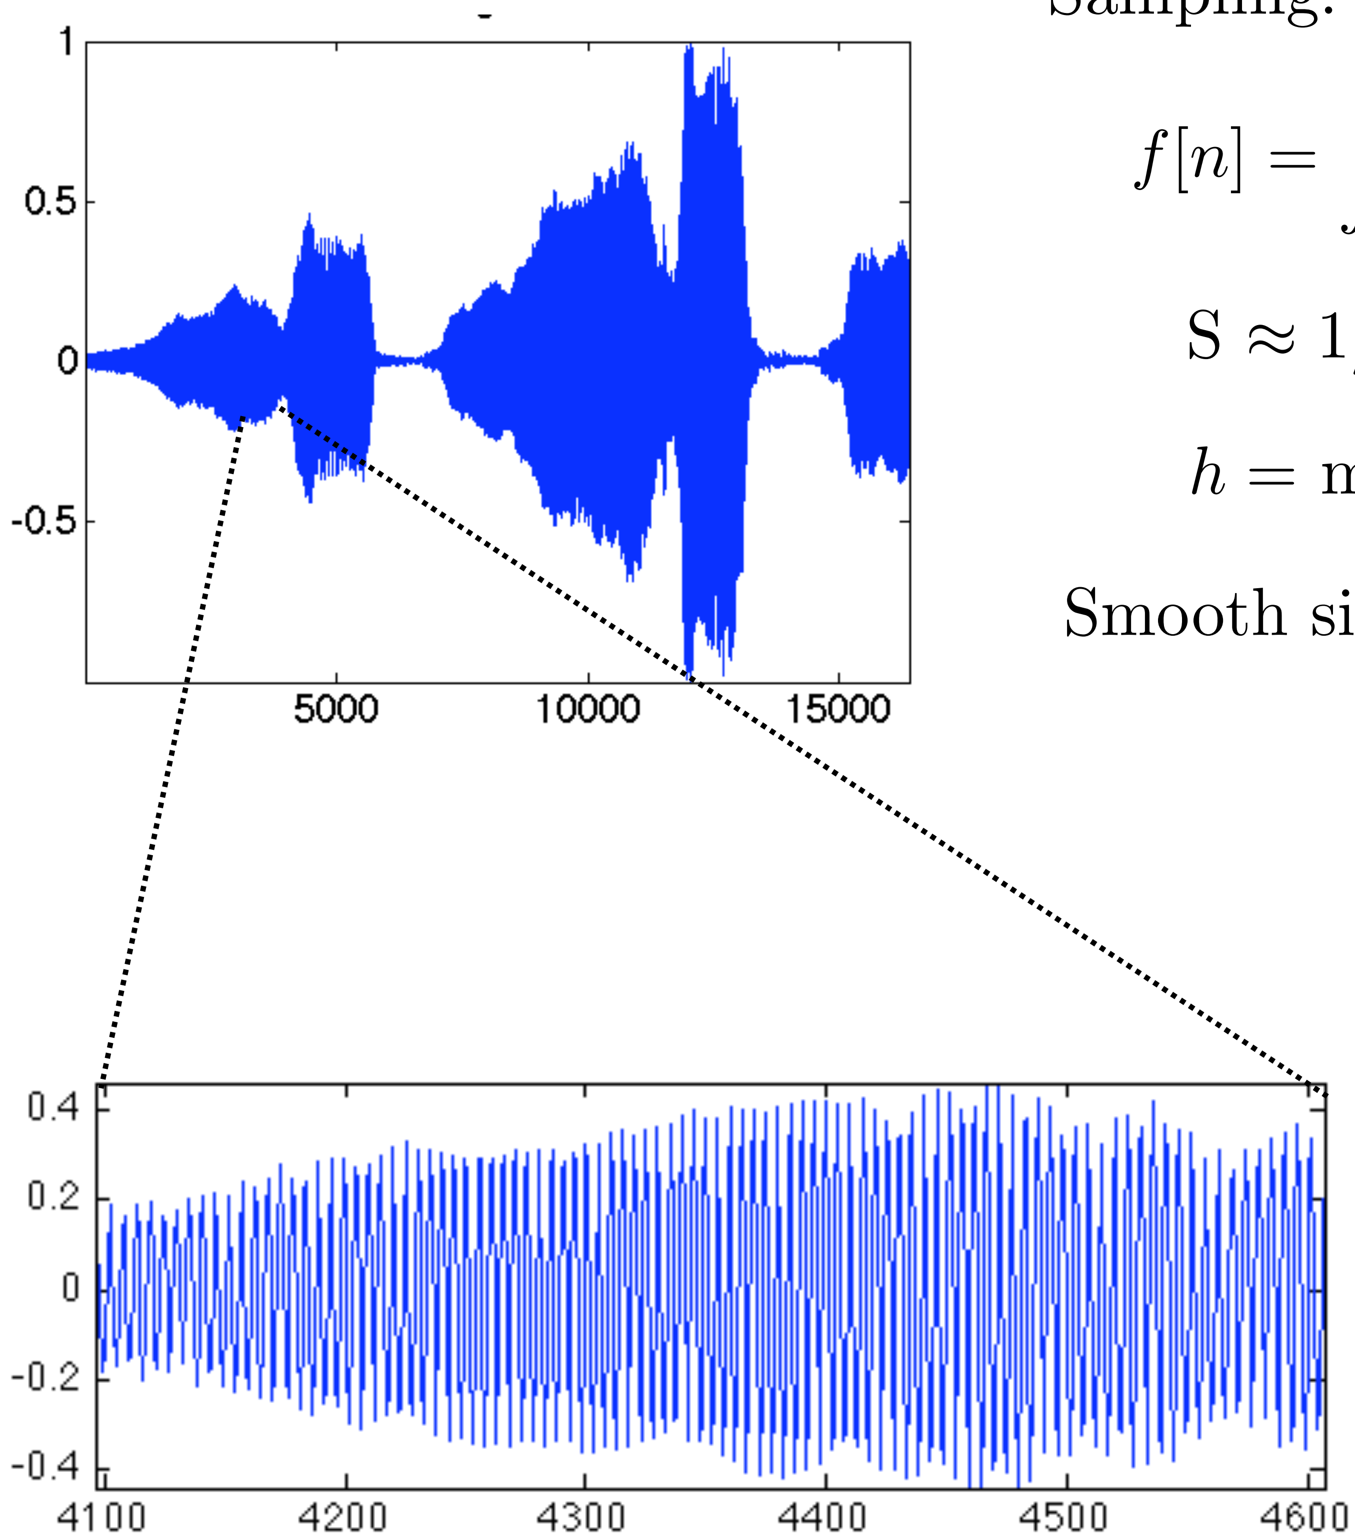

Sampling: $f_0 \in \mathrm{L}^2([0,1]) \mapsto f \in \mathbb{C}^N.$

$$f[n] = \int_{-S/2}^{S/2} f_0(n/N + x) h(x) \mathrm{d}x = f_0 \star \tilde{h}(n/N).$$

S $\approx 1/N$ kernel suport size. $\tilde{h}(x) = h(-x)$

h = micro/optics impulse response.

Smooth signals: $f[n] \approx f_0(n/N)$.

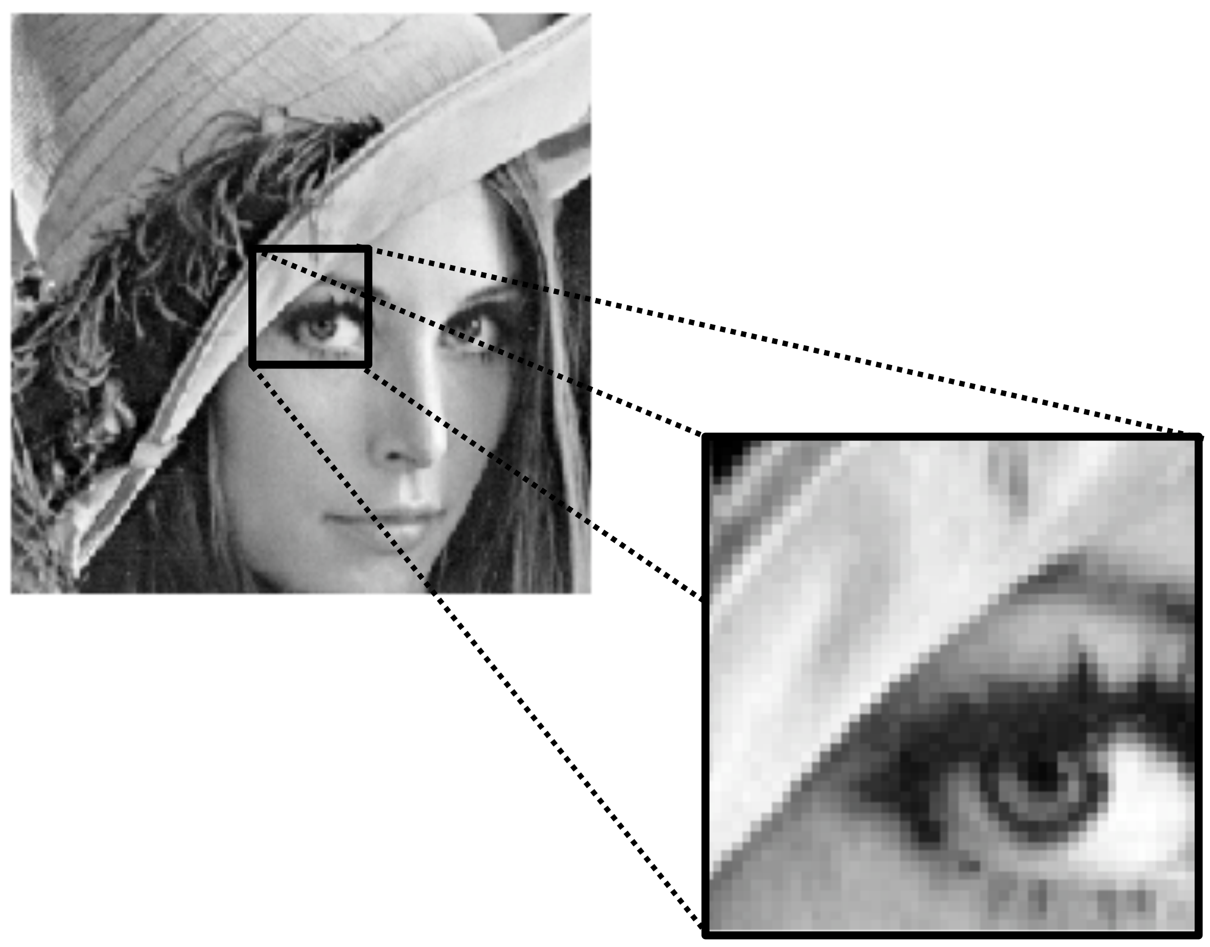

Uniform Sampling and Smoothness

Uniform sampling: $f[i] = f(i/N) = \langle f, \varphi_i \rangle$ where $\varphi_i = \delta_{i/N}$ Dirac.

Shannon Theorem: f is recovered from $\{\langle f, \varphi_i \rangle\}_{i=0}^{P-1}$ if $\text{Supp}(\hat{f}) \subset [-N\pi, N\pi]$.

Linear reconstruction formula: $$f(t) = \sum_i f[i] h(t - i/N) \quad \text{where} \quad h(t) = \frac{\sin(\pi N t)}{\pi N t}$$

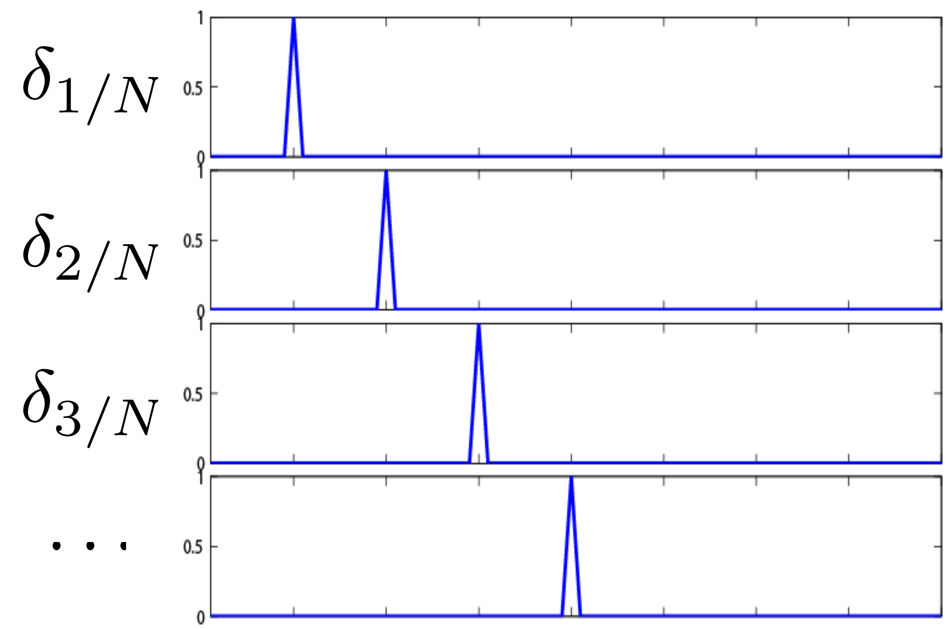

pointwise sampling

Uniform Sampling and Smoothness

Uniform sampling: $f[i] = f(i/N) = \langle f, \varphi_i \rangle$ where $\varphi_i = \delta_{i/N}$ Dirac.

Shannon Theorem: f is recovered from $\{\langle f, \varphi_i \rangle\}_{i=0}^{P-1}$ if $\text{Supp}(\hat{f}) \subset [-N\pi, N\pi]$.

