Adaptive Geometric Representations

Gabriel Peyré

www.numerical-tours.com

CNRS
NatImages
CEREMADE
Dauphine Université Paris

Natural Image Priors

"Typical" image drawn at random: (denoising noise)

Small $\|f\|_{\text{Sob}} = \int \|\nabla f\|^2$
Fourier decomposition

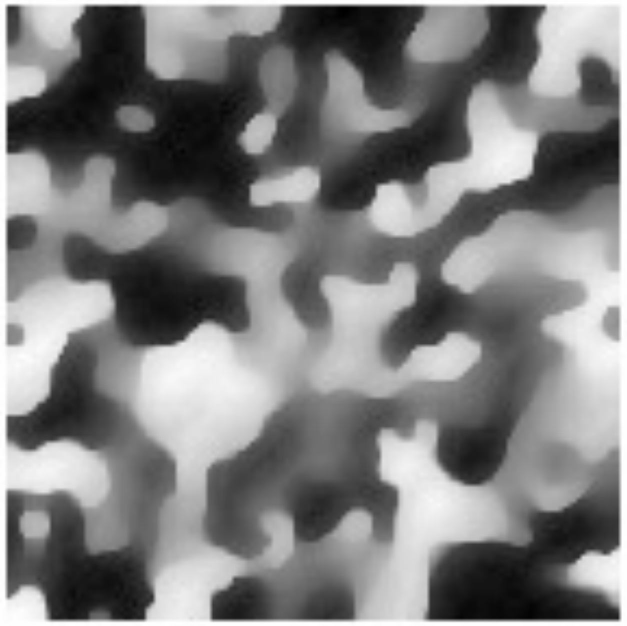

Small $\|f\|_{\text{TV}} = \int \|\nabla f\|$
Wavelet decomposition

Natural Image Priors

"Typical" image drawn at random: (denoising noise)

Small $\|f\|_{\text{Sob}} = \int \|\nabla f\|^2$
Fourier decomposition

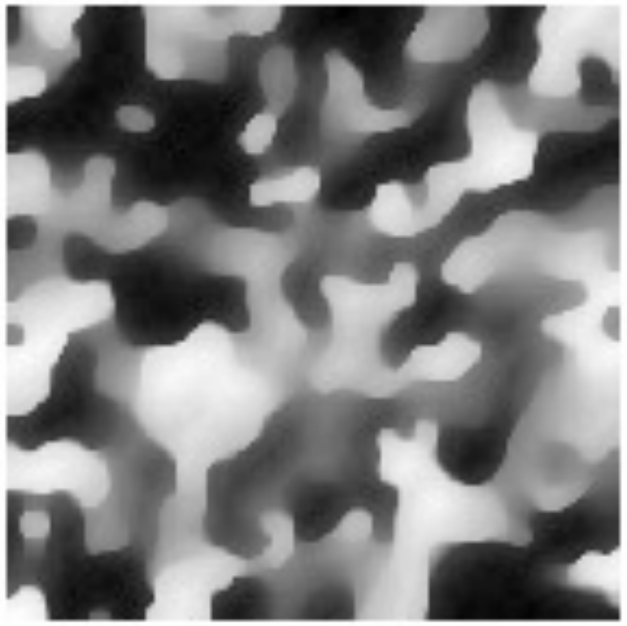

Small $\|f\|_{\text{TV}} = \int \|\nabla f\|$
Wavelet decomposition

Natural images: structure + texture + noise + ...

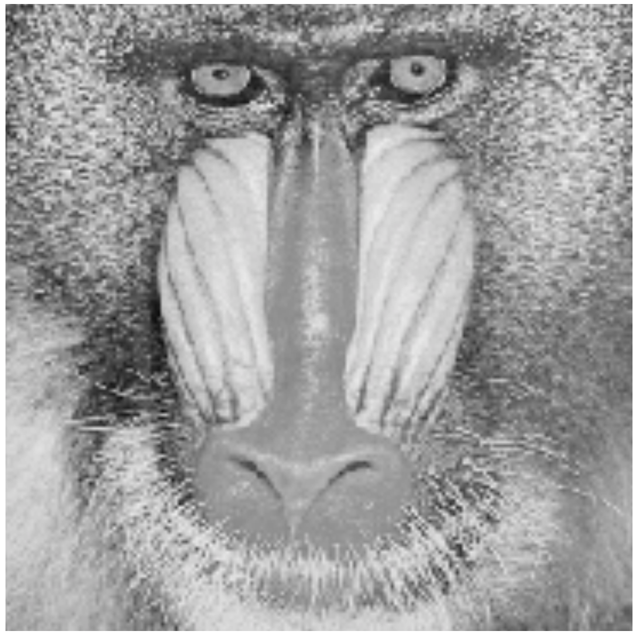

Overview

- Sparsity for Approximation
- Sparsity for Processing
- Geometric Images
- Adaptive Geometric Processing
- Adaptive Inverse Problems Regularization
- Geometric Texture Synthesis

Sparse Approximation in a Basis

Orthogonal basis $\mathcal{B} = (\psi_m)_m$

$$f = \sum_m \langle f, \psi_m \rangle \psi_m \longrightarrow f_M = \sum_{m \in I_T} \langle f, \psi_m \rangle \psi_m$$

Sparse Approximation in a Basis

Orthogonal basis $\mathcal{B} = (\psi_m)_m$

$$f = \sum_m \langle f,\, \psi_m \rangle \psi_m \longrightarrow f_M = \sum_{m \in I_T} \langle f,\, \psi_m \rangle \psi_m$$

Linear approximation:

I_T does not depend on f.

e.g.: $I_T = \{0, 1, \ldots, M-1\}$

(low frequencies)

Image f

Linear approximation

Sparse Approximation in a Basis

Orthogonal basis $\mathcal{B} = (\psi_m)_m$

$$f = \sum_m \langle f, \psi_m \rangle \psi_m \quad \longrightarrow \quad f_M = \sum_{m \in I_T} \langle f, \psi_m \rangle \psi_m$$

Linear approximation:

I_T does not depend on f.

e.g.: $I_T = \{0, 1, \ldots, M-1\}$

(low frequencies)

Image f

Linear approximation

Non-linear approximation:

minimize $\|f - f_M\|$ for a given M.

$I_T = \{m \setminus |\langle f, \psi_m \rangle| > T\}$

and $\quad M = \# I_T$.

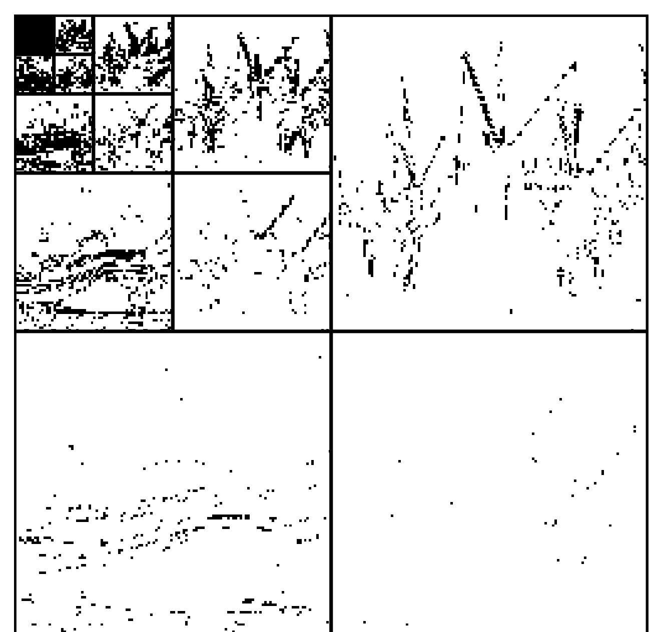

Coefficients I_T

Non-linear approximation

Image and Texture Models

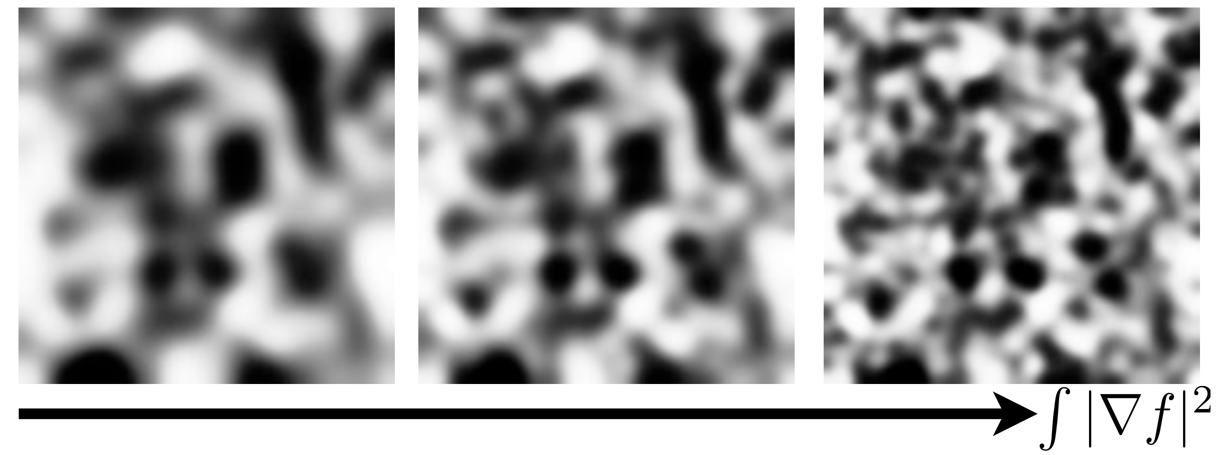

Uniformly smooth C^α *image.*

Fourier, Wavelets: $\|f - f_M\|^2 = O(M^{-\alpha})$.

Image and Texture Models

Uniformly smooth C^α *image.*

Fourier, Wavelets: $\|f - f_M\|^2 = O(M^{-\alpha})$.

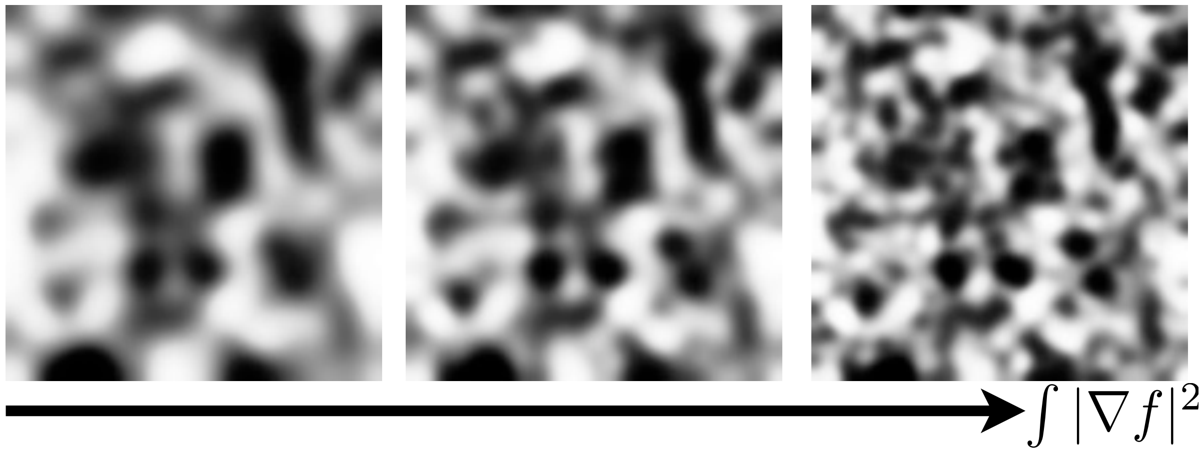

Discontinuous image with bounded variation.