Linear reconstruction formula: $$f(t) = \sum_i f[i] h(t - i/N) \quad \text{where} \quad h(t) = \frac{\sin(\pi N t)}{\pi N t}$$

$\longrightarrow$ works for smooth signals/images.

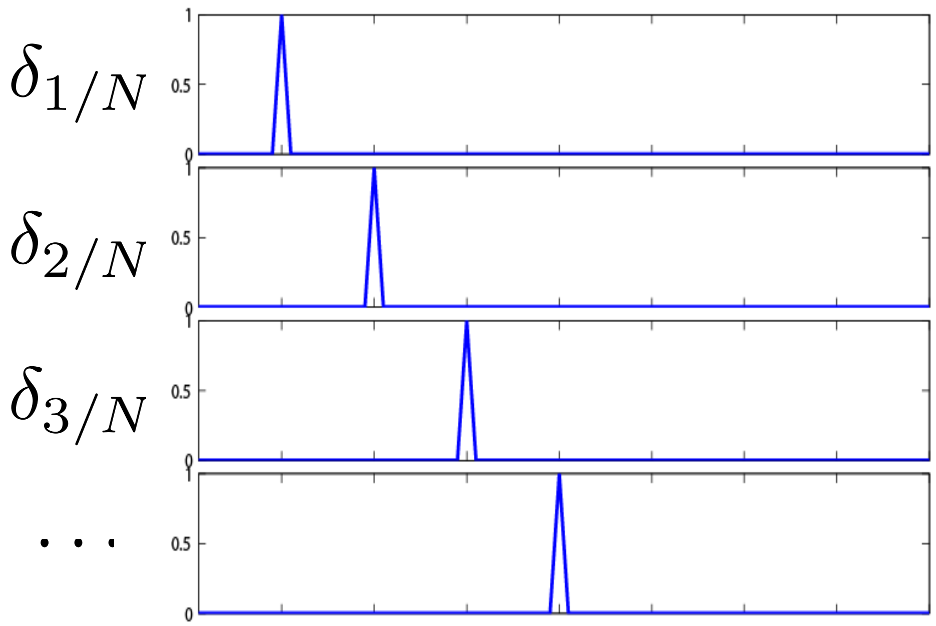

pointwise sampling

Uniform Sampling and Smoothness

Uniform sampling: $f[i] = f(i/N) = \langle f, \varphi_i \rangle$ where $\varphi_i = \delta_{i/N}$ Dirac.

Shannon Theorem: f is recovered from $\{\langle f, \varphi_i \rangle\}_{i=0}^{P-1}$ if $\text{Supp}(\hat{f}) \subset [-N\pi, N\pi]$.

Linear reconstruction formula: $$f(t) = \sum_i f[i] h(t - i/N) \quad \text{where} \quad h(t) = \frac{\sin(\pi N t)}{\pi N t}$$

$\longrightarrow$ works for smooth signals/images.

$\longrightarrow$ discontinuities / edges requires a large P.

$\delta_{1/N}$

$\delta_{2/N}$

$\delta_{3/N}$

$\ldots$

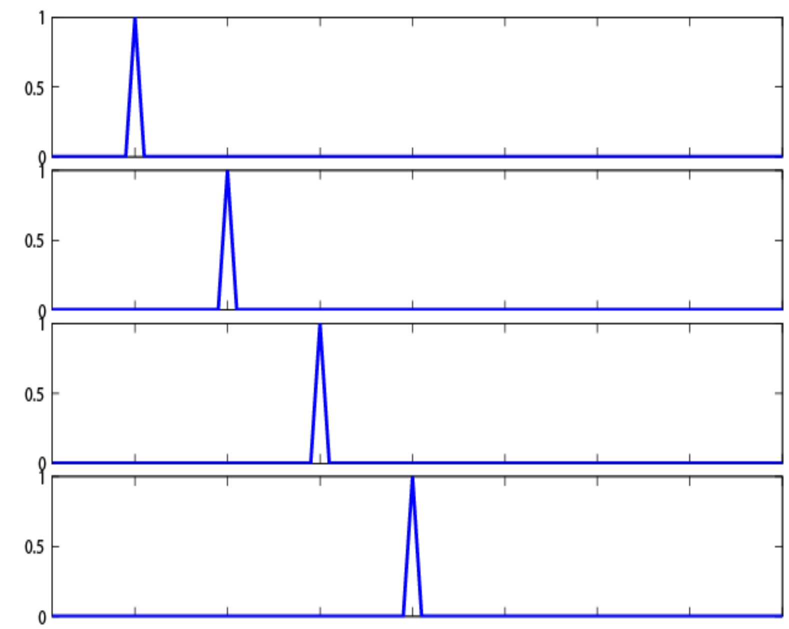

pointwise sampling

Natural signals / images are not smooth.

Uniform Sampling and Smoothness

Uniform sampling: $f[i] = f(i/N) = \langle f, \varphi_i \rangle$ where $\varphi_i = \delta_{i/N}$ Dirac.

Shannon Theorem: f is recovered from $\{\langle f, \varphi_i \rangle\}_{i=0}^{P-1}$ if $\mathrm{Supp}(\hat{f}) \subset [-N\pi, N\pi]$.

Linear reconstruction formula: $$f(t) = \sum_i f[i] h(t - i/N) \quad \text{where} \quad h(t) = \frac{\sin(\pi N t)}{\pi N t}$$

$\rightarrow$ works for smooth signals/images.

$\rightarrow$ discontinuities / edges requires a large P.

$\delta_{1/N}$

$\delta_{2/N}$

$\delta_{3/N}$

$\ldots$

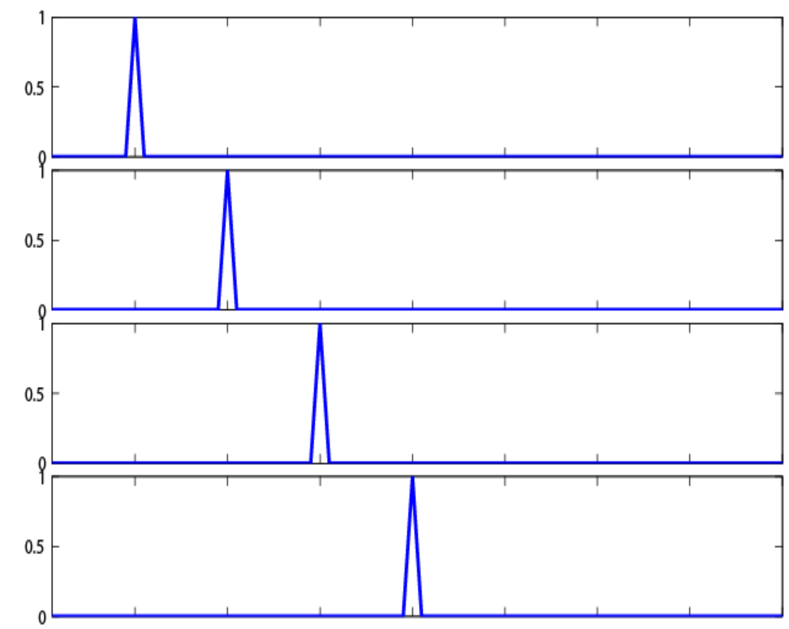

pointwise sampling

Natural signals / images are not smooth.

Possible solution: compressive sensing.

Discrete Ortho-bases

Discrete signal $f \in \mathbb{C}^N$.

1D signals: $f \in \mathbb{C}^N$, $N = \#\text{samples}$.

2D Images: $f \in \mathbb{C}^N \simeq \mathbb{C}^{N_0 \times N_0}$, $N = N_0 \times N_0 = \#\text{pixels}$.

Discrete inner product: $$\langle f,\, g \rangle = \sum_{n=0}^{N-1} f[n] \bar{g}[n]$$

Discrete orthogonal basis $\{\psi_m\}_{0 \leqslant m < N}$ of $\mathbb{C}^N$.

Decomposition: $$f = \sum_{m=0}^{N-1} \langle f,\, \psi_m \rangle \psi_m$$

Energy conservation: $$\|f\|^2 = \sum_{n=0}^{N-1} |f[n]|^2 = \sum_{m=0}^{N-1} |\langle f,\, \psi_m \rangle|^2$$

Discrete Ortho-bases

Discrete signal $f \in \mathbb{C}^N$.

1D signals: $f \in \mathbb{C}^N$, $N = \#\text{samples}$.

2D Images: $f \in \mathbb{C}^N \simeq \mathbb{C}^{N_0 \times N_0}$, $N = N_0 \times N_0 = \#\text{pixels}$.

Discrete inner product: $$\langle f,\, g \rangle = \sum_{n=0}^{N-1} f[n]\bar{g}[n]$$

Discrete orthogonal basis $\{\psi_m\}_{0 \leqslant m < N}$ of $\mathbb{C}^N$.

Decomposition: $$f = \sum_{m=0}^{N-1} \langle f,\, \psi_m \rangle \psi_m$$

Energy conservation: $$\|f\|^2 = \sum_{n=0}^{N-1} |f[n]|^2 = \sum_{m=0}^{N-1} |\langle f,\, \psi_m \rangle|^2$$

Discrete Fourier basis: $$e_m[n] = \frac{1}{\sqrt{N}} e^{\frac{2i\pi}{N} mn} = \frac{1}{\sqrt{N}} e_m(n/N)$$

Discrete Wavelets: implicitely defined (filtering cascade).

Discrete Ortho-bases

Discrete signal $f \in \mathbb{C}^N$.

1D signals: $f \in \mathbb{C}^N$, $N = \#\text{samples}$.

2D Images: $f \in \mathbb{C}^N \simeq \mathbb{C}^{N_0 \times N_0}$, $N = N_0 \times N_0 = \#\text{pixels}$.

Discrete inner product: $\langle f,\, g\rangle = \sum_{n=0}^{N-1} f[n]\bar{g}[n]$

Discrete orthogonal basis $\{\psi_m\}_{0 \leqslant m < N}$ of $\mathbb{C}^N$.

Decomposition: $f = \sum_{m=0}^{N-1} \langle f,\, \psi_m\rangle \psi_m$

Energy conservation: $\|f\|^2 = \sum_{n=0}^{N-1} |f[n]|^2 = \sum_{m=0}^{N-1} |\langle f,\, \psi_m\rangle|^2$

Discrete Fourier basis: $e_m[n] = \frac{1}{\sqrt{N}} e^{\frac{2i\pi}{N} mn} = \frac{1}{\sqrt{N}} e_m(n/N)$

Discrete Wavelets: implicitely defined (filtering cascade).

Computing the $\{\langle f,\, \psi_m\rangle\}_{0 \leqslant m < N}$: brute force $O(N^2)$.

Discrete Fourier: FFT algorithm $O(N\log(N))$.

Discrete Wavelets: FWT algorithm $O(N)$.

Overview

- 1D Orthogonal Decompositions
- Discretization
- 2D Bases
- Linear and Non-linear Approximations
- Decompositions for Processing

2D Extensions of 1D Bases

Image of $N = N_0 \times N_0$ pixel (regular grid).

Separable bases (tensor product):

$$\psi_{m_1,m_2}[n_1,n_2] = \psi_{m_1}[n_1]\psi_{m_2}[n_2]$$

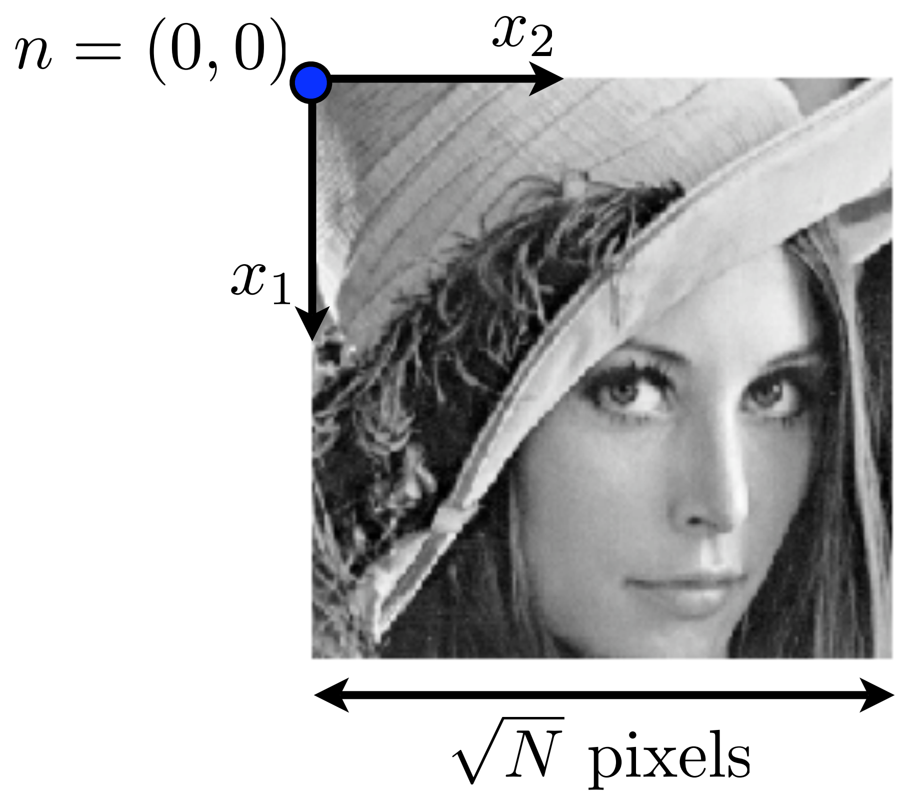

2D Extensions of 1D Bases

Image of $N = N_0 \times N_0$ pixel (regular grid).

Separable bases (tensor product):

$$\psi_{m_1,m_2}[n_1,n_2] = \psi_{m_1}[n_1]\psi_{m_2}[n_2]$$

2D Fourier: $\quad e_{m_1,m_2}[n] = e^{\frac{2i\pi}{N}(m_1 n_1 + m_2 n_2)}$

2D wave orthogonal to (m_1, m_2).

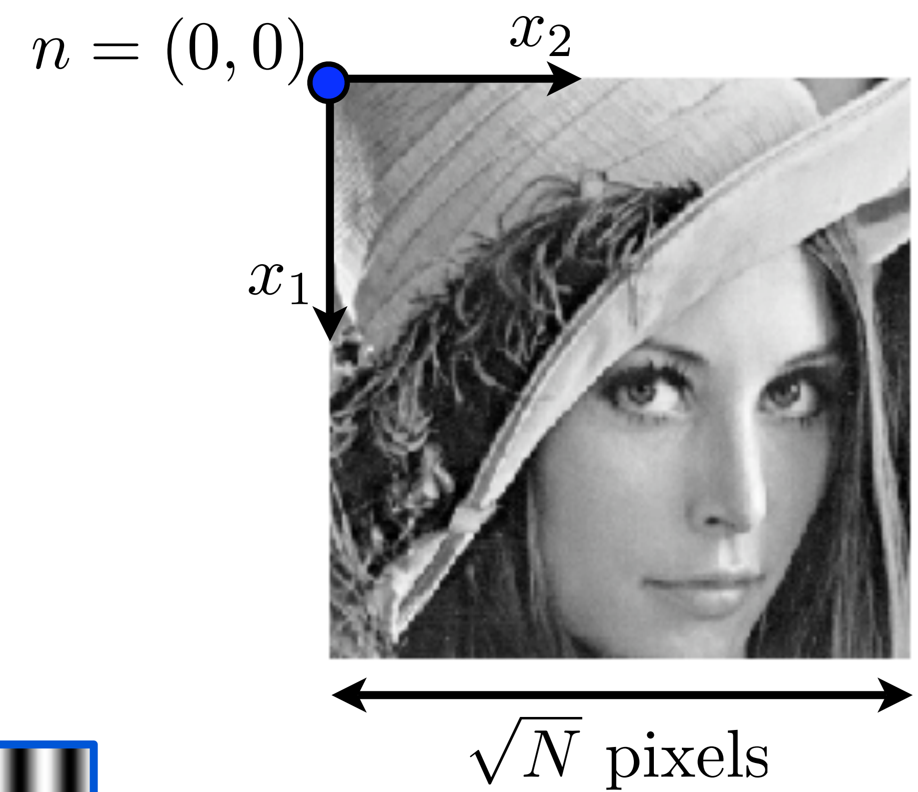

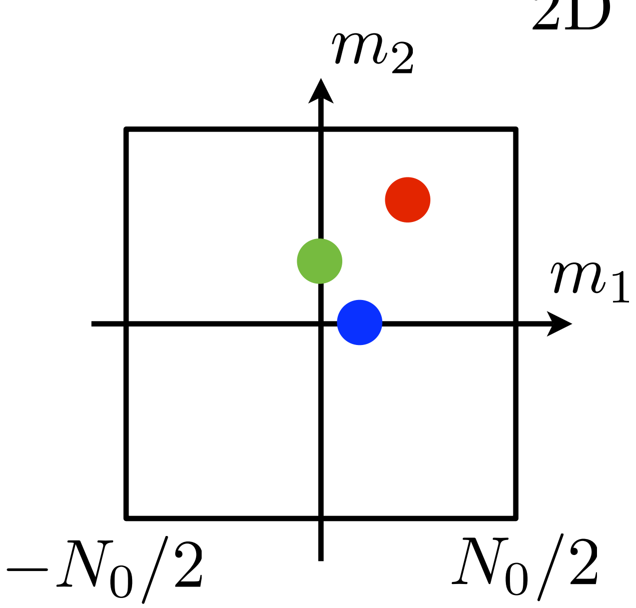

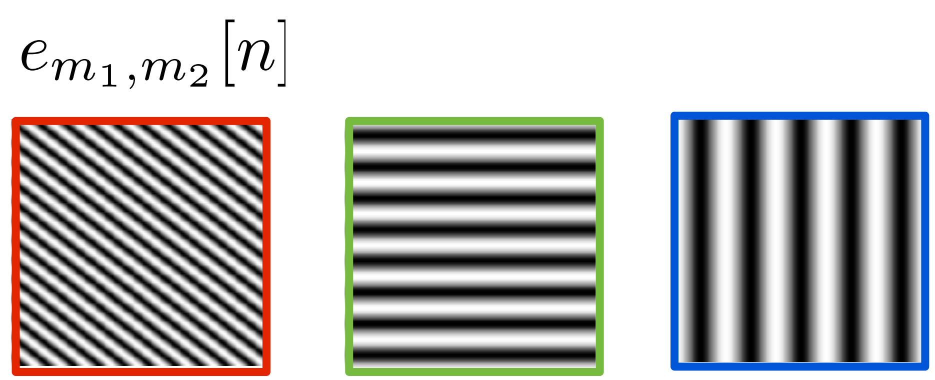

$(m_1, m_2)^\perp =$ orientation.

2D Extensions of 1D Bases

Image of $N = N_0 \times N_0$ pixel (regular grid).

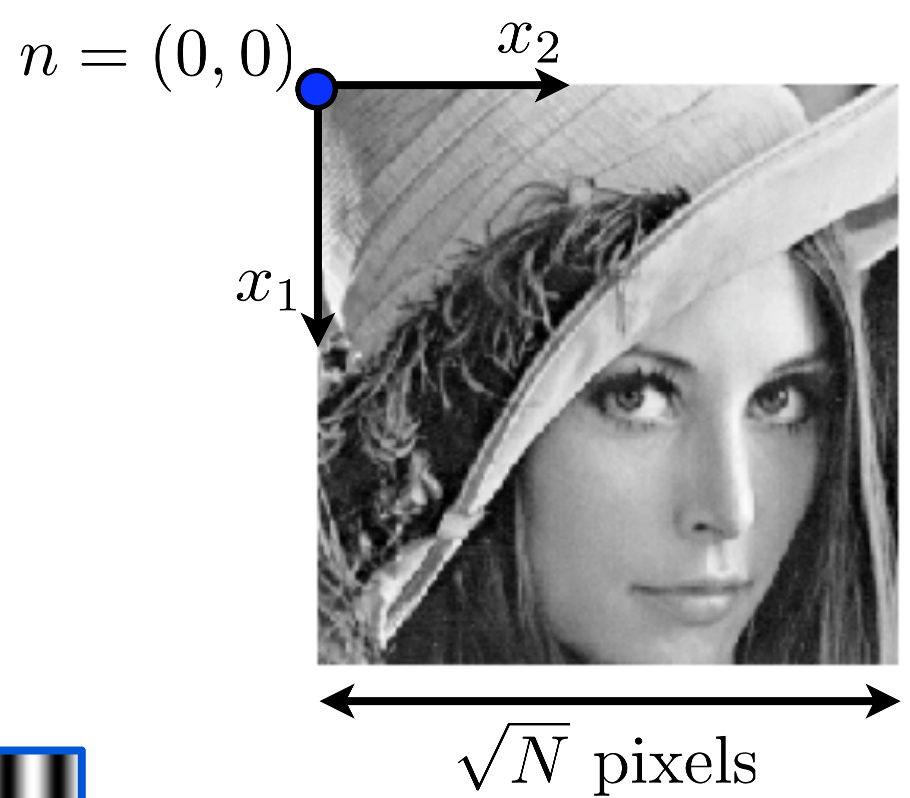

Separable bases (tensor product):

$$\psi_{m_1,m_2}[n_1,n_2] = \psi_{m_1}[n_1]\psi_{m_2}[n_2]$$

2D Fourier: $\mathrm{e}_{m_1,m_2}[n] = e^{\frac{2i\pi}{N}(m_1 n_1 + m_2 n_2)}$

2D wave orthogonal to (m_1, m_2).

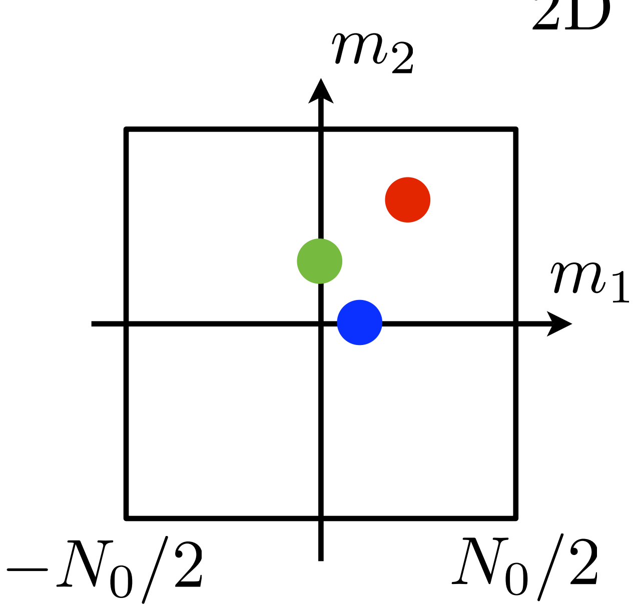

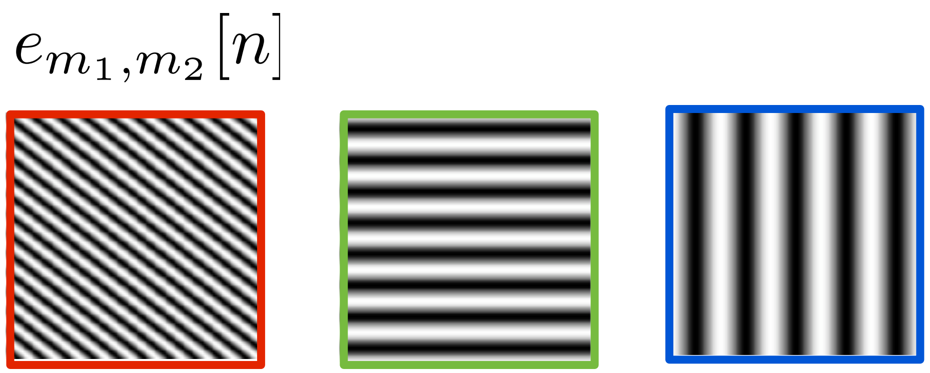

$(m_1, m_2)^\perp =$ orientation.

Anisotropic wavelets:

$$\psi_{j_1,j_2,n_1,n_2}[x_1,x_2] = \psi_{j_1,n_1}[x_1]\psi_{j_2,n_2}[x_2]$$

2D atoms stretched in horizontal/vertical directions.

2D Extensions of 1D Bases

Image of $N = N_0 \times N_0$ pixel (regular grid).

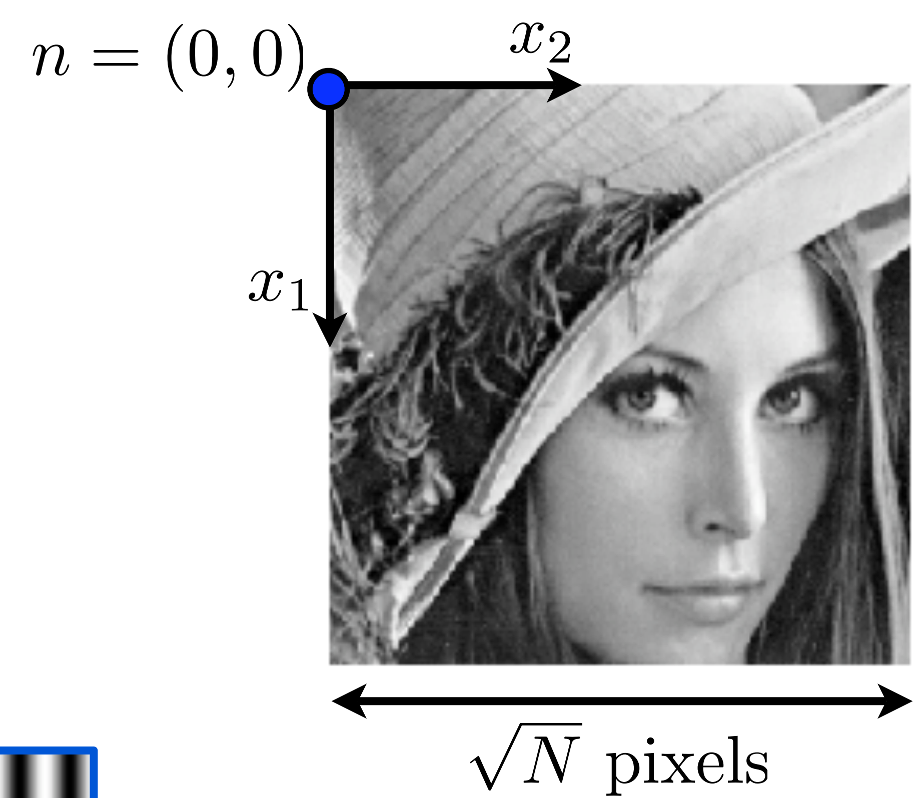

Separable bases (tensor product):

$$\psi_{m_1,m_2}[n_1, n_2] = \psi_{m_1}[n_1]\psi_{m_2}[n_2]$$

2D Fourier: $\quad e_{m_1,m_2}[n] = e^{\frac{2i\pi}{N}(m_1 n_1 + m_2 n_2)}$

2D wave orthogonal to (m_1, m_2).