Wavelets: $\|f - f_M\|^2 = O(M^{-1})$.

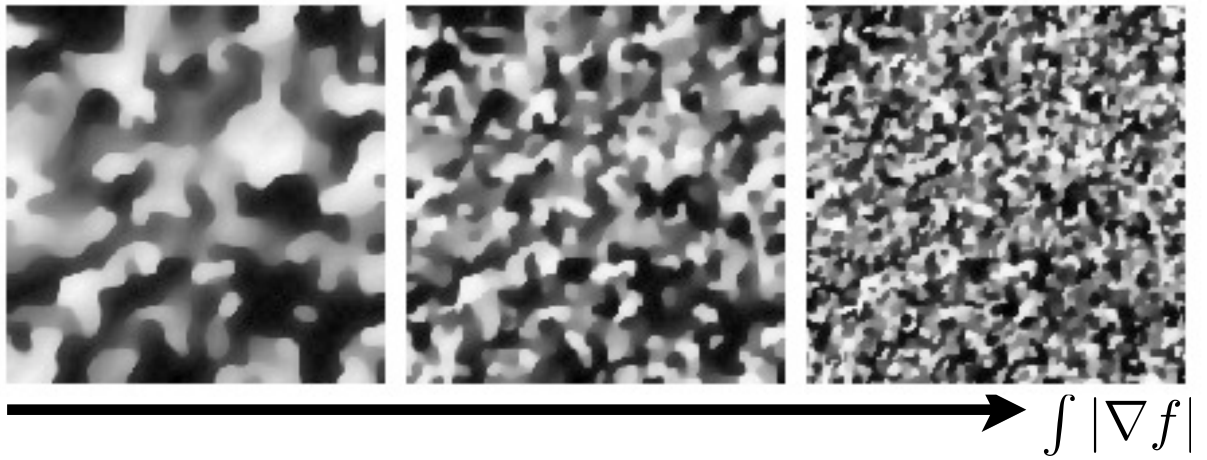

Image and Texture Models

Uniformly smooth C^α image.

Fourier, Wavelets: $\|f - f_M\|^2 = O(M^{-\alpha})$.

Discontinuous image with bounded variation.

Wavelets: $\|f - f_M\|^2 = O(M^{-1})$.

C^2-geometrically regular image.

Curvelets: $\|f - f_M\|^2 = O(\log^3(M) M^{-2})$.

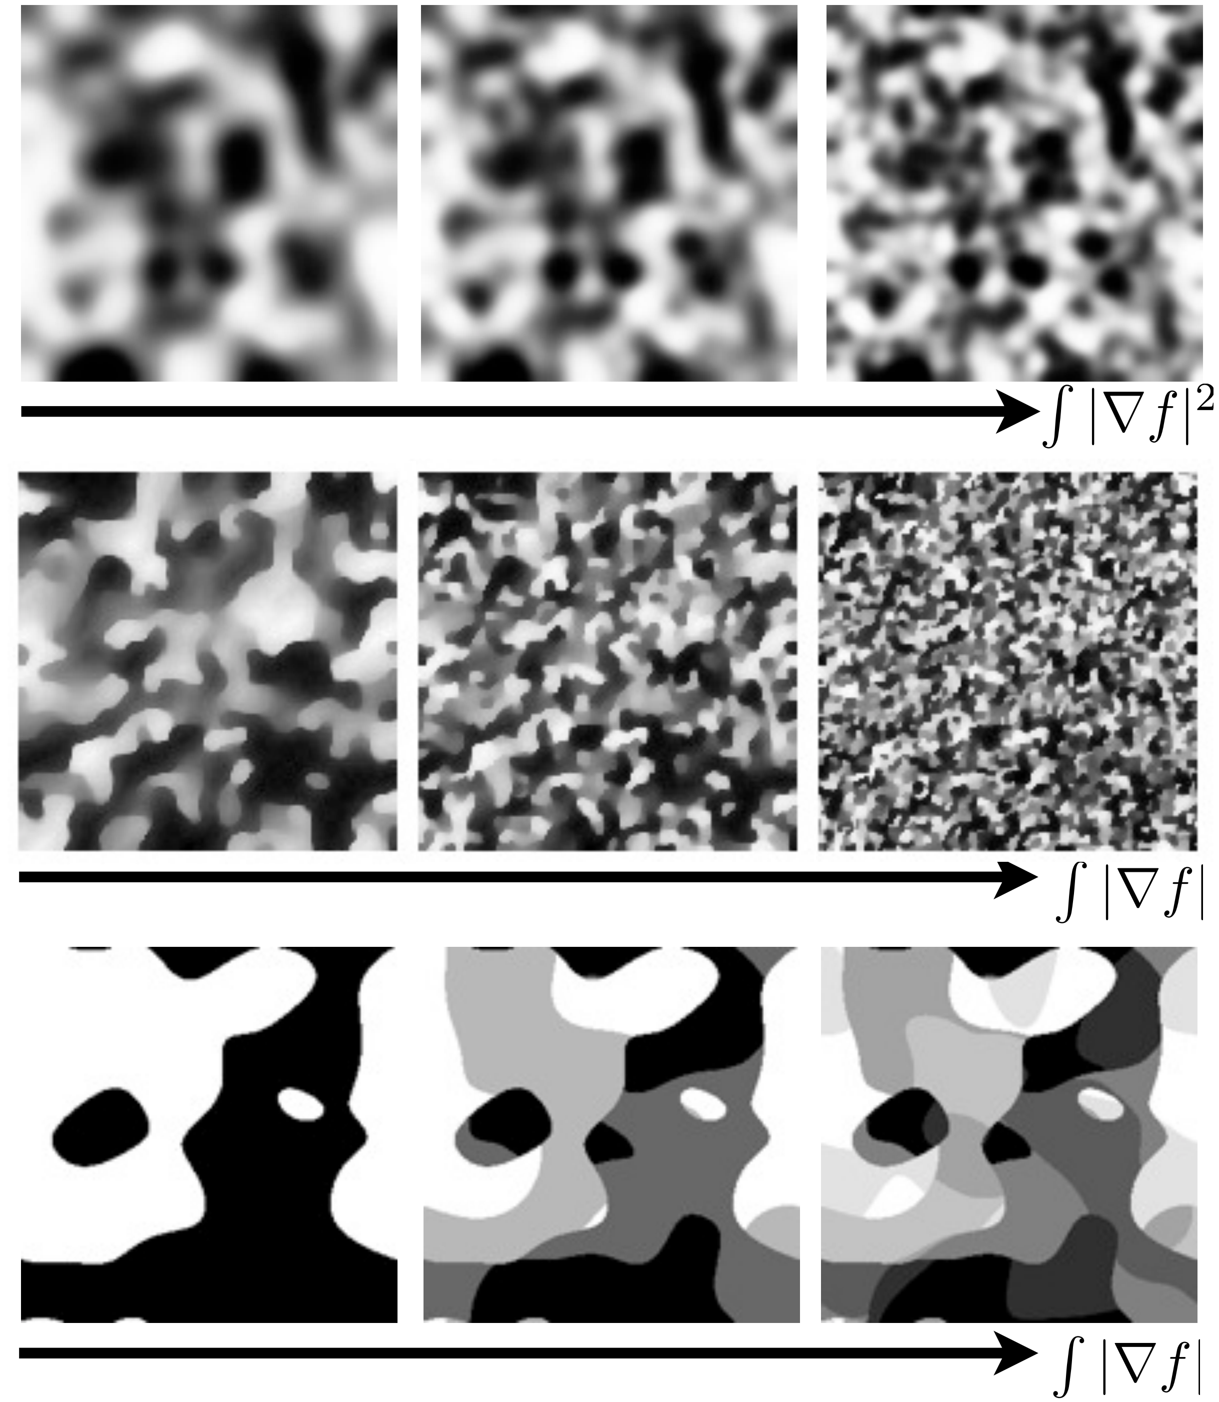

Image and Texture Models

Uniformly smooth C^α *image.*

Fourier, Wavelets: $\|f - f_M\|^2 = O(M^{-\alpha})$.

Discontinuous image with bounded variation.

Wavelets: $\|f - f_M\|^2 = O(M^{-1})$.

C^2*-geometrically regular image.*

Curvelets: $\|f - f_M\|^2 = O(\log^3(M) M^{-2})$.

C^α*-geometrically regular image.*

Bandlets: $\|f - f_M\|^2 = O(M^{-\alpha})$.

$\longrightarrow$ Adaptivity to the edge orientation.