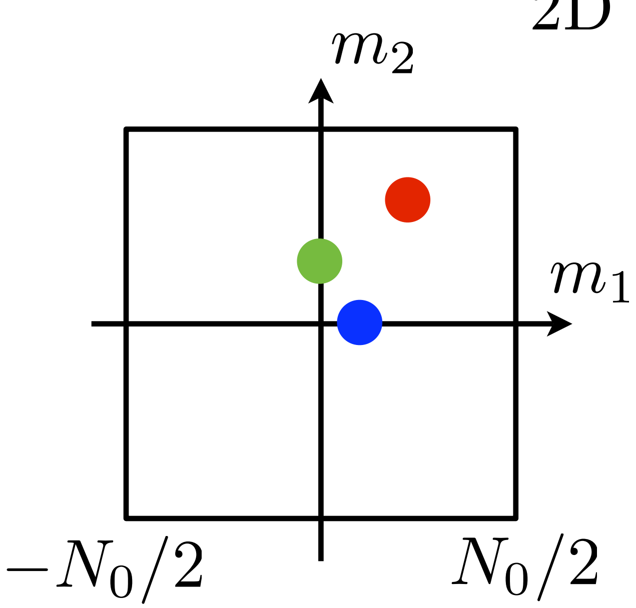

$e_{m_1,m_2}[n]$

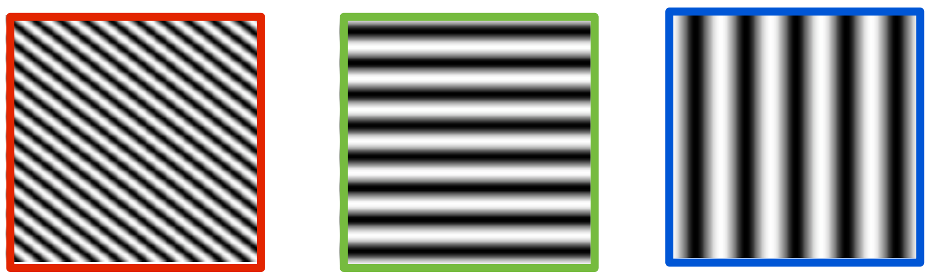

$(m_1, m_2)^\perp$ = orientation.

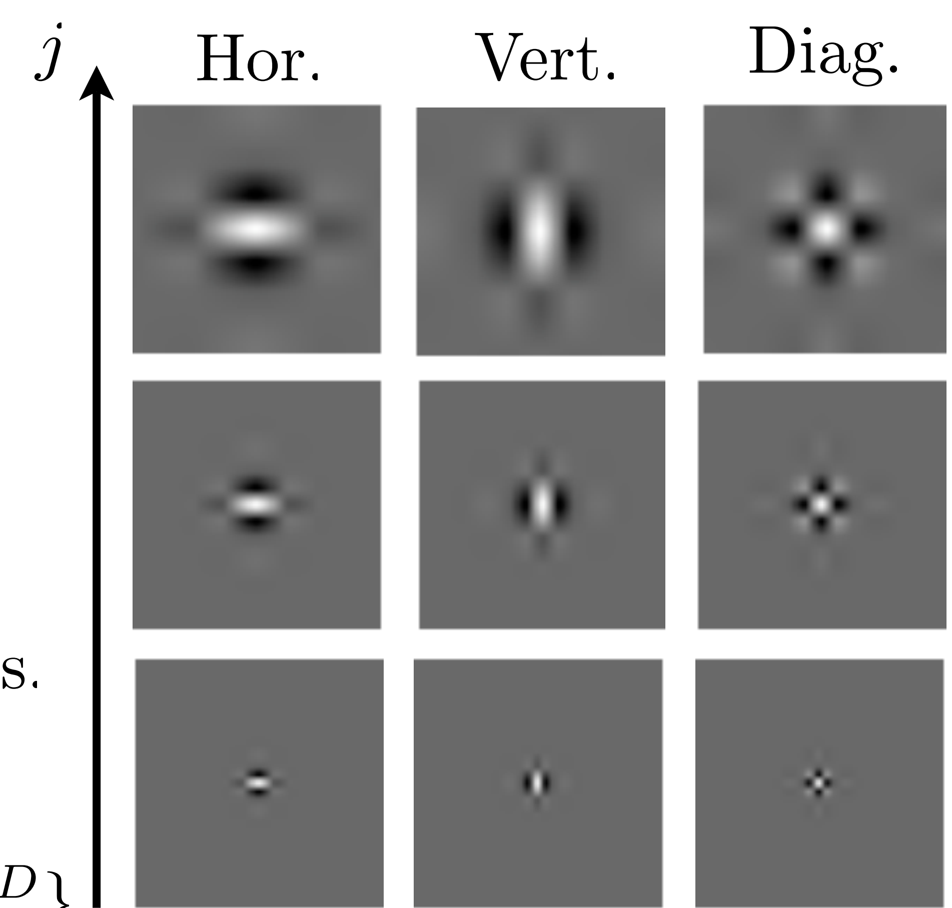

Anisotropic wavelets:

$$\psi_{j_1,j_2,n_1,n_2}[x_1, x_2] = \psi_{j_1,n_1}[x_1]\psi_{j_2,n_2}[x_2]$$

2D atoms stretched in horizontal/vertical directions.

Non-separable bases:

Isotropic wavelets: three mothers wavelets $\{\psi^H, \psi^V, \psi^D\}$

Overview

- 1D Orthogonal Decompositions
- Discretization
- 2D Bases
- Linear and Non-linear Approximations
- Decompositions for Processing

Linear Approximation

Decomposition: $f = \sum_{m\in\mathbb{Z}} \langle f,\, \psi_m \rangle \psi_m$ Error$= 0$

$$f \xrightarrow[\text{transform}]{\text{forward}} \{\langle f,\, \psi_m\rangle\}_{m\in\mathbb{Z}} \xrightarrow[\text{transform}]{\text{backward}} f$$

Image f

Linear Approximation

Decomposition: $$f = \sum_{m \in \mathbb{Z}} \langle f,\, \psi_m \rangle \psi_m$$ Error$= 0$

$$f \xrightarrow[\text{transform}]{\text{forward}} \{\langle f,\, \psi_m \rangle\}_{m \in \mathbb{Z}} \xrightarrow[\text{transform}]{\text{backward}} f$$

Approximation: $$f_M = \sum_{m \in I_M} \langle f,\, \psi_m \rangle \psi_m, \quad M = \# I_M.$$

$$f \xrightarrow[\text{transform}]{\text{forward}} \{\langle f,\, \psi_m \rangle\}_{m \in \mathbb{Z}} \xrightarrow[\text{selection}]{\text{coefficients}} \{\langle f,\, \psi_m \rangle\}_{m \in I_M} \xrightarrow[\text{transform}]{\text{backward}} f_M$$

Approximation error: $$\|f - f_M\|^2 = \sum_{m \notin I_M} |\langle f,\, \psi_m \rangle|^2$$

Image f

Linear Approximation

Decomposition: $$f = \sum_{m \in \mathbb{Z}} \langle f, \psi_m \rangle \psi_m$$ Error$= 0$

$$f \xrightarrow[\text{transform}]{\text{forward}} \{\langle f, \psi_m \rangle\}_{m \in \mathbb{Z}} \xrightarrow[\text{transform}]{\text{backward}} f$$

Approximation: $$f_M = \sum_{m \in I_M} \langle f, \psi_m \rangle \psi_m, \quad M = \# I_M.$$

$$f \xrightarrow[\text{transform}]{\text{forward}} \{\langle f, \psi_m \rangle\}_{m \in \mathbb{Z}} \xrightarrow[\text{selection}]{\text{coefficients}} \{\langle f, \psi_m \rangle\}_{m \in I_M} \xrightarrow[\text{transform}]{\text{backward}} f_M$$

Approximation error: $$\|f - f_M\|^2 = \sum_{m \notin I_M} |\langle f, \psi_m \rangle|^2$$

Linear selection: I_M is fixed. $$(f+g)_M = f_M + g_M$$

Ortho-projection on $V_M = \text{Span}(\psi_m)_{m \in I_M}$: $\quad f_M = \text{Proj}_{V_M}(f)$

Fourier low-frequency approximation: $I_M = \{-M/2+1, \ldots, M/2\}$.

Wavelet coase-scale approximation: $I_M = \{m = (j,m) \setminus j \geqslant j_0\}$.

Image f

Linear approximation
$M = N/16$

Non-Linear Approximation

Approximation: $$f_M = \sum_{m \in I_M} \langle f, \psi_m \rangle \psi_m, \quad M = \# I_M.$$

Non-linear approximation: I_M depends on f.

Non-Linear Approximation

Approximation: $$f_M = \sum_{m \in I_M} \langle f,\, \psi_m \rangle \psi_m, \quad M = \# I_M.$$

Non-linear approximation: I_M depends on f.

Optimal L^2 choice: minimize $\|f - f_M\|^2 \quad \Longrightarrow \quad I_M = \underset{|I| \leqslant M}{\operatorname{argmin}} \sum_{m \notin I} |\langle f,\, \psi_m \rangle|^2$

Closed form solution: $$I_M = \{M \text{ largest coefficients } |\langle f,\, \psi_m \rangle|\}.$$

Image f

Linear approximation

Non-linear approximation

Non-Linear Approximation

Approximation: $$f_M = \sum_{m \in I_M} \langle f,\, \psi_m \rangle \psi_m, \quad M = \# I_M.$$

Non-linear approximation: I_M depends on f.

Optimal L^2 choice: minimize $\|f - f_M\|^2 \quad \Longrightarrow \quad I_M = \operatorname*{argmin}_{|I| \leqslant M} \sum_{m \notin I} |\langle f,\, \psi_m \rangle|^2$

Closed form solution: $I_M = \{M \text{ largest coefficients } |\langle f,\, \psi_m \rangle|\}.$

Computation with thresholding: $I_M = \{m \setminus |\langle f,\, \psi_m \rangle| > T\}$

$$f_M = \sum_{|\langle f, \psi_m \rangle| > T} \langle f,\, \psi_m \rangle \psi_m$$

Image f

Linear approximation

Non-linear approximation

Non-Linear Approximation

Approximation: $\displaystyle f_M = \sum_{m \in I_M} \langle f,\, \psi_m \rangle \psi_m, \quad M = \# I_M.$

Non-linear approximation: I_M depends on f.

Optimal L^2 choice: minimize $\|f - f_M\|^2 \quad \Longrightarrow \quad \displaystyle I_M = \underset{|I| \leqslant M}{\operatorname{argmin}} \sum_{m \notin I} |\langle f,\, \psi_m \rangle|^2$

Closed form solution: $I_M = \{M \text{ largest coefficients } |\langle f,\, \psi_m \rangle|\}.$

Computation with thresholding: $I_M = \{m \setminus |\langle f,\, \psi_m \rangle| > T\}$

$$f_M = \sum_{|\langle f,\, \psi_m \rangle| > T} \langle f,\, \psi_m \rangle \psi_m$$

1:1 mapping $M \leftrightarrow T$

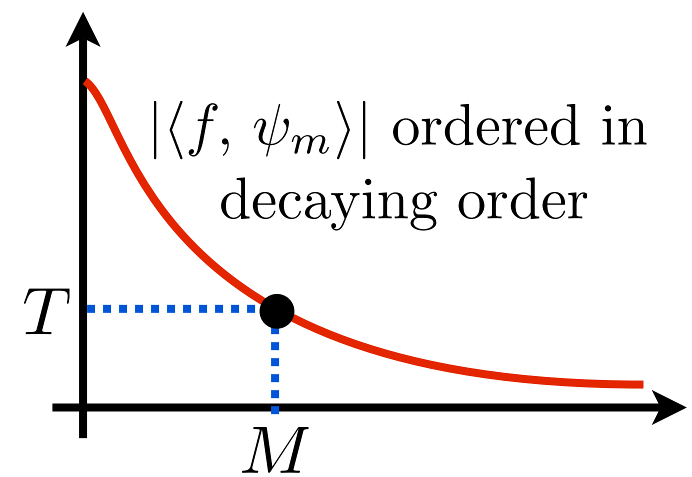

Image f

Linear approximation

Non-linear approximation

Efficient Approximation

Signal prior: $f \in \Theta \subset \mathrm{L}^2([0,1]^d)$.

Efficient Approximation

Signal prior: $f \in \Theta \subset \mathrm{L}^2([0,1]^d)$.

Smooth: bounded derivatives $\|f\|_{\mathrm{C}^\alpha} \leqslant C$.

Efficient Approximation

Signal prior: $f \in \Theta \subset \mathrm{L}^2([0,1]^d)$.

Smooth: bounded derivatives $\|f\|_{\mathrm{C}^\alpha} \leqslant C$.

1D piecewise smooth: less than K discontinuities.

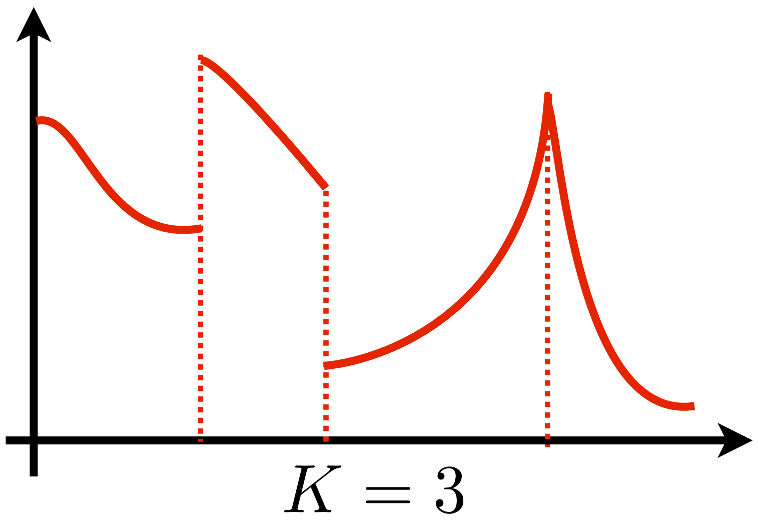

Efficient Approximation

Signal prior: $f \in \Theta \subset \mathrm{L}^2([0,1]^d)$.