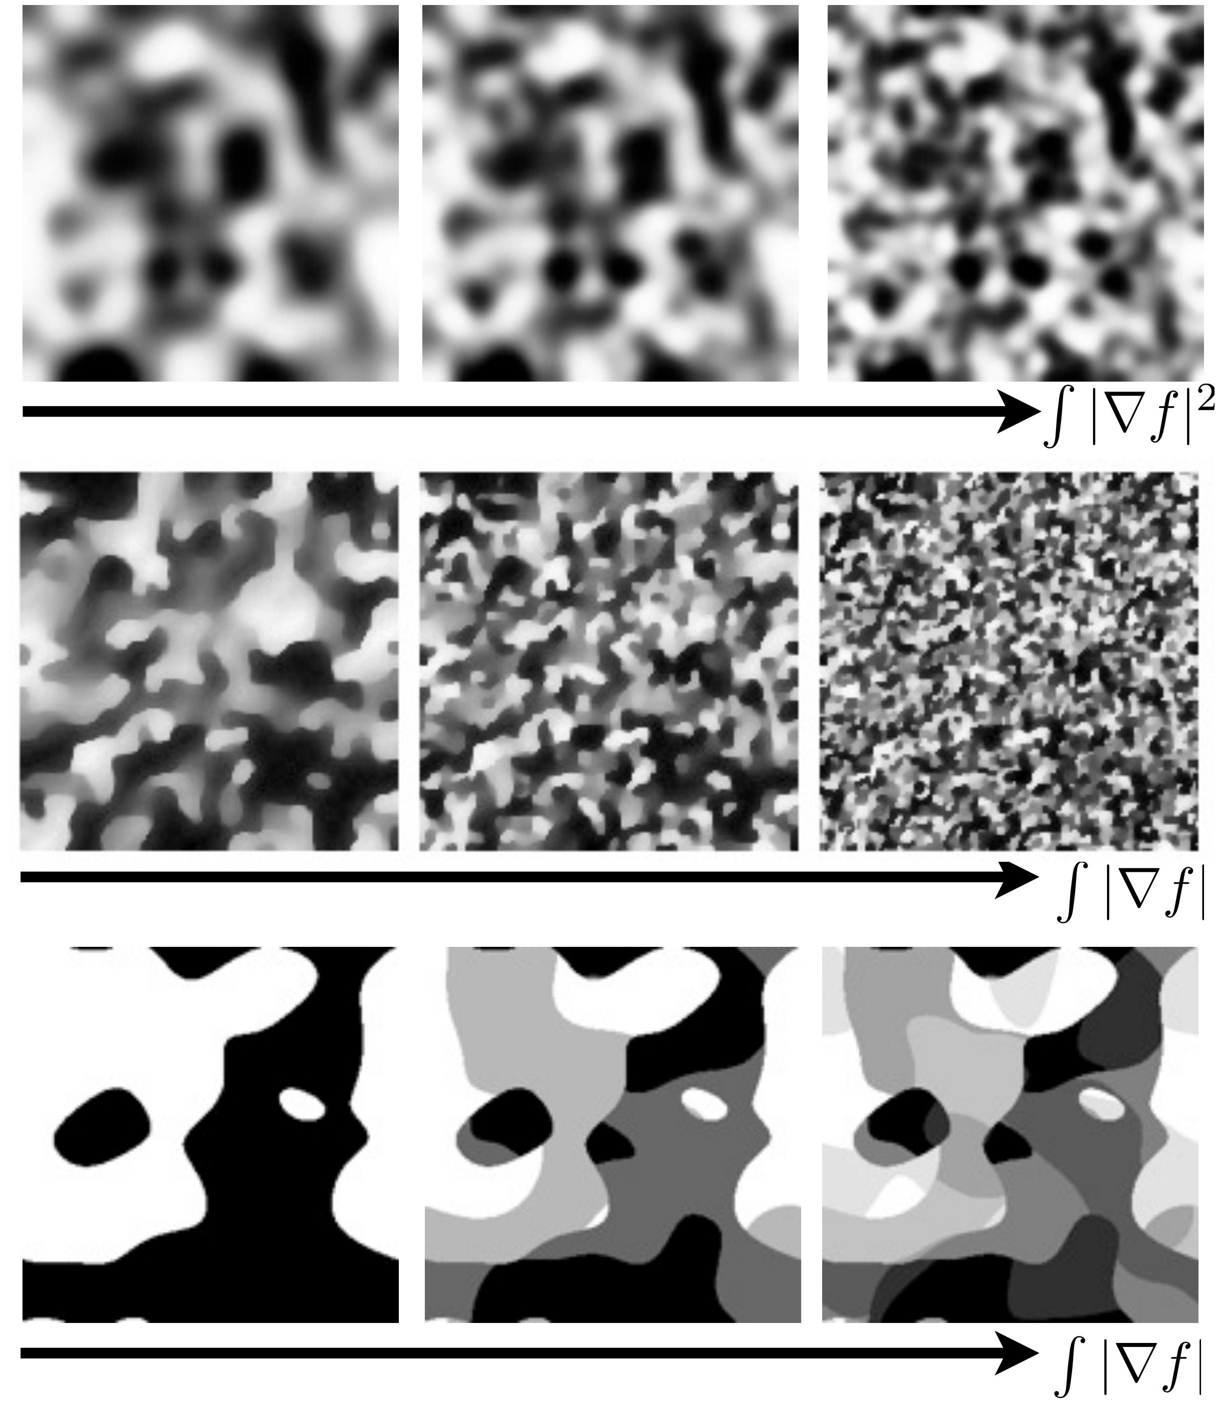

Image and Texture Models

Uniformly smooth C^α *image.*

Fourier, Wavelets: $\|f - f_M\|^2 = O(M^{-\alpha})$.

Discontinuous image with bounded variation.

Wavelets: $\|f - f_M\|^2 = O(M^{-1})$.

C^2*-geometrically regular image.*

Curvelets: $\|f - f_M\|^2 = O(\log^3(M) M^{-2})$.

C^α*-geometrically regular image.*

Bandlets: $\|f - f_M\|^2 = O(M^{-\alpha})$.

$\longrightarrow$ Adaptivity to the edge orientation.

More complex images: needs adaptivity.

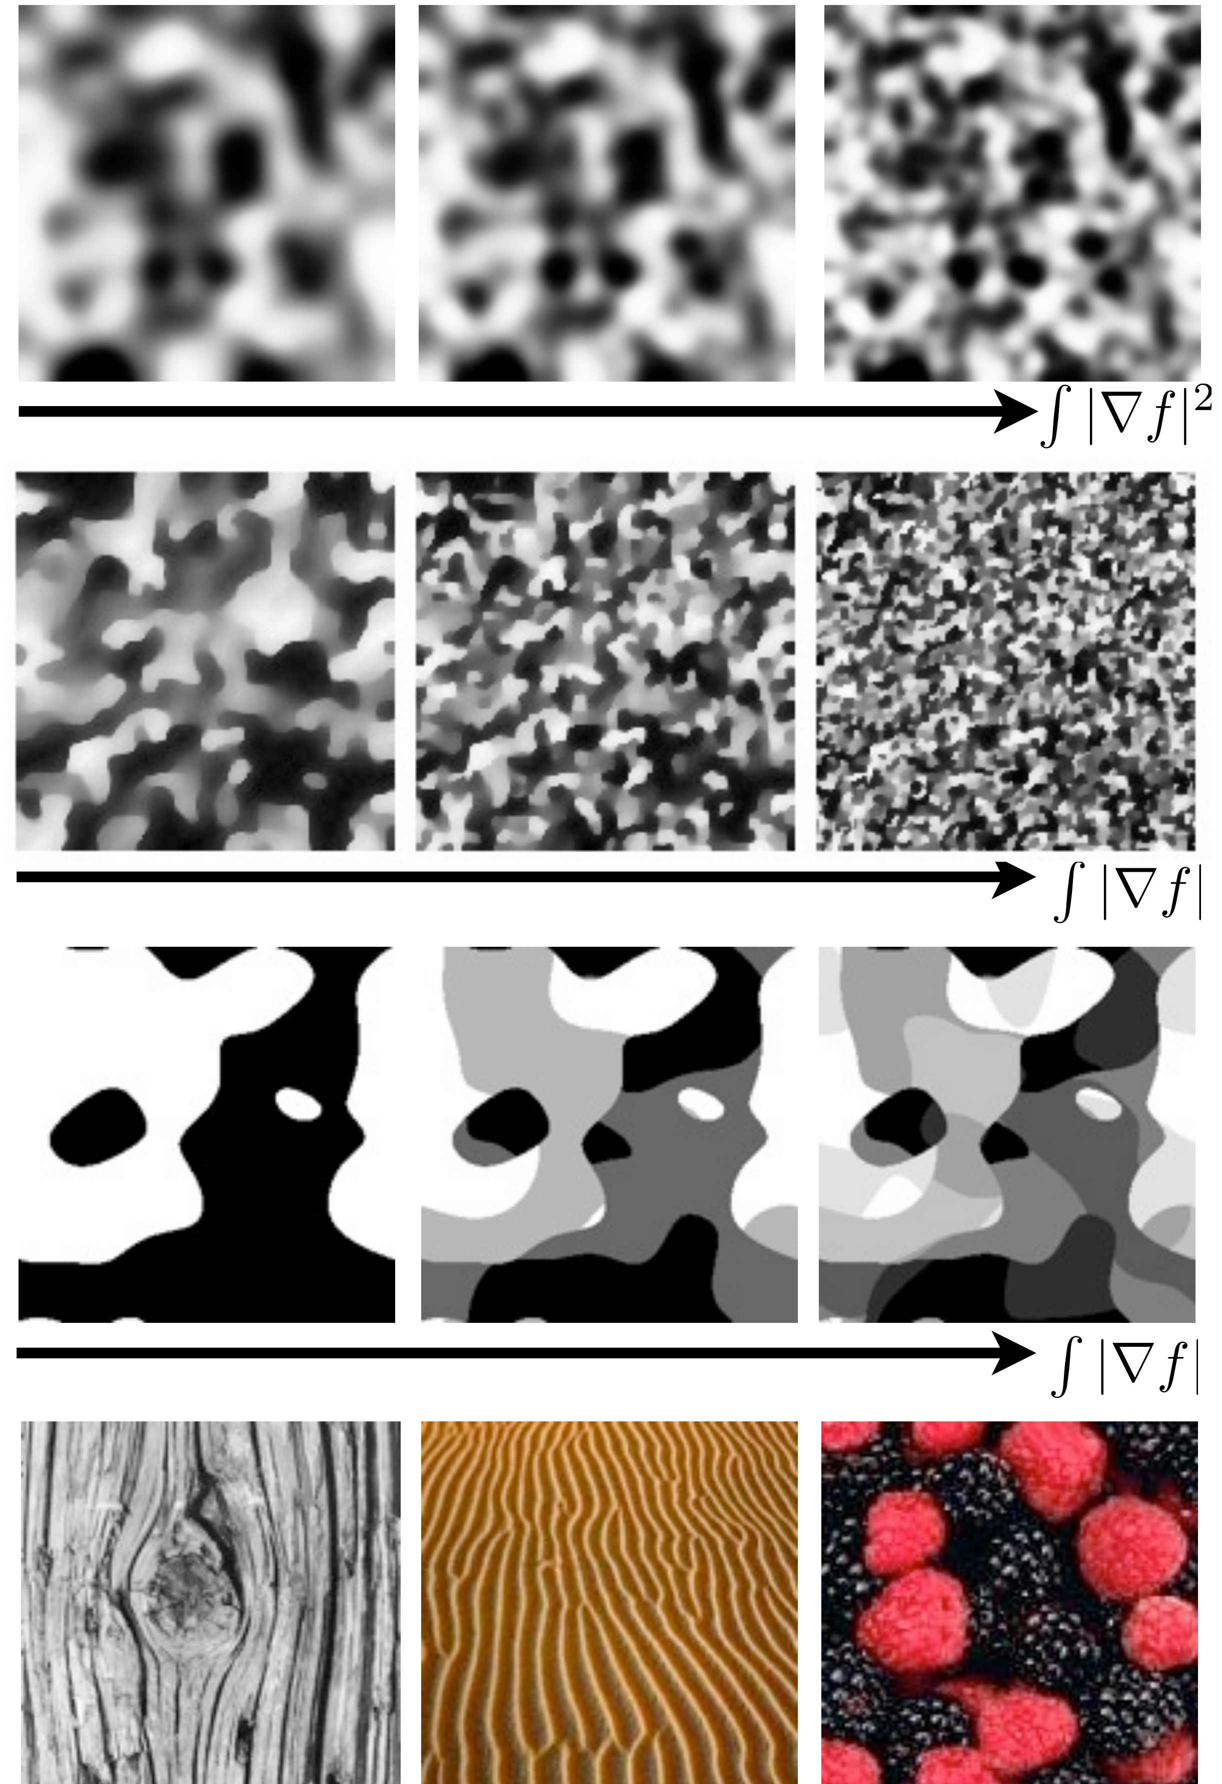

Overview

- Sparsity for Approximation
- Sparsity for Processing
- Geometric Images
- Adaptive Geometric Processing
- Adaptive Inverse Problems Regularization
- Geometric Texture Synthesis