Smooth: bounded derivatives $\|f\|_{\mathrm{C}^\alpha} \leqslant C$.

1D piecewise smooth: less than K discontinuities.

Bounded variations: $\|f\|_{TV}$ "=" $\int \|\nabla f\| \leqslant C$.

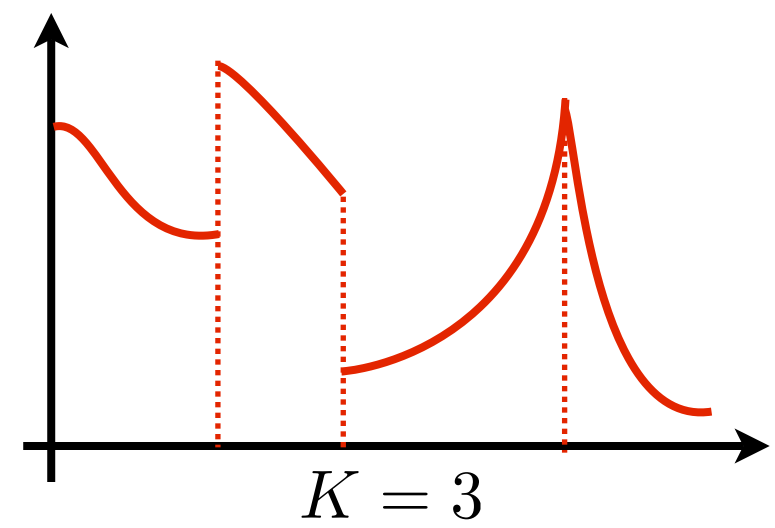

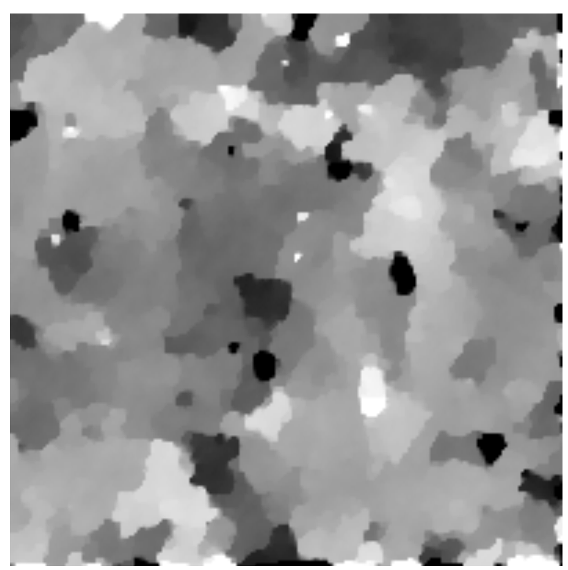

Efficient Approximation

Signal prior: $f \in \Theta \subset \mathrm{L}^2([0,1]^d)$.

Smooth: bounded derivatives $\|f\|_{\mathrm{C}^\alpha} \leqslant C$.

1D piecewise smooth: less than K discontinuities.

Bounded variations: $\|f\|_{TV}$ "=" $\int \|\nabla f\| \leqslant C$.

2D cartoon: smooth discontinuity curves.

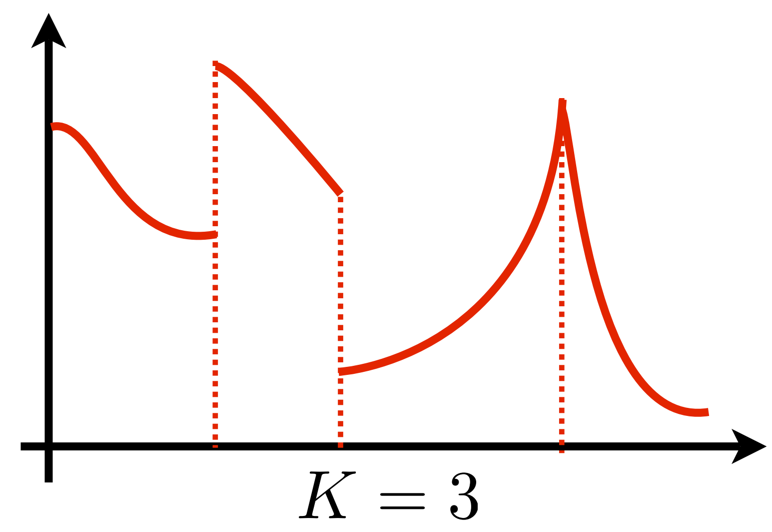

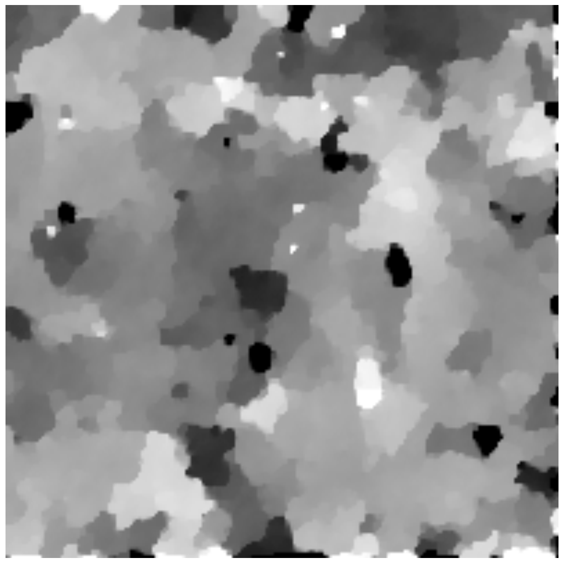

Efficient Approximation

Signal prior: $f \in \Theta \subset \mathrm{L}^2([0,1]^d)$.

Smooth: bounded derivatives $\|f\|_{C^\alpha} \leqslant C$.

1D piecewise smooth: less than K discontinuities.

Bounded variations: $\|f\|_{TV}$ " $=$ " $\int \|\nabla f\| \leqslant C$.

2D cartoon: smooth discontinuity curves.

Natural images: ??

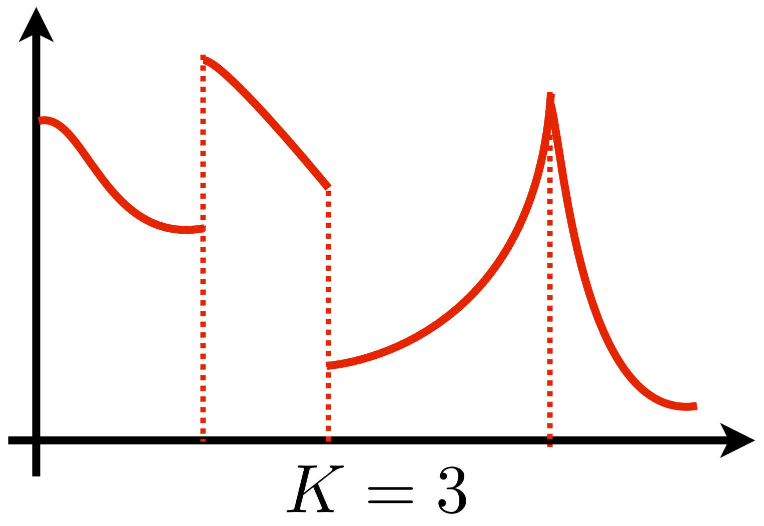

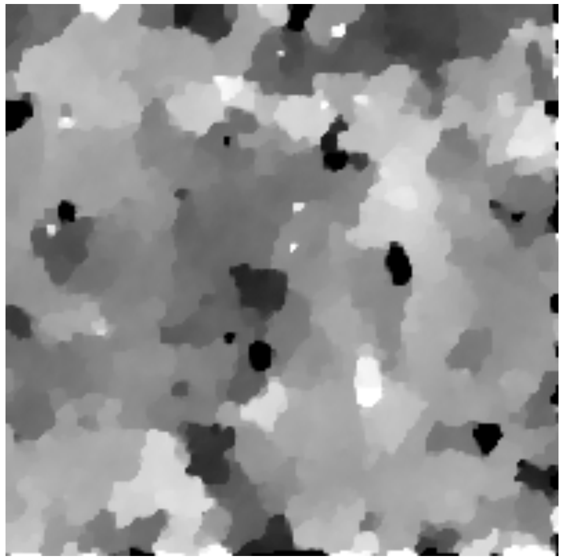

Efficient Approximation

Signal prior: $f \in \Theta \subset \mathrm{L}^2([0,1]^d)$.

Smooth: bounded derivatives $\|f\|_{\mathrm{C}^\alpha} \leqslant C$.

1D piecewise smooth: less than K discontinuities.

Bounded variations: $\|f\|_{TV}$ " = " $\int \|\nabla f\| \leqslant C$.

2D cartoon: smooth discontinuity curves.

Natural images: ??

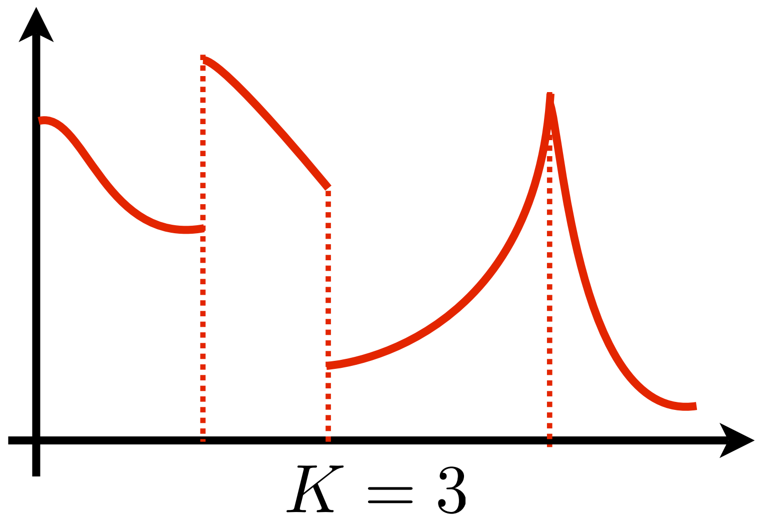

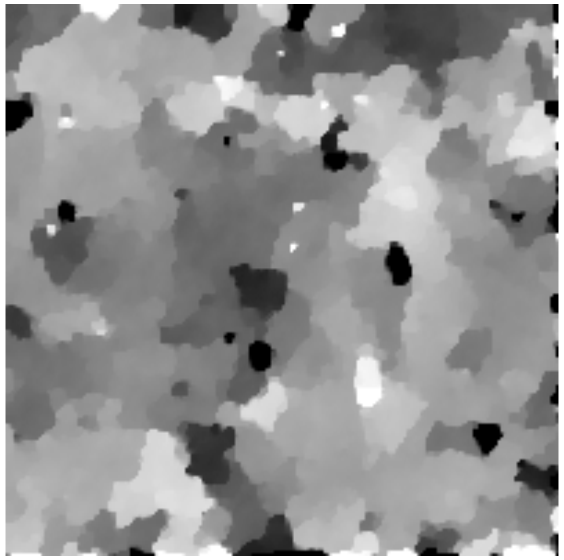

Efficient representation: $\|f - f_M\|$ decay fast with M for all $f \in \Theta$.

$$\|f - f_M\|^2 \leqslant C_f M^{-\alpha} \quad \text{for a large } \alpha.$$

Efficient Approximation

Signal prior: $f \in \Theta \subset \mathrm{L}^2([0,1]^d)$.

Smooth: bounded derivatives $\|f\|_{C^\alpha} \leqslant C$.

1D piecewise smooth: less than K discontinuities.

Bounded variations: $\|f\|_{TV}$ "$=$" $\int \|\nabla f\| \leqslant C$.

2D cartoon: smooth discontinuity curves.

Natural images: ??

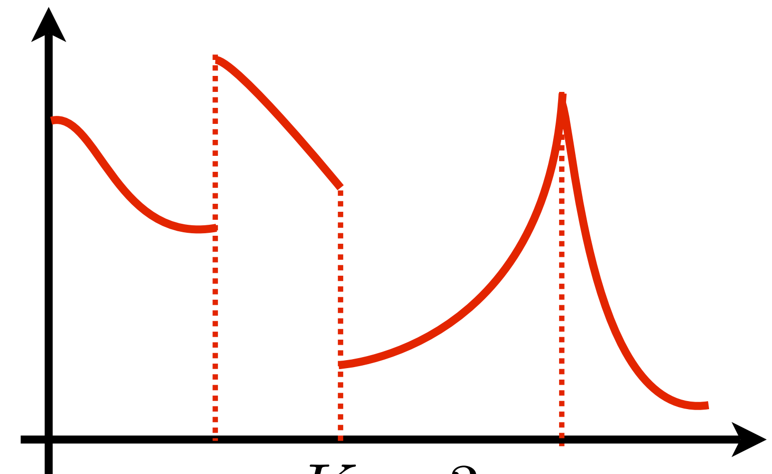

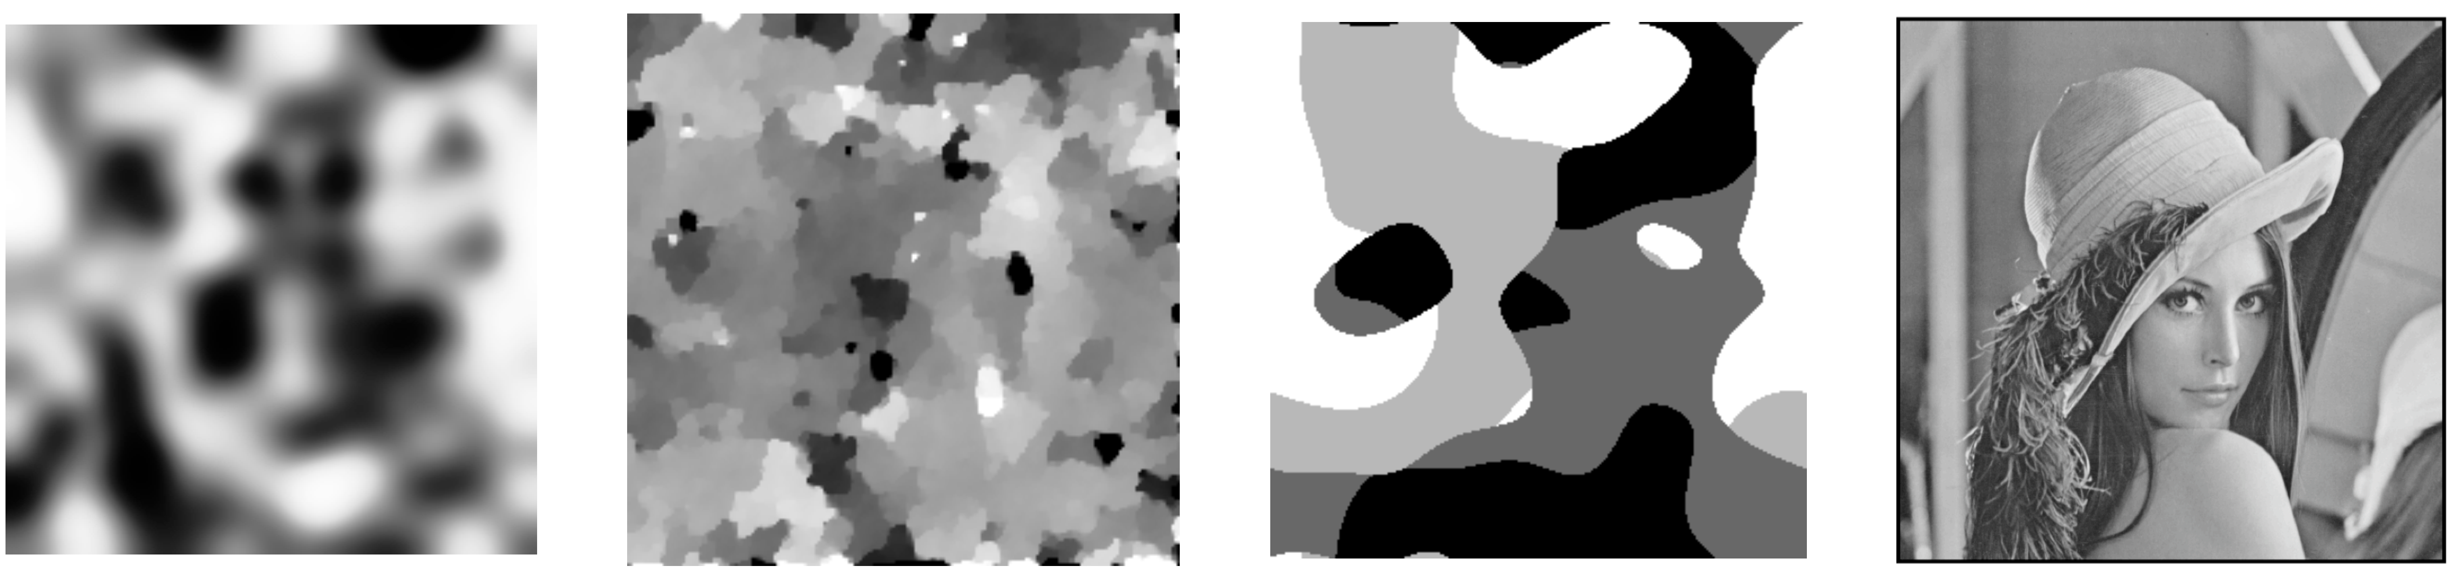

Efficient representation: $\|f - f_M\|$ decay fast with M for all $f \in \Theta$.

$$\|f - f_M\|^2 \leqslant C_f M^{-\alpha} \quad \text{for a large } \alpha.$$

Goal: design efficient basis $\{\psi_m\}_m$ to maximize α.

Overview

- 1D Orthogonal Decompositions
- Discretization
- 2D Bases
- Linear and Non-linear Approximations
- **Decompositions for Processing**

Compression by Transform-coding

$$f \xrightarrow[\text{transform}]{\text{forward}} a[m] = \langle f,\, \psi_m \rangle \in \mathbb{R}$$

Image f

Zoom on f

Compression by Transform-coding

$$f \xrightarrow[\text{transform}]{\text{forward}} a[m] = \langle f,\, \psi_m \rangle \in \mathbb{R} \xrightarrow[\text{bin } T]{\text{quantization}} q[m] \in \mathbb{Z}$$

Quantization: $$q[m] = \text{sign}(a[m]) \left\lfloor \frac{|a[m]|}{T} \right\rfloor \in \mathbb{Z}$$

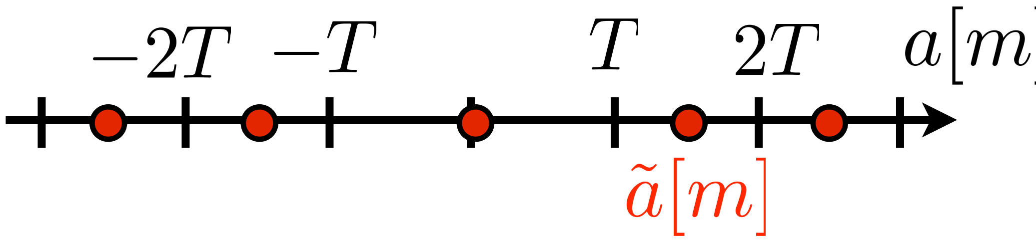

Image f

Zoom on f

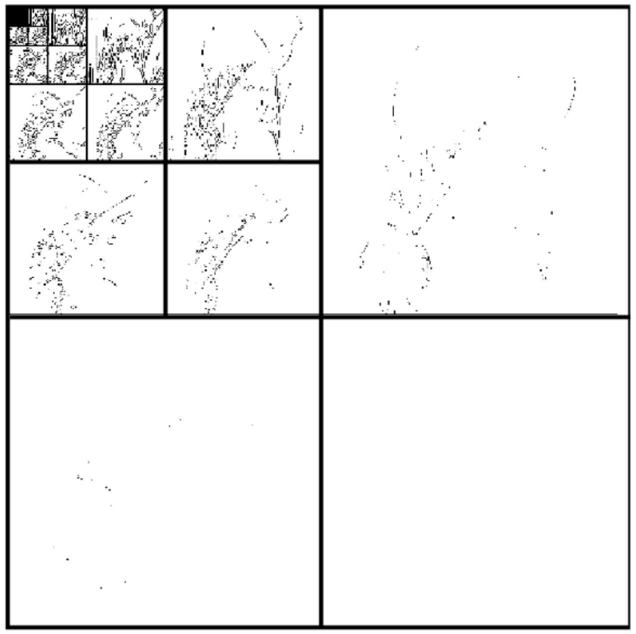

Quantized $q[m]$

Compression by Transform-coding

$$f \xrightarrow[\text{transform}]{\text{forward}} a[m] = \langle f,\, \psi_m \rangle \in \mathbb{R} \xrightarrow[\text{bin } T]{\text{quantization}} q[m] \in \mathbb{Z} \xrightarrow{\text{coding}} \underbrace{\{\, 0,1,0,0,1,\ldots \}}_{R \text{ bits}}$$

Quantization: $\quad q[m] = \mathrm{sign}(a[m]) \left\lfloor \dfrac{|a[m]|}{T} \right\rfloor \in \mathbb{Z}$

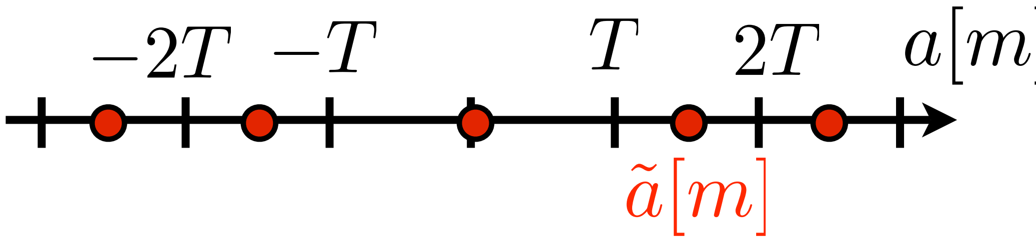

Entropic coding: use statistical redundancy (many 0's).

Image f

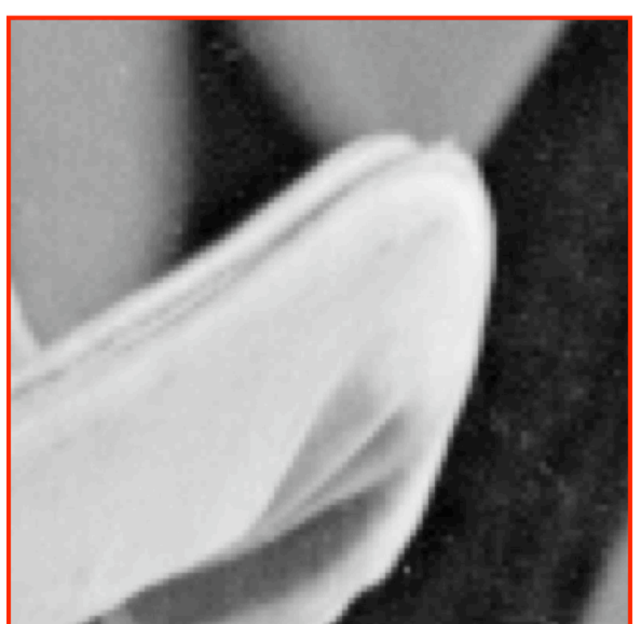

Zoom on f

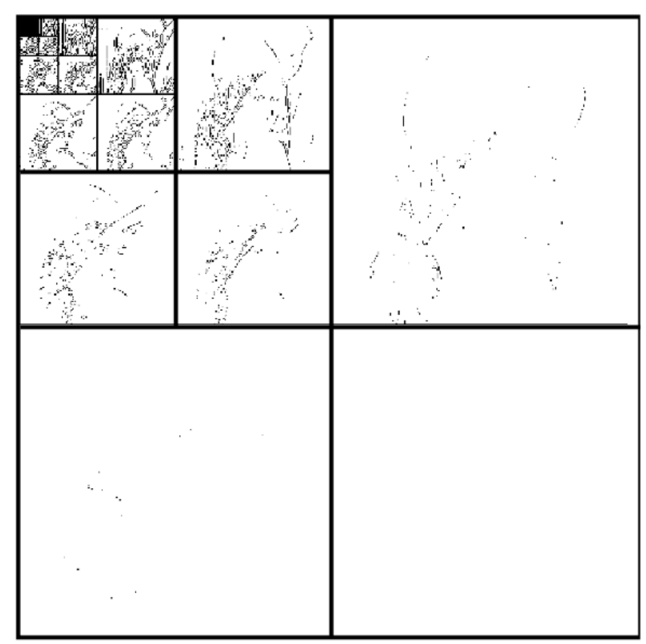

Quantized $q[m]$

Compression by Transform-coding

$$f \xrightarrow[\text{transform}]{\text{forward}} a[m] = \langle f,\, \psi_m \rangle \in \mathbb{R} \xrightarrow[\text{bin } T]{\text{quantization}} q[m] \in \mathbb{Z} \xrightarrow{\text{coding}} \underbrace{\{\,0,1,0,0,1,\ldots\}}_{R \text{ bits}} \xrightarrow{\text{decoding}} q[m] \in \mathbb{Z}$$

Quantization: $q[m] = \mathrm{sign}(a[m]) \left\lfloor \frac{|a[m]|}{T} \right\rfloor \in \mathbb{Z}$

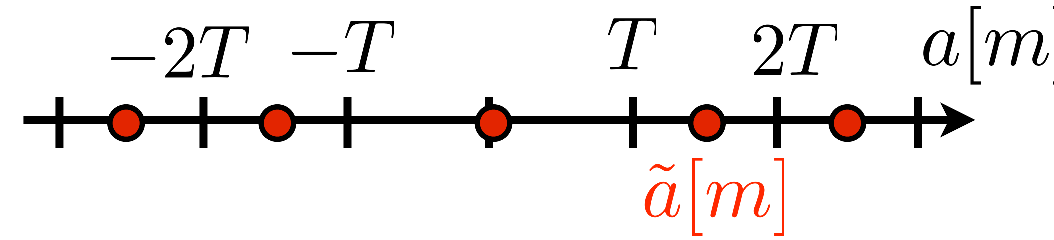

Entropic coding: use statistical redundancy (many 0's).