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

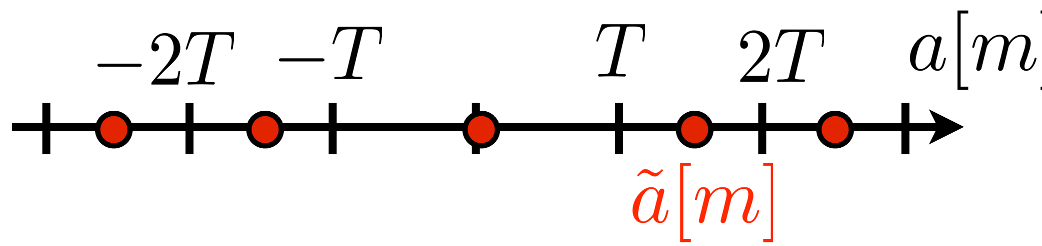

Image f

Zoom on f

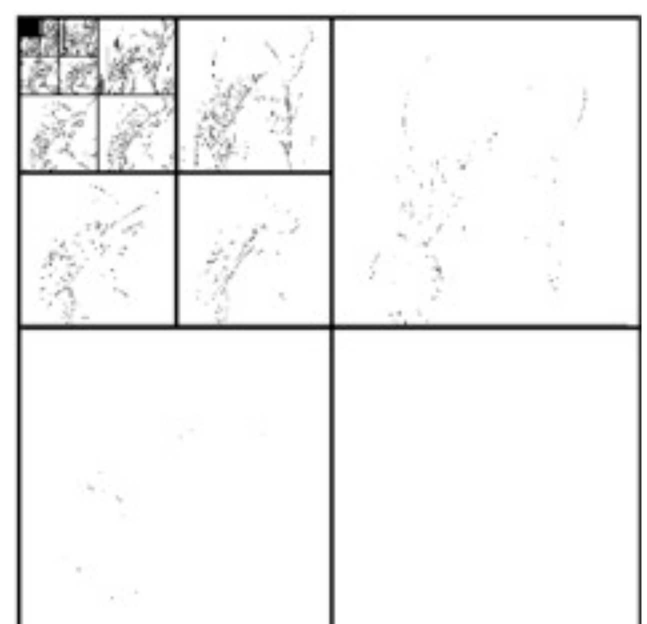

Quantized $q[m]$

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

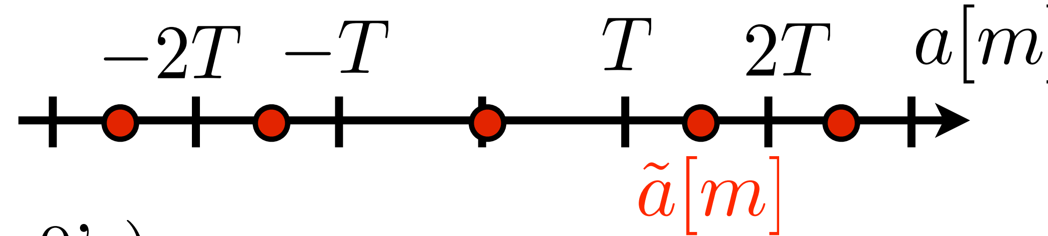

Entropic coding: use statistical redundancy (many 0's).

Image f

Zoom on f

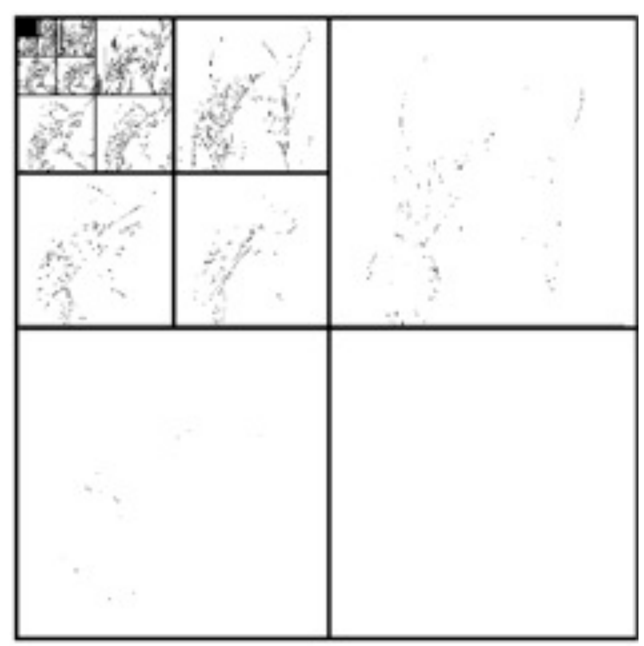

Quantized $q[m]$

Compression by Transform-coding

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

$$\underbrace{\{\, 0,1,0,0,1,\ldots \}}_{R \text{ bits}} \xrightarrow{\text{decoding}} q[m] \in \mathbb{Z}$$

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

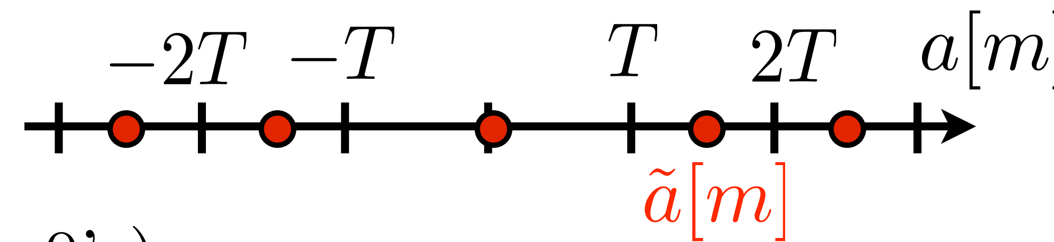

Entropic coding: use statistical redundancy (many 0's).

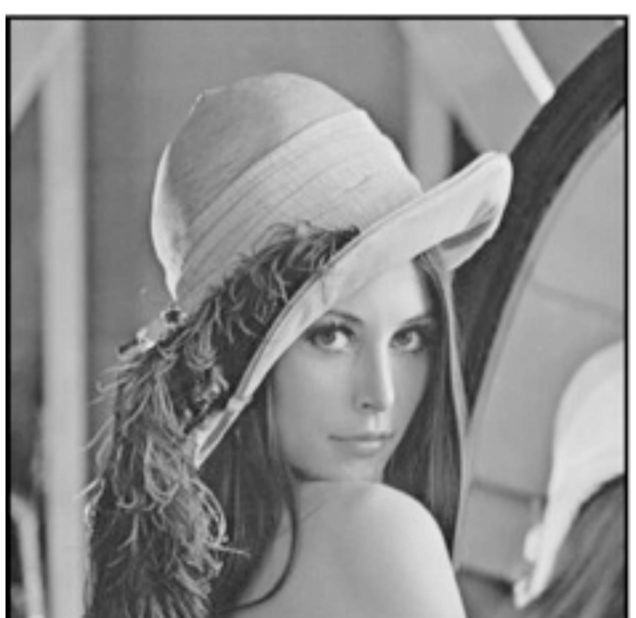

Image f

Zoom on f

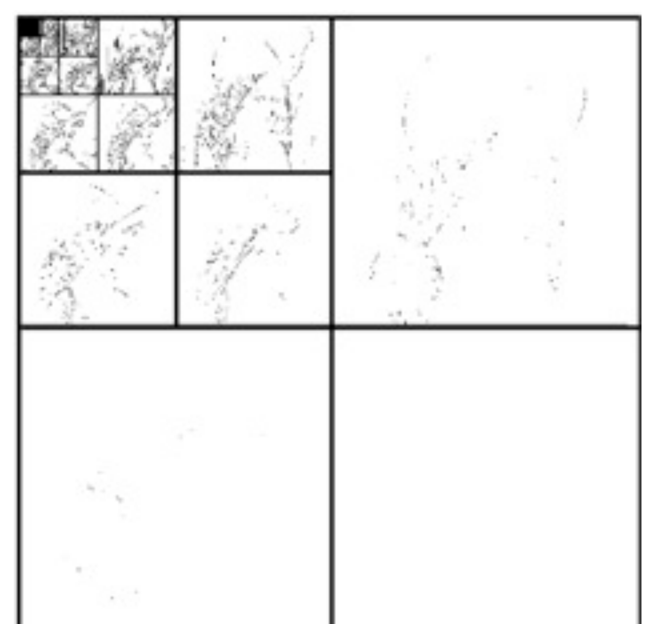

Quantized $q[m]$

Compression by Transform-coding

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

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

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

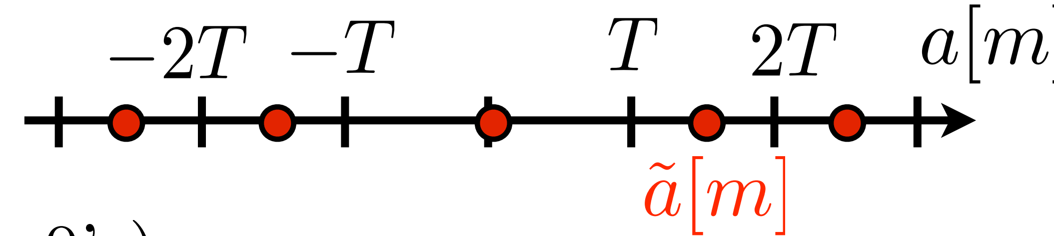

Entropic coding: use statistical redundancy (many 0's).

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

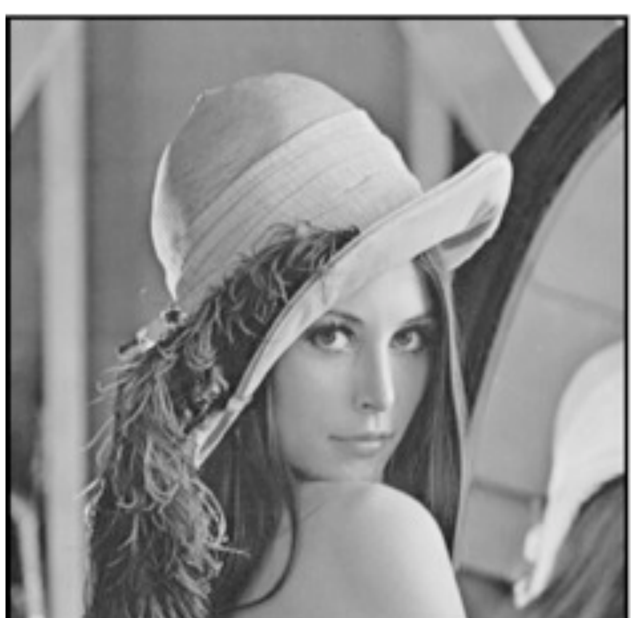
Image f


Zoom on f

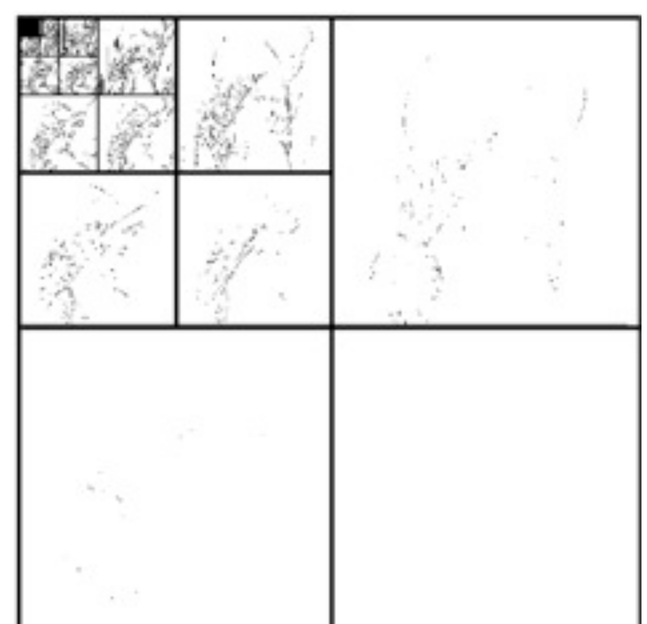
Quantized $q[m]$

Compression by Transform-coding

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

$$\underbrace{\{\, 0,1,0,0,1,\ldots \}}_{R \text{ bits}} \xrightarrow{\text{decoding}} q[m] \in \mathbb{Z}$$

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

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

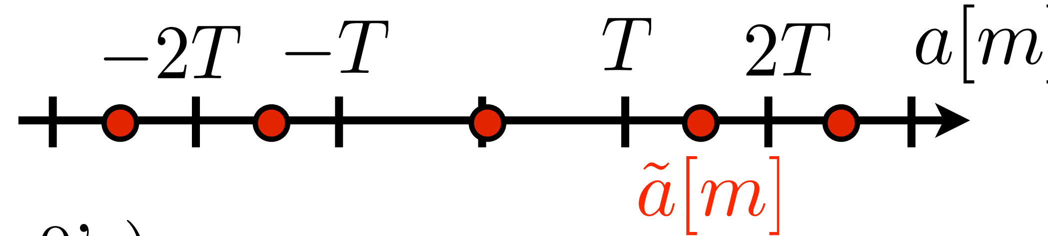

Entropic coding: use statistical redundancy (many 0's).

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

Image f

Zoom on f

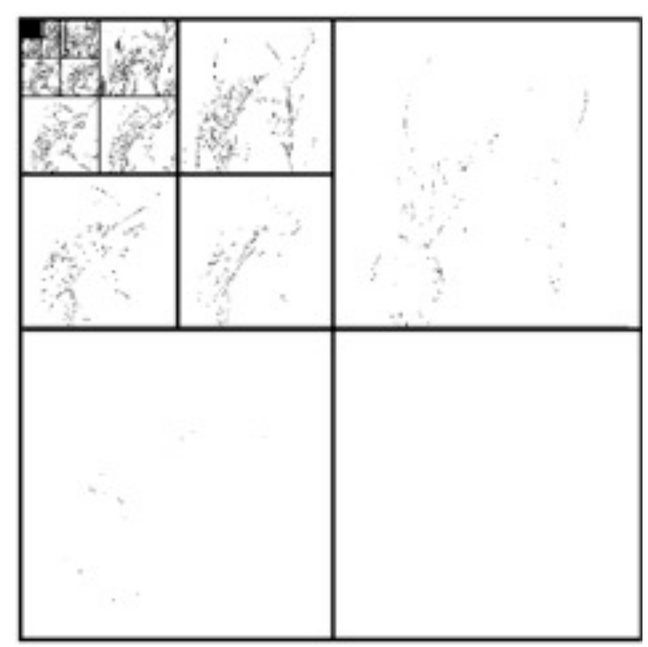

Quantized $q[m]$

f_R, R =0.2 bit/pixel

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

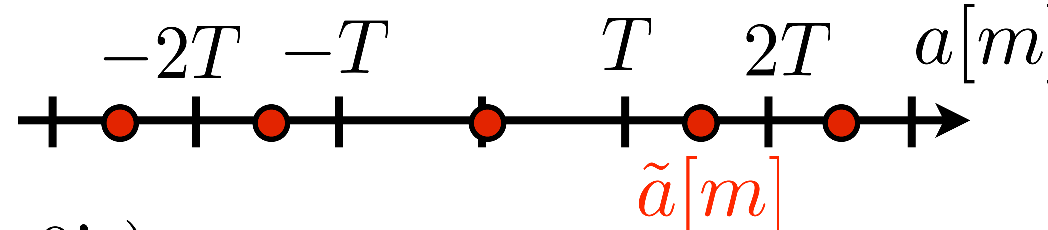

Entropic coding: use statistical redundancy (many 0's).

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

"*Theorem:*" $\quad \|f - f_M\|^2 = O(M^{-\alpha}) \implies \|f - f_R\|^2 = O(\log^\alpha(R) R^{-\alpha})$