Image f

Zoom on f

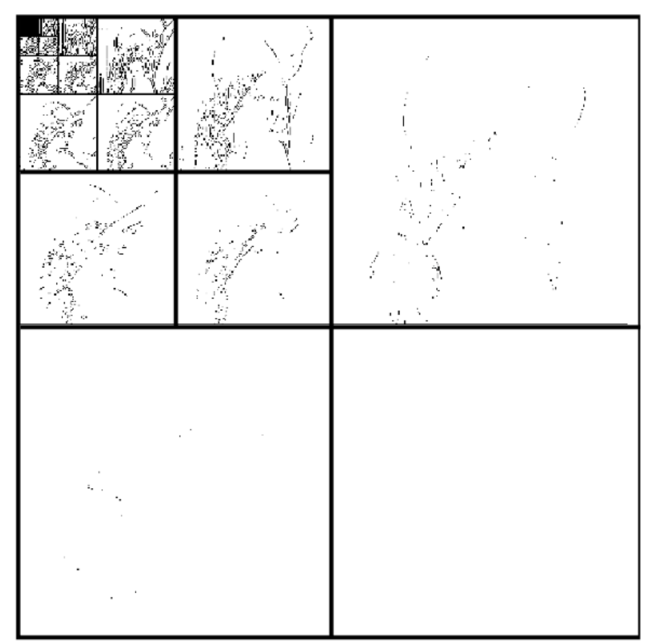

Quantized $q[m]$

Compression by Transform-coding

$$f \xrightarrow[\text{transform}]{\text{forward}} a[m] = \langle f,\, \psi_m \rangle \in \mathbb{R} \xrightarrow[\text{bin } T]{\text{quantization}} q[m] \in \mathbb{Z} \xrightarrow{\text{coding}} \{\, 0,1,0,0,1,\ldots \}$$

$\longleftrightarrow$ R bits

$$\{\, 0,1,0,0,1,\ldots \} \xrightarrow{\text{decoding}} q[m] \in \mathbb{Z} \xrightarrow{\text{dequantization}} \tilde{a}[m]$$

Quantization: $q[m] = \text{sign}(a[m]) \left\lfloor \dfrac{|a[m]|}{T} \right\rfloor \in \mathbb{Z}$

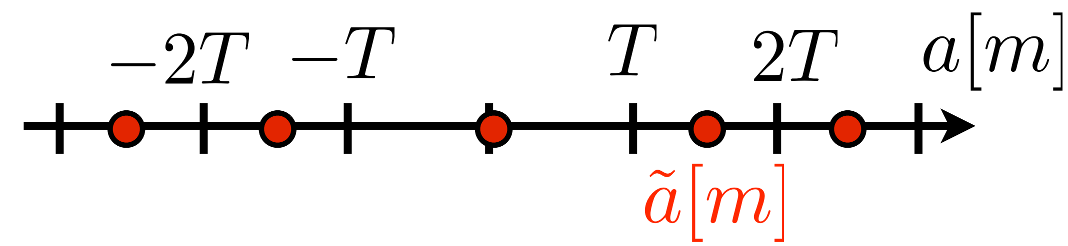

Entropic coding: use statistical redundancy (many 0's).

Dequantization: $\tilde{a}[m] = \text{sign}(q[m]) \left(|q[m] + \dfrac{1}{2} \right) T$

Image f

Zoom on f

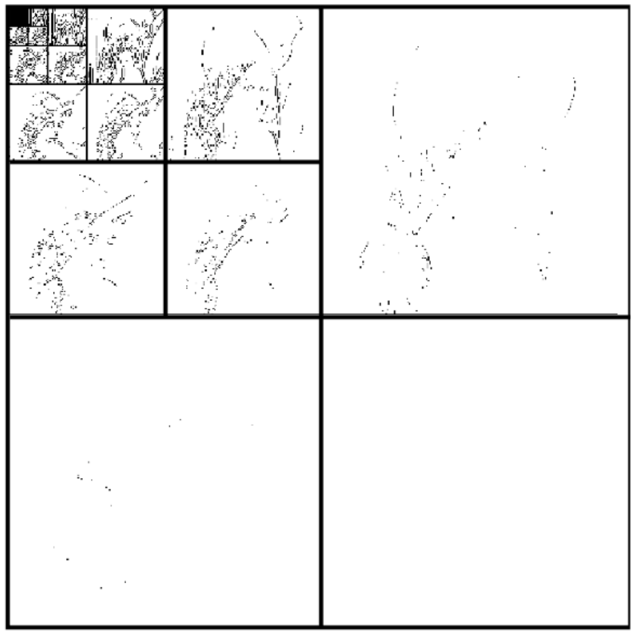

Quantized $q[m]$

Compression by Transform-coding

$$f \xrightarrow[\text{transform}]{\text{forward}} a[m] = \langle f,\, \psi_m \rangle \in \mathbb{R} \xrightarrow[\text{bin } T]{\text{quantization}} q[m] \in \mathbb{Z} \xrightarrow{\text{coding}} \{\ 0,1,0,0,1,\ldots\ \}$$

$\longleftrightarrow$ R bits

$$f_R = \sum_{m \in I_T} \tilde{a}[m]\psi_m \xleftarrow[\text{transform}]{\text{backward}} \tilde{a}[m] \xleftarrow{\text{dequantization}} q[m] \in \mathbb{Z} \xleftarrow{\text{decoding}} \{\ 0,1,0,0,1,\ldots\ \}$$

Quantization: $\quad q[m] = \mathrm{sign}(a[m]) \left\lfloor \dfrac{|a[m]|}{T} \right\rfloor \in \mathbb{Z}$

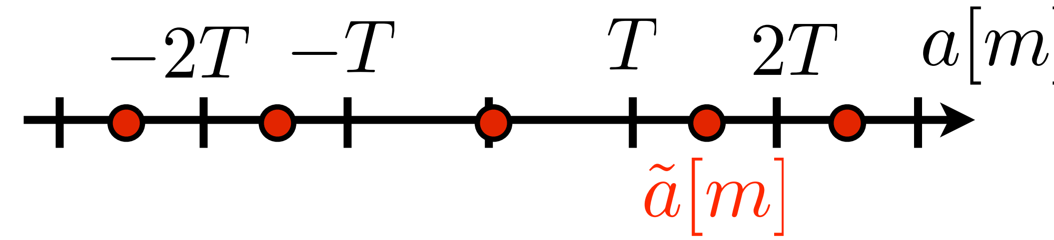

Entropic coding: use statistical redundancy (many 0's).

Dequantization: $\quad \tilde{a}[m] = \mathrm{sign}(q[m]) \left(|q[m] + \dfrac{1}{2} \right) T$

Fast decay of $\|f - f_M\|$ with M. $\iff$ Fast decay of $\|f - f_R\|$ with R.

Image f

Zoom on f

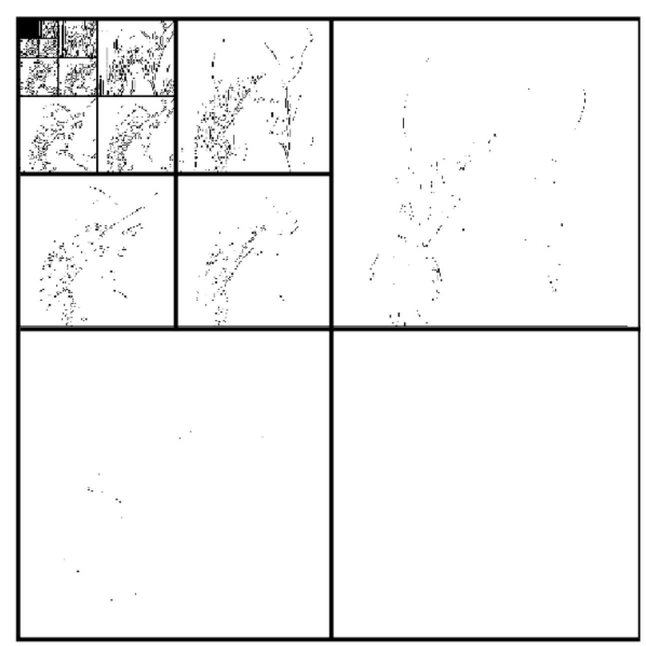

Quantized $q[m]$

f_R, R =0.2 bit/pixel

Denoising

Noisy image $f = f_0 + w$, $w \sim \mathcal{N}(0, \sigma)$ white noise.

Denoised: $\tilde{f}$ depends only on f.

Clean f_0

Noisy $f = f_0 + w$

Denoising

Noisy image $f = f_0 + w$, $w \sim \mathcal{N}(0,\sigma)$ white noise.

Denoised: $\tilde{f}$ depends only on f.

Denoising by approximation:

$$f = \sum_{m=0}^{N-1} \langle f,\, \psi_m \rangle \psi_m \xrightarrow{\text{denoising} \,=\, \text{approximation}} \tilde{f} = \sum_{m \in I} \langle f,\, \psi_m \rangle \psi_m$$

Linear/non linear denoising: choice of I (depends or not on f).

Clean f_0

Noisy $f = f_0 + w$

Denoised $\tilde{f}$

Denoising

Noisy image $f = f_0 + w$, $w \sim \mathcal{N}(0, \sigma)$ white noise.

Denoised: $\tilde{f}$ depends only on f.

Denoising by approximation:

$$f = \sum_{m=0}^{N-1} \langle f,\, \psi_m \rangle \psi_m \xrightarrow{\text{denoising = approximation}} \tilde{f} = \sum_{m \in I} \langle f,\, \psi_m \rangle \psi_m$$

Linear/non linear denoising: choice of I (depends or not on f).

Denoising performance: $\mathbb{E}_w(\|f_0 - \tilde{f}\|)$ function of σ.

Fast decay of $\|f_0 - f_{0,M}\|$ with M. $\iff$ Fast decay of $\mathbb{E}_w(\|f_0 - \tilde{f}\|)$ with σ.

Clean f_0

Noisy $f = f_0 + w$

Denoised $\tilde{f}$

Inverse Problems

Recovering f_0 from P noisy measurements $y = \Phi f_0 + w \in \mathbb{R}^P$.

$\Phi : \mathbb{R}^N \mapsto \mathbb{R}^P$ with $P \ll N$ (missing information)

$w[n] \sim \mathcal{N}(0, \sigma)$ white noise.

Inverse Problems

Recovering f_0 from P noisy measurements $y = \Phi f_0 + w \in \mathbb{R}^P$.

$\Phi : \mathbb{R}^N \mapsto \mathbb{R}^P$ with $P \ll N$ (missing information)

$w[n] \sim \mathcal{N}(0, \sigma)$ white noise.

Denoising: $\Phi = \mathrm{Id}_N$, $P = N$.

Inverse Problems

Recovering f_0 from P noisy measurements $y = \Phi f_0 + w \in \mathbb{R}^P$.

$\Phi : \mathbb{R}^N \mapsto \mathbb{R}^P$ with $P \ll N$ (missing information)

$w[n] \sim \mathcal{N}(0, \sigma)$ white noise.

Denoising: $\Phi = \mathrm{Id}_N$, $P = N$.

Inpainting: set $\Omega \subset \{0, \ldots, N-1\}$ of missing pixels, $P = N - |\Omega|$.

$\xrightarrow{\Phi}$

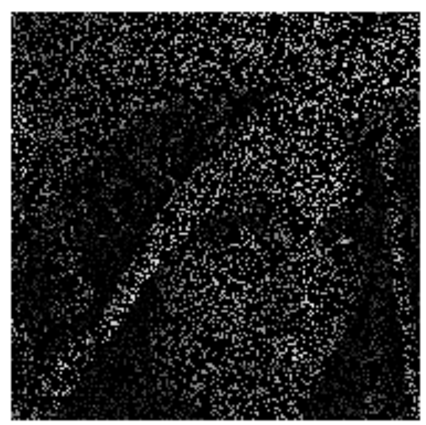

$$(\Phi f)(x) = \begin{cases} 0 & \text{if} \quad x \in \Omega, \\ f(x) & \text{if} \quad x \notin \Omega. \end{cases}$$

Inverse Problems

Recovering f_0 from P noisy measurements $y = \Phi f_0 + w \in \mathbb{R}^P$.

$\Phi : \mathbb{R}^N \mapsto \mathbb{R}^P$ with $P \ll N$ (missing information)

$w[n] \sim \mathcal{N}(0, \sigma)$ white noise.

Denoising: $\Phi = \mathrm{Id}_N$, $P = N$.

Inpainting: set $\Omega \subset \{0, \ldots, N-1\}$ of missing pixels, $P = N - |\Omega|$.

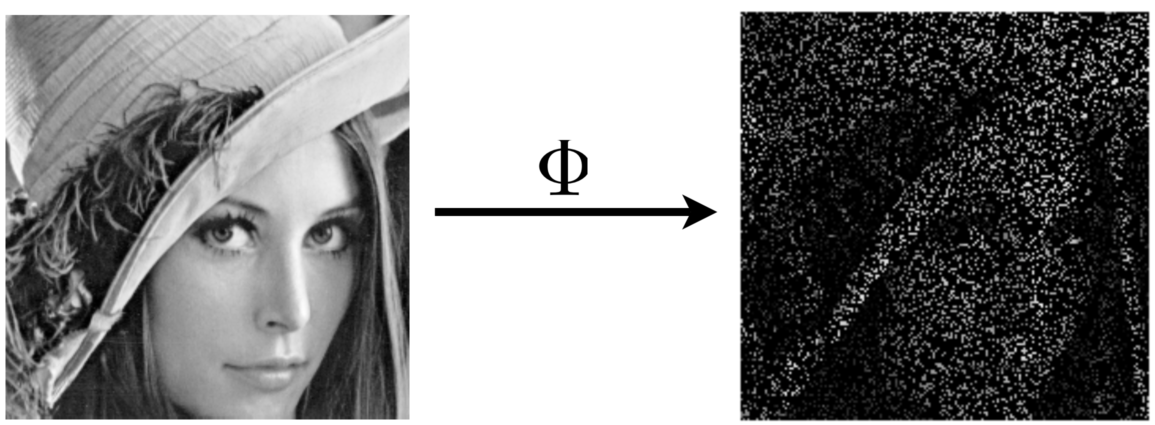

$$(\Phi f)(x) = \begin{cases} 0 & \text{if} \quad x \in \Omega, \\ f(x) & \text{if} \quad x \notin \Omega. \end{cases}$$

Super-resolution: $\Phi f = (f * \varphi) \downarrow_k$, $P = N/k$.