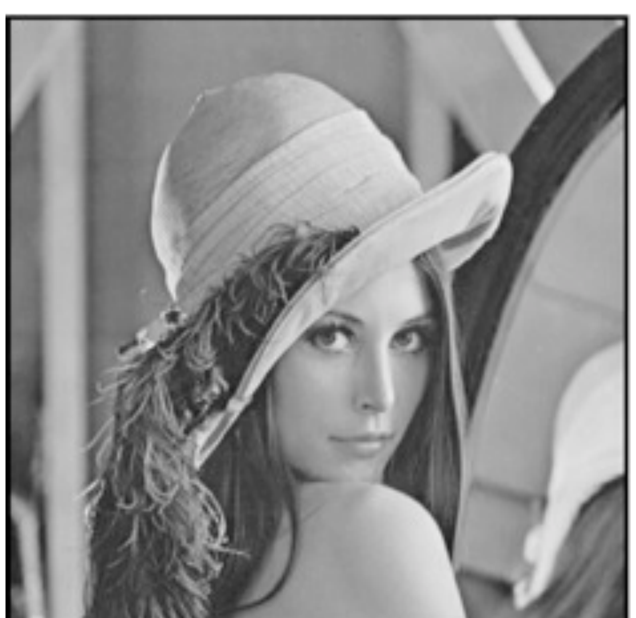
Image f

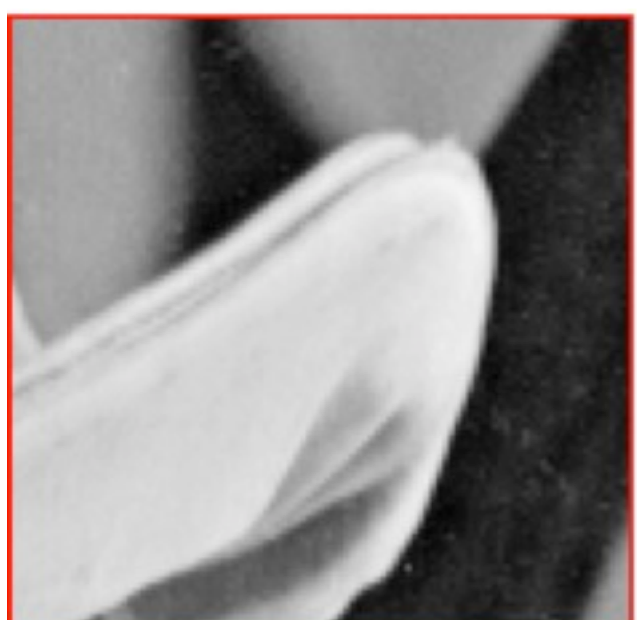
Zoom on f

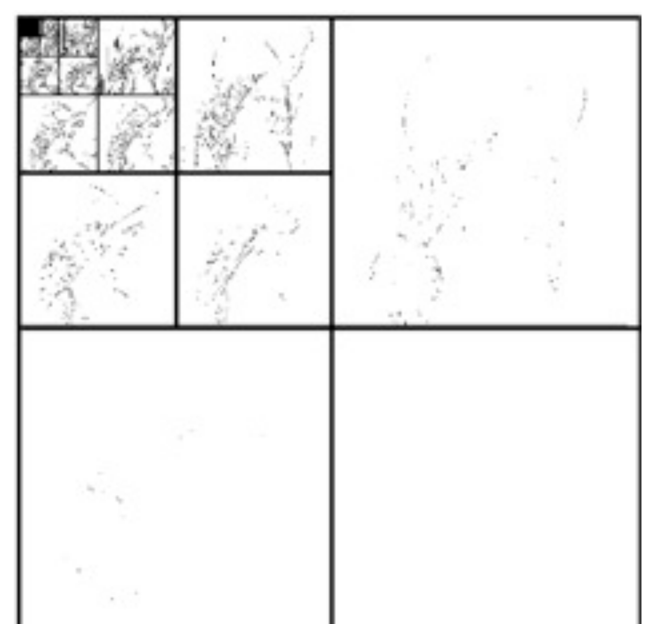
Quantized $q[m]$


f_R, R =0.2 bit/pixel

JPEG-2000 vs. JPEG

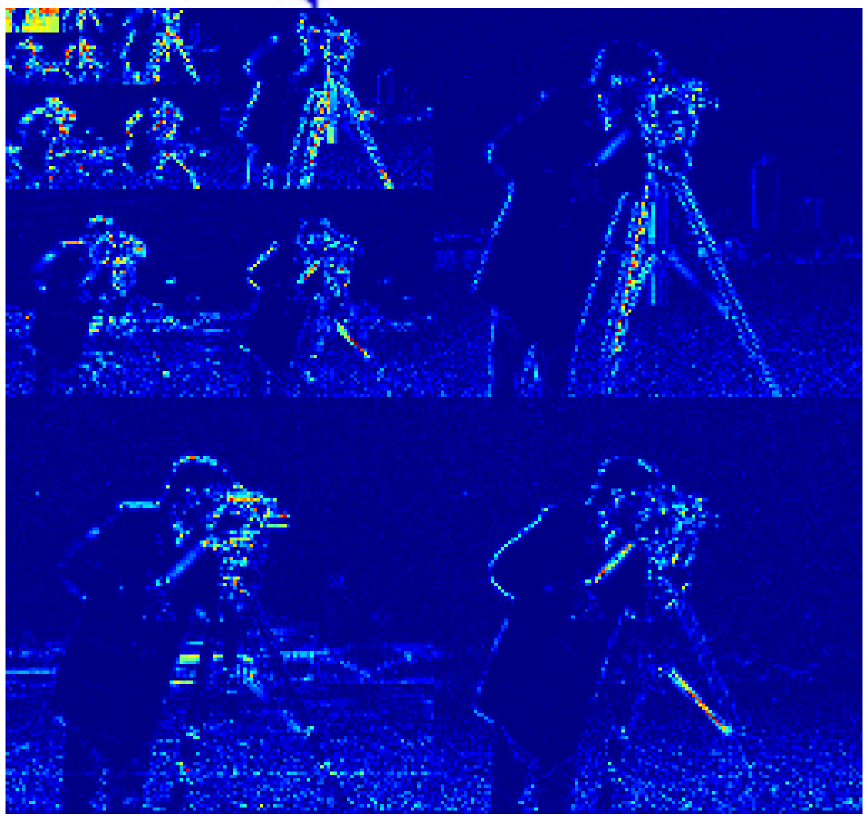

JPEG2k: exploit the statistical redundancy of coefficients.

+ embedded coder.

$\rightarrow$ chunks of large coefficients.

$\rightarrow$ neighboring coefficients are *not* independent.

Image f

JPEG, $R = .19$bit/pxl

JPEG2k, $R = .15$bit/pxl

Denoising (Donoho/Johnstone)

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 $f_0 = f + w$

Denoising (Donoho/Johnstone)

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

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

Denoising by approximation: $$f = \sum_{m=0}^{N-1} \langle f,\, \psi_m \rangle \psi_m \xrightarrow{\text{thresh.}} \tilde{f} = \sum_{|\langle f,\, \psi_m \rangle| > T} \langle f,\, \psi_m \rangle \psi_m$$

Clean f_0

Noisy $f_0 = f + w$

Denoised f

Denoising (Donoho/Johnstone)

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

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

Denoising by approximation: $f = \sum_{m=0}^{N-1} \langle f,\, \psi_m \rangle \psi_m \xrightarrow{\text{thresh.}} \tilde{f} = \sum_{|\langle f,\, \psi_m \rangle| > T} \langle f,\, \psi_m \rangle \psi_m$

Theorem: if $\|f_0 - f_{0,M}\|^2 = O(M^{-\alpha})$,

$$\|\tilde{f} - f_0\|^2 = O(\sigma^{\frac{2\alpha}{\alpha+1}}) \quad \text{for} \quad T = \sqrt{2\log(N)}\sigma$$

In practice:
$T \approx 3\sigma$

Clean f_0

Noisy $f_0 = f + w$

Denoised 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, $P = N - |\Omega|$.

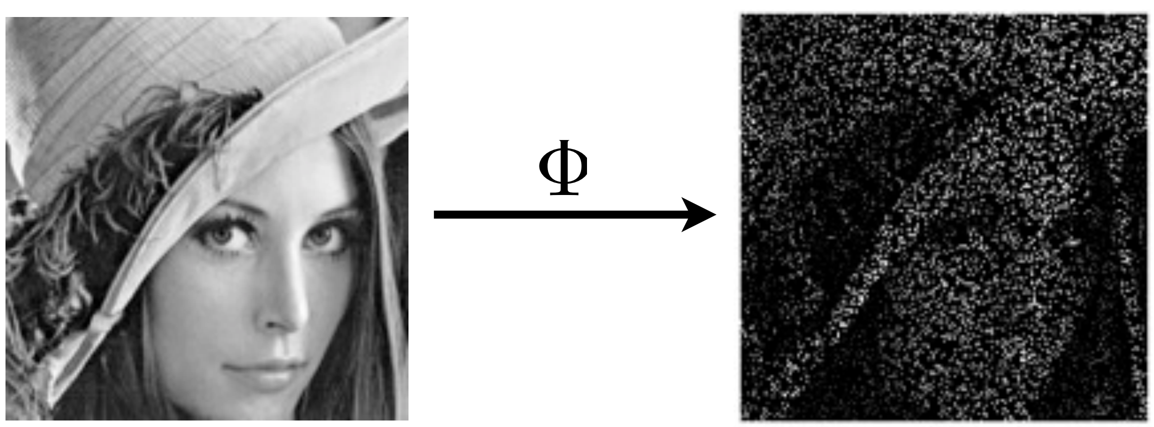

$$(\Phi f)(x) = \begin{cases} 0 & \text{if} \quad x \in \Omega, \\ f(x) & \text{if} \quad x \notin \Omega. \end{cases}$$

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, $P = N - |\Omega|$.

$\xrightarrow{\Phi}$

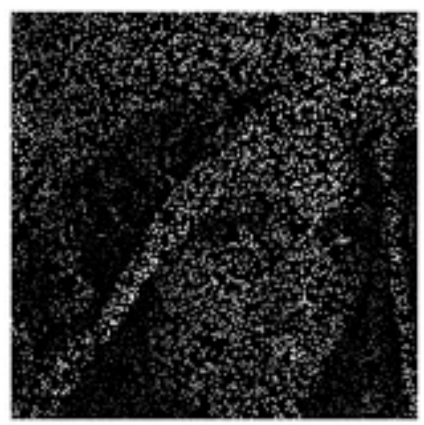

$$(\Phi f)(x) = \begin{cases} 0 & \text{if} \quad x \in \Omega, \\ f(x) & \text{if} \quad x \notin \Omega. \end{cases}$$

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

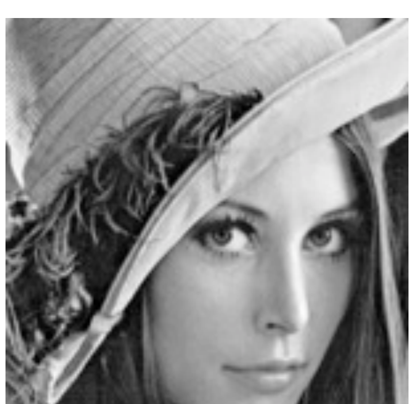

$\xrightarrow{\Phi}$

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.

Image f

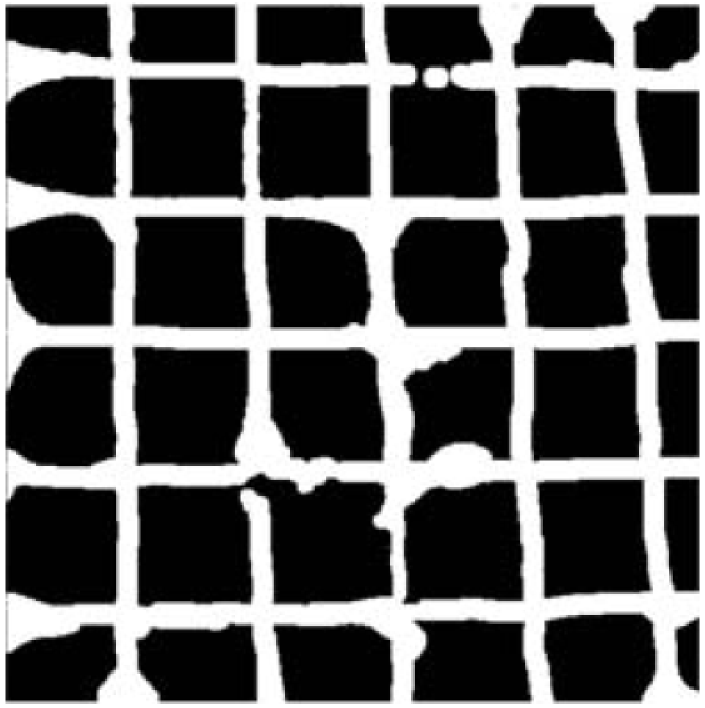

Mask Ω

Inpainted $\tilde{f}$

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.

Image f

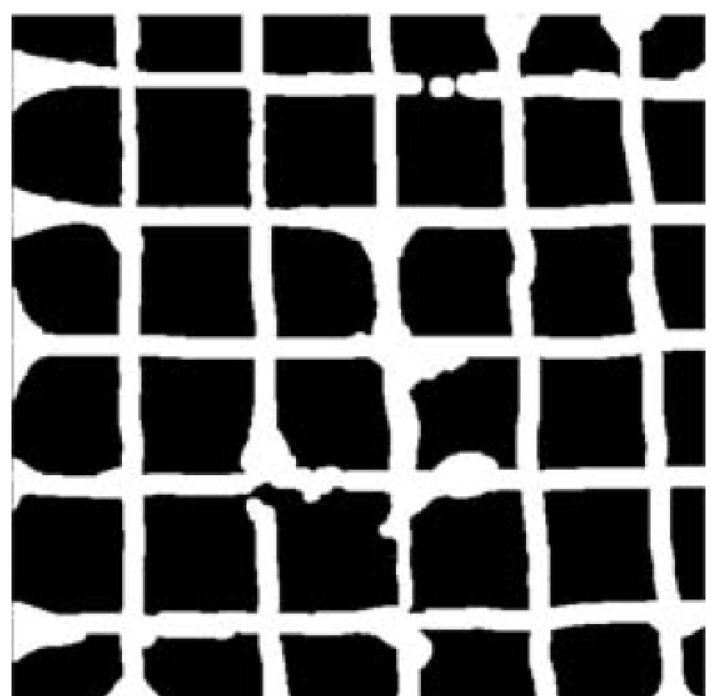

Mask Ω

Inpainted $\tilde{f}$

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

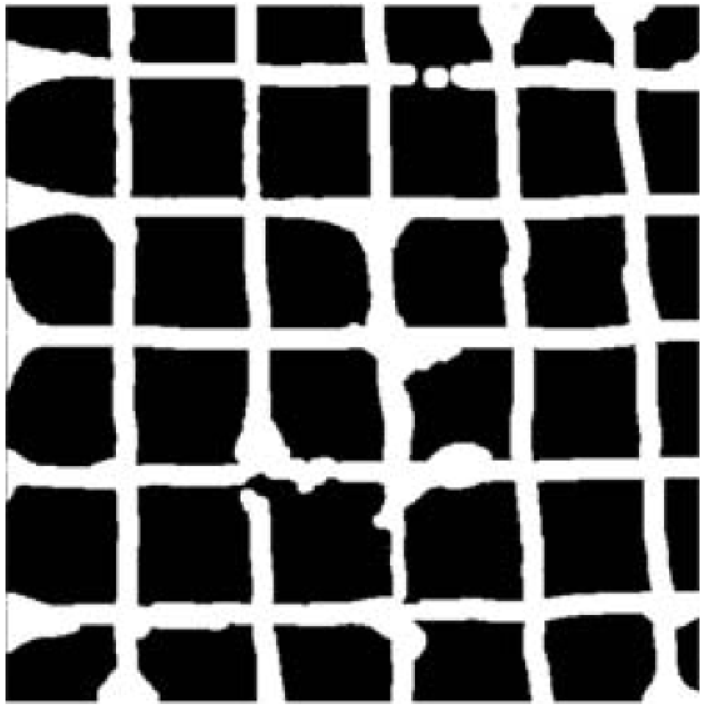

Mask Ω

Inpainted $\tilde{f}$

Overview

- Sparsity for Approximation
- Sparsity for Processing
- Geometric Images
- Adaptive Geometric Processing
- Adaptive Inverse Problems Regularization
- Geometric Texture Synthesis

Piecewise Regular Functions in 1D

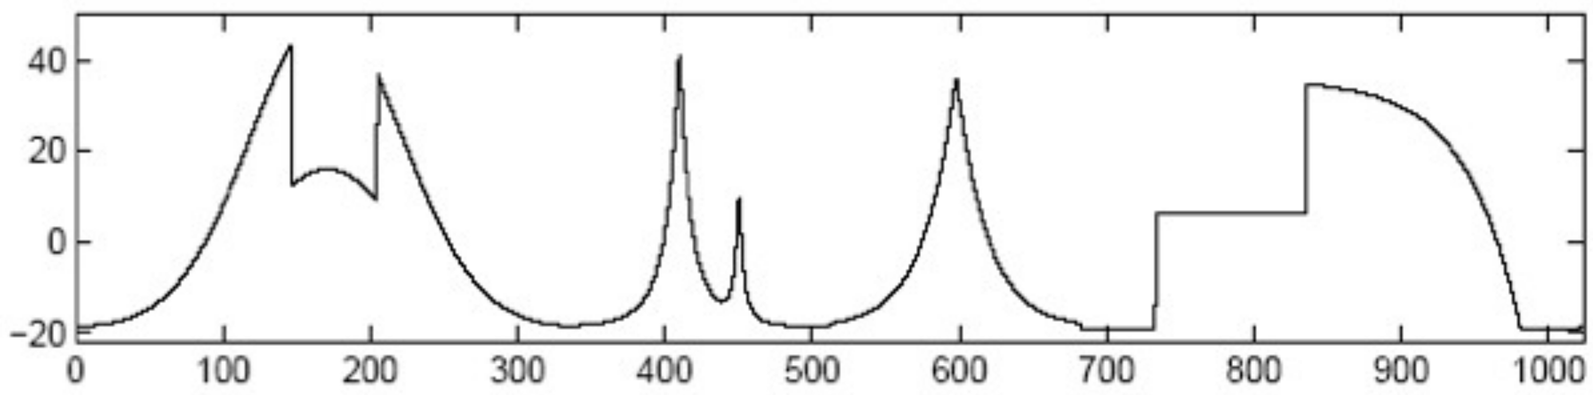

ψ

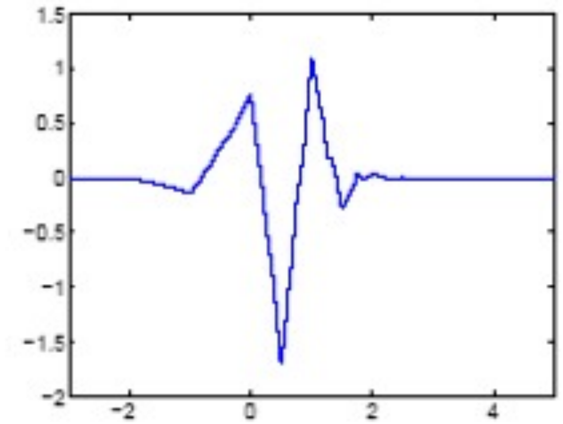

Theorem: If f is C^α outside a finite set of discontinuities:

$$\|f - f_M\|^2 = \begin{cases} O(M^{-1}) & \text{(Fourier)}, \\ O(M^{-2\alpha}) & \text{(wavelets)}. \end{cases}$$

For Fourier, linear$\approx$non-linear, sub-optimal.

For wavelets, linear$\ll$non-linear, optimal.

Piecewise Regular Functions in 2D

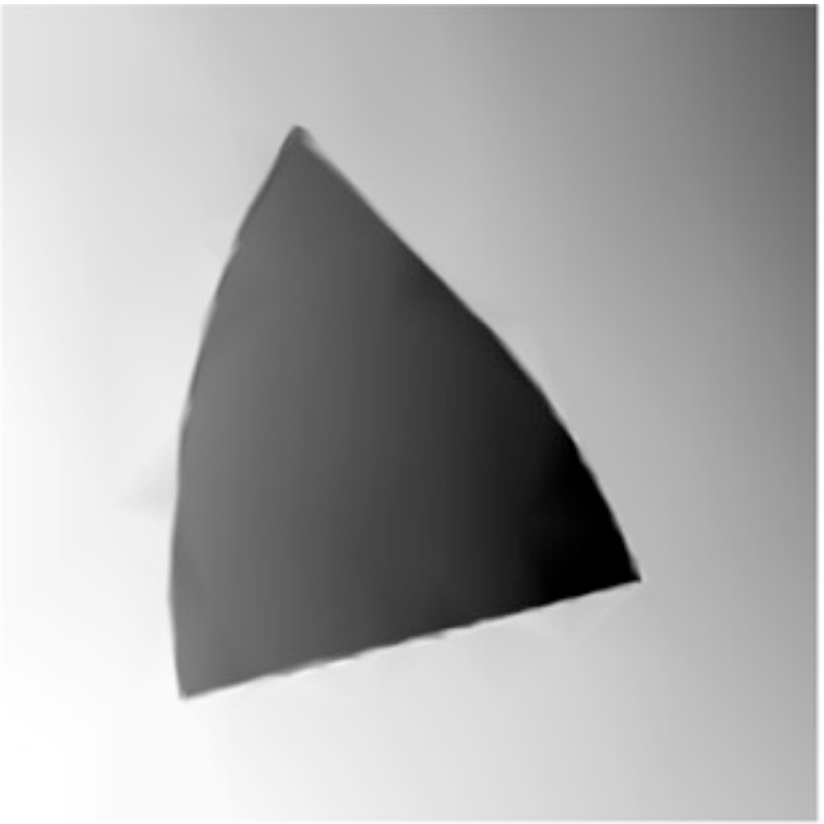

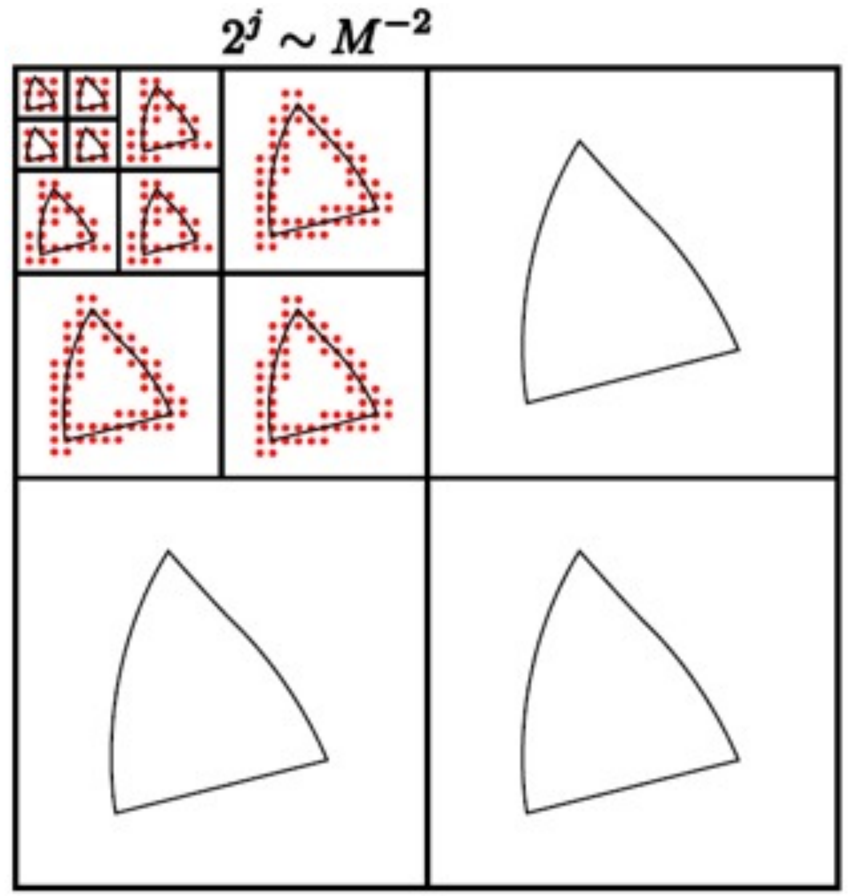

Theorem: If f is C^α outside a set of finite length edge curves,

$$\|f - f_M\|^2 = \begin{cases} O(M^{-1/2}) & \text{(Fourier)}, \\ O(M^{-1}) & \text{(wavelets)}. \end{cases}$$

Fourier $\ll$ Wavelets.

Wavelets: same result for BV functions (optimal).