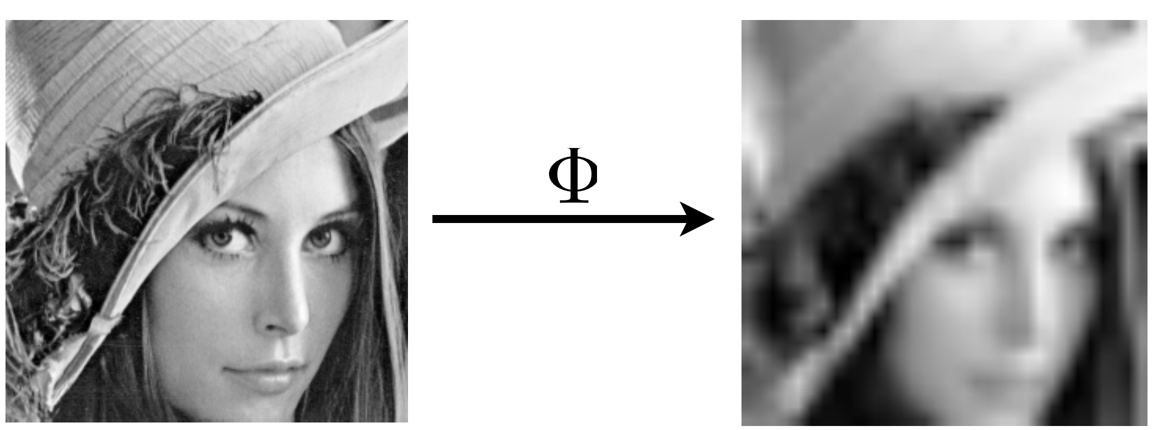

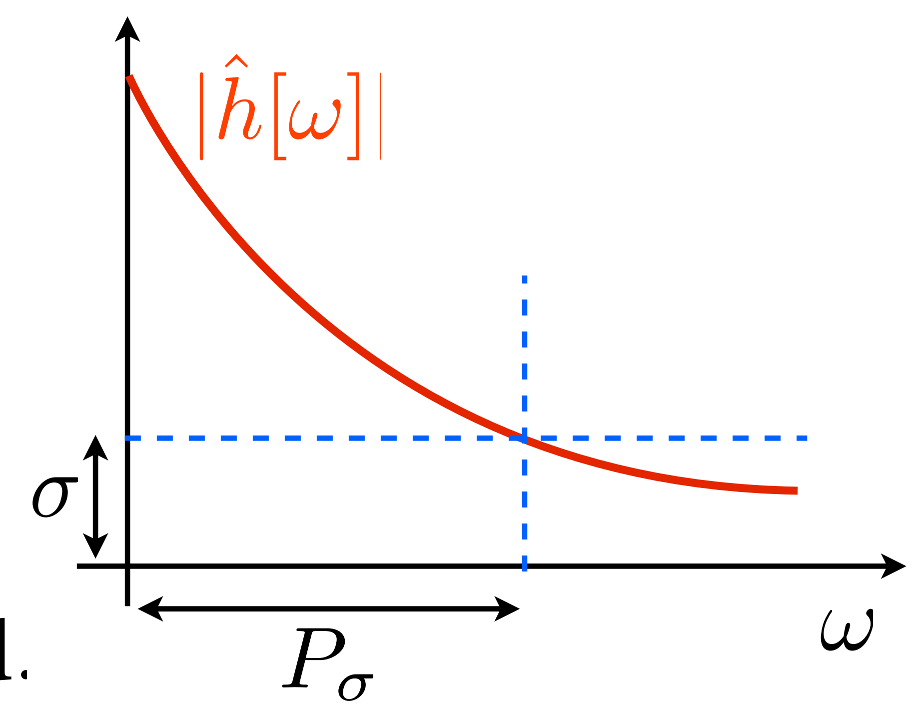

De-blurring: $\Phi f = f \star \varphi$, $P = N$ but ill-posed.

Restoration with Sparsity

Measurements: $y = \Phi f_0 + w$

Compute $\tilde{f}$ such that:

- *Fit measures:* $\Phi \tilde{f} \approx y$
- *Sparsity*: only few $\{\langle \tilde{f}, \psi_m \rangle\}_m$ are large.

Restoration with Sparsity

Measurements: $y = \Phi f_0 + w$

Compute $\tilde{f}$ such that:

- *Fit measures:* $\Phi \tilde{f} \approx y$
- *Sparsity*: only few $\{\langle \tilde{f}, \psi_m \rangle\}_m$ are large.

Performance measure:

- super-resolution error $\|f_0 - \tilde{f}\|$.
- visual quality.

Restoration with Sparsity

Measurements: $y = \Phi f_0 + w$

Compute $\tilde{f}$ such that:

- *Fit measures:* $\Phi \tilde{f} \approx y$
- *Sparsity*: only few $\{\langle \tilde{f},\, \psi_m \rangle\}_m$ are large.

Performance measure:

- super-resolution error $\|f_0 - \tilde{f}\|$.
- visual quality.

Efficient restoration:

- Fast decay of $\|f_0 - f_M\|$ with M.
- $\|\Phi \psi_m\|$ not too small if $\langle f_0,\, \psi_m \rangle$ is large.

Restoration with Sparsity

Measurements: $y = \Phi f_0 + w$

Compute $\tilde{f}$ such that:

- *Fit measures:* $\Phi \tilde{f} \approx y$
- *Sparsity*: only few $\{\langle \tilde{f}, \psi_m \rangle\}_m$ are large.

Performance measure:

- super-resolution error $\|f_0 - \tilde{f}\|$.
- visual quality.

Efficient restoration:

- Fast decay of $\|f_0 - f_M\|$ with M.
- $\|\Phi \psi_m\|$ not too small if $\langle f_0, \psi_m \rangle$ is large.

Image f

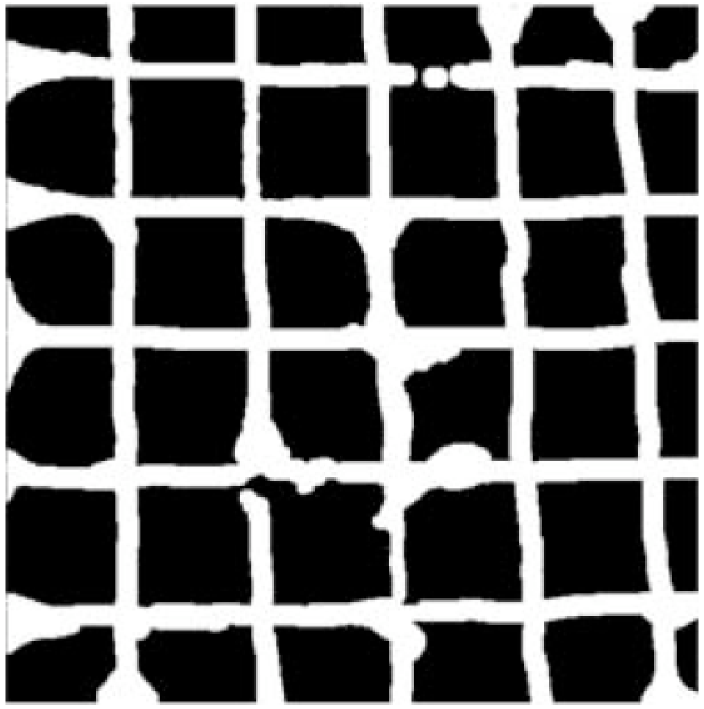

Mask Ω

Inpainted $\tilde{f}$

Conclusion

Orthogonal decompositions:

Fourier, wavelets, etc.

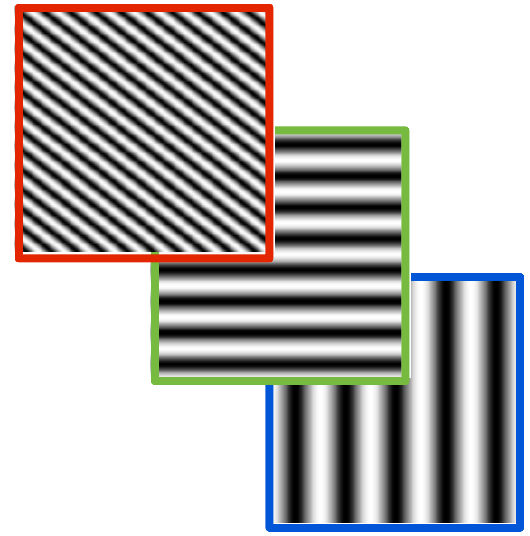

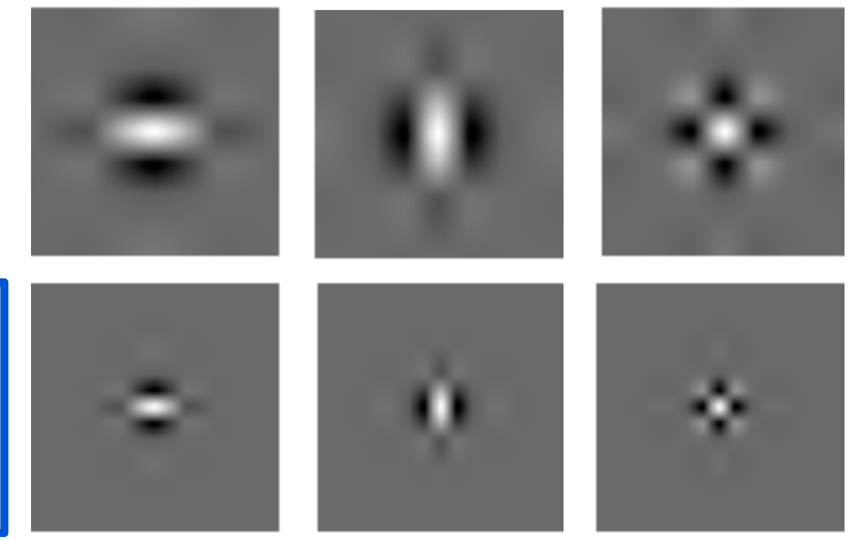

Conclusion

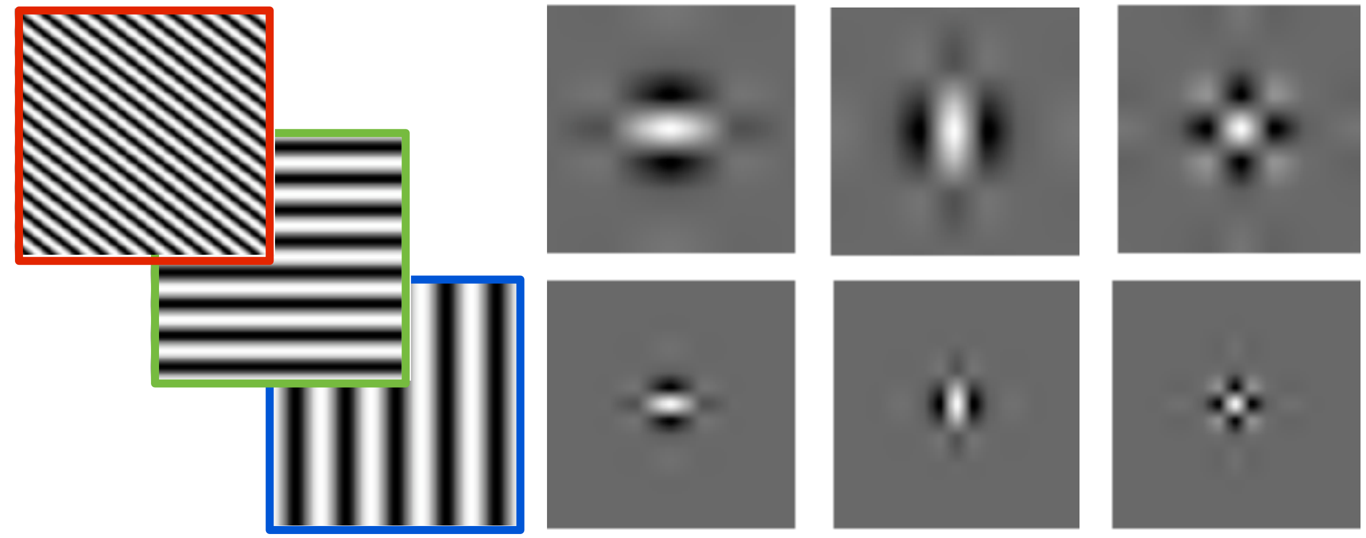

Orthogonal decompositions:

Fourier, wavelets, etc.

Linear vs. non-linear approximation:

Linear = filtering.

Non-linear = thresholding.

Conclusion

Orthogonal decompositions:

Fourier, wavelets, etc.

Linear vs. non-linear approximation:

Linear = filtering.

Non-linear = thresholding.

Compression, denoising: $\approx$ approximation.

Fast decay of approximation error

$\Longrightarrow$ efficient processing.