Piecewise Regular Functions in 2D

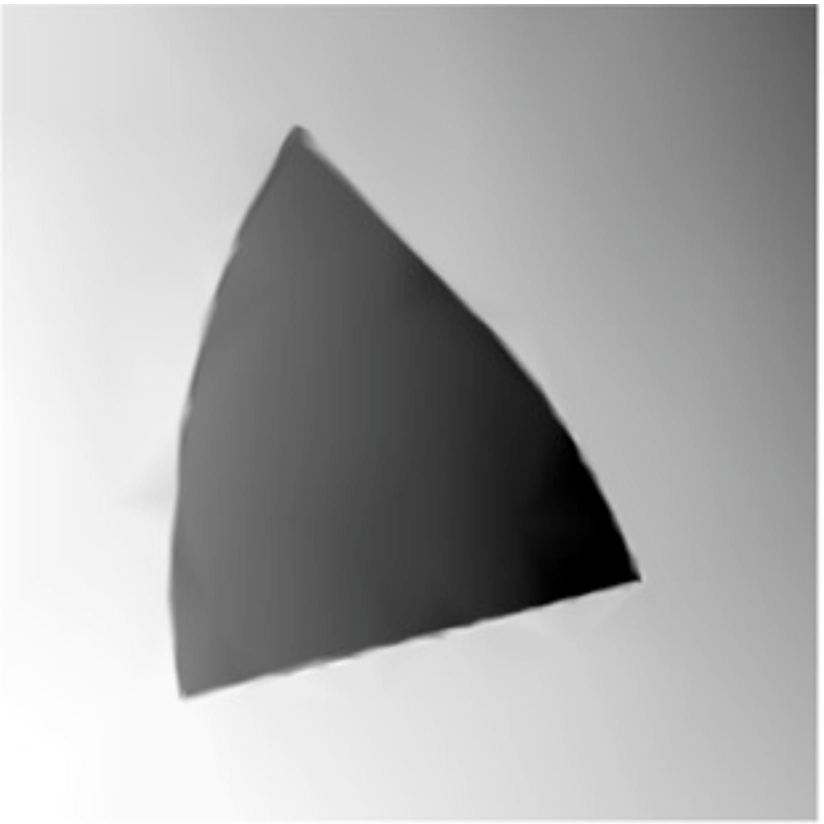

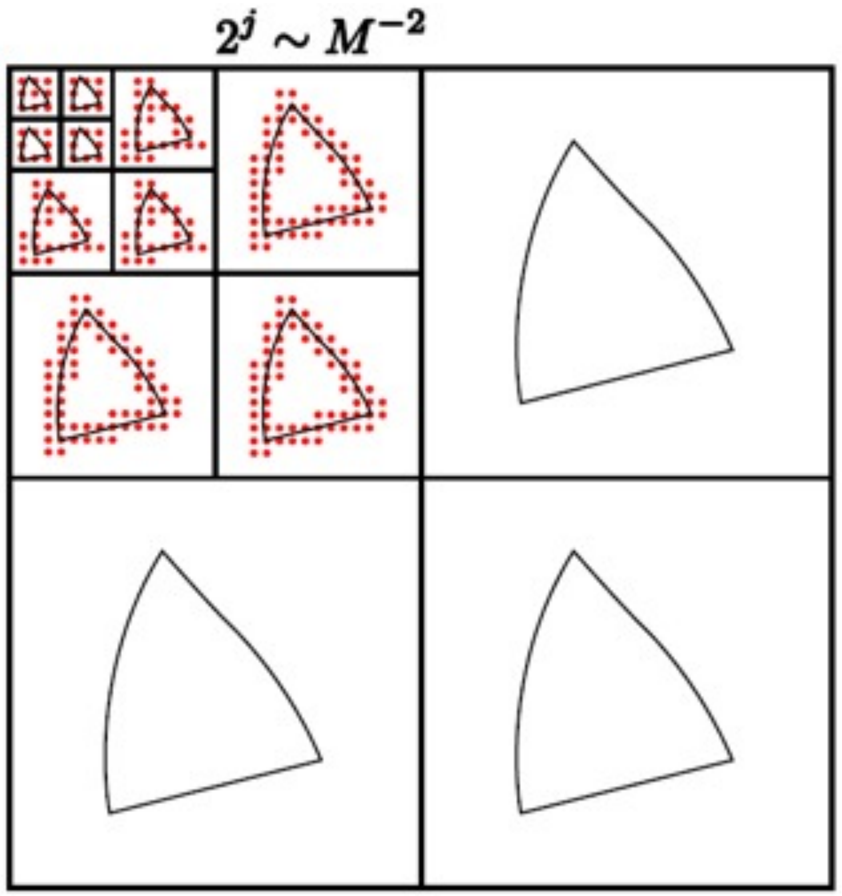

Theorem: If f is C^α outside a set of finite length edge curves,

$$\|f - f_M\|^2 = \begin{cases} O(M^{-1/2}) & \text{(Fourier)}, \\ O(M^{-1}) & \text{(wavelets)}. \end{cases}$$

Fourier $\ll$ Wavelets.

Wavelets: same result for BV functions (optimal).

Regular C^α edges: sub-optimal (requires anisotropy).

Geometrically Regular Images

Geometric image model: f is C^α outside a set of C^α edge curves.

BV image: level sets have finite lengths.

Geometric image: level sets are regular.


Geometry = cartoon image


Sharp edges


Smoothed edges

Curvelets for Cartoon Images

Curvelets: [Candes, Donoho] [Candes, Demanet, Ying, Donoho]

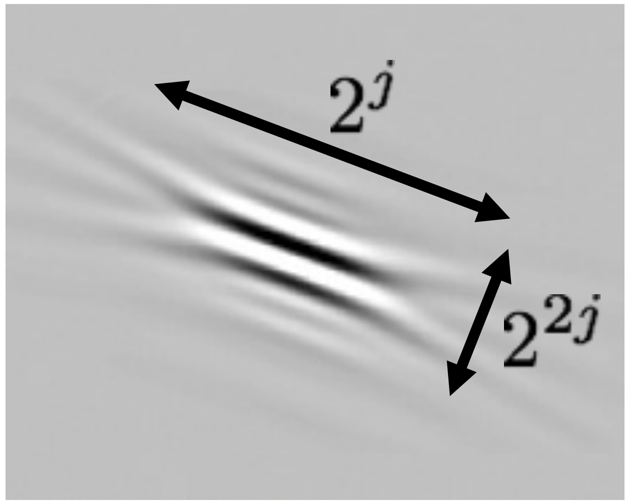

If f is C^α outside C^α edges, for $\alpha \geq 2$

$$\|f - f_M\|^2 = O(\log^3(M) M^{-2}).$$

`www.curvelet.org`

Curvelets for Cartoon Images

Curvelets: [Candes, Donoho] [Candes, Demanet, Ying, Donoho]

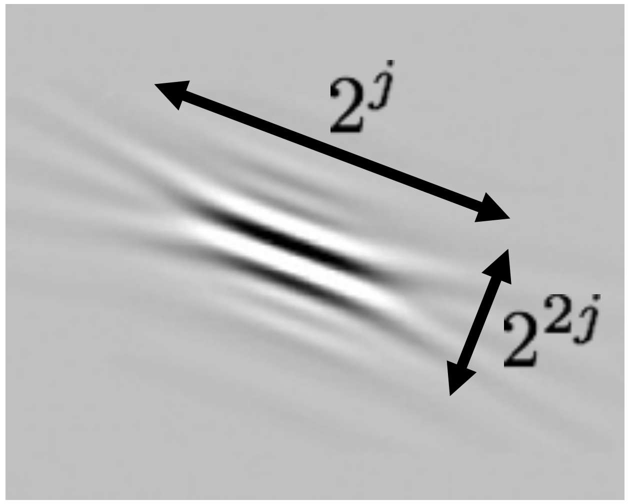

If f is C^α outside C^α edges, for $\alpha \geq 2$

$$\|f - f_M\|^2 = O(\log^3(M) M^{-2}).$$

`www.curvelet.org`

Redundant tight frame (redundancy ≈ 5): not efficient for compression.

Denoising by curvelet thresholding: recovers edges and geometric textures.

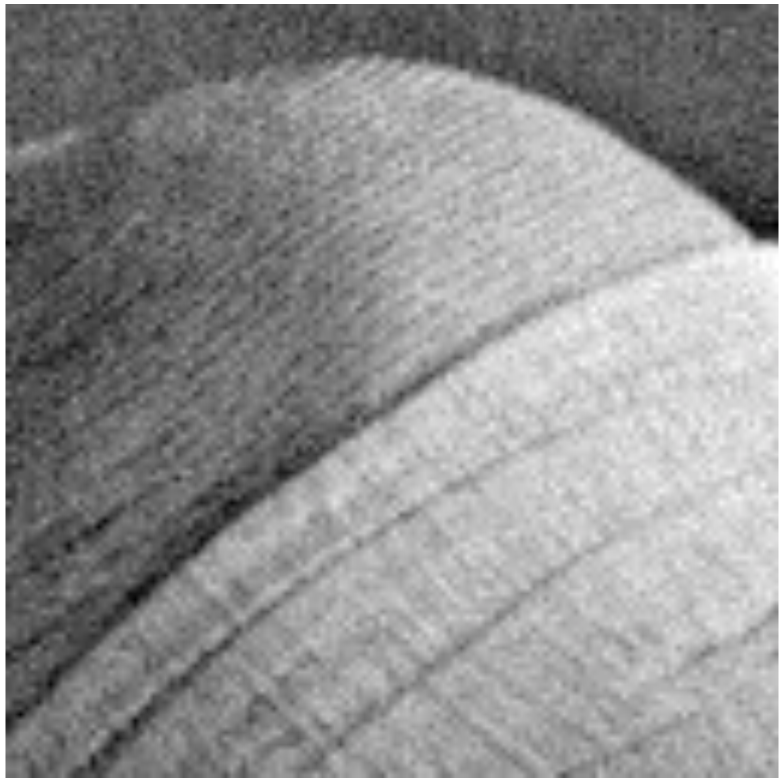

Noisy

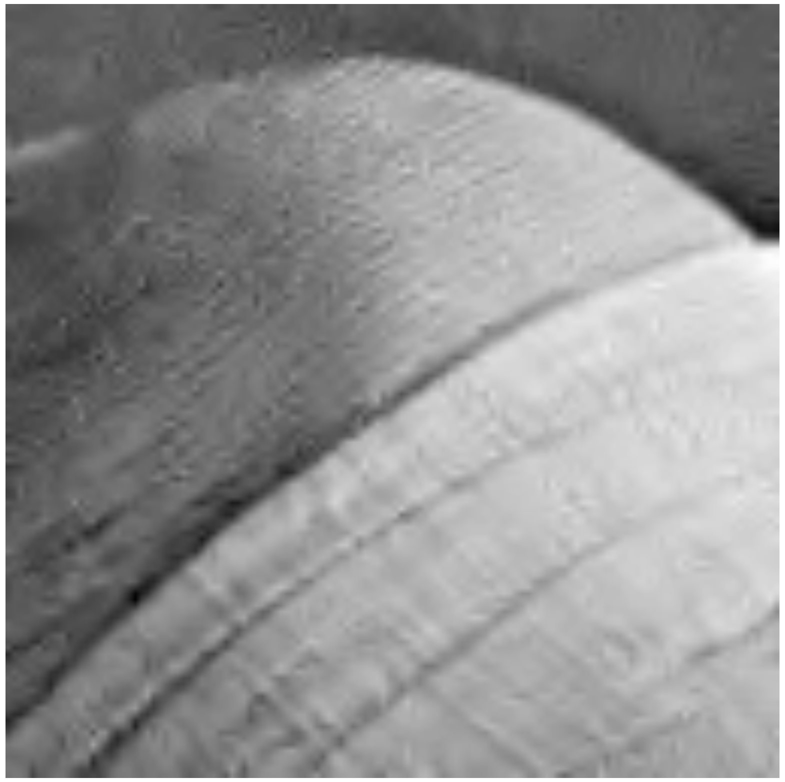

Wavelets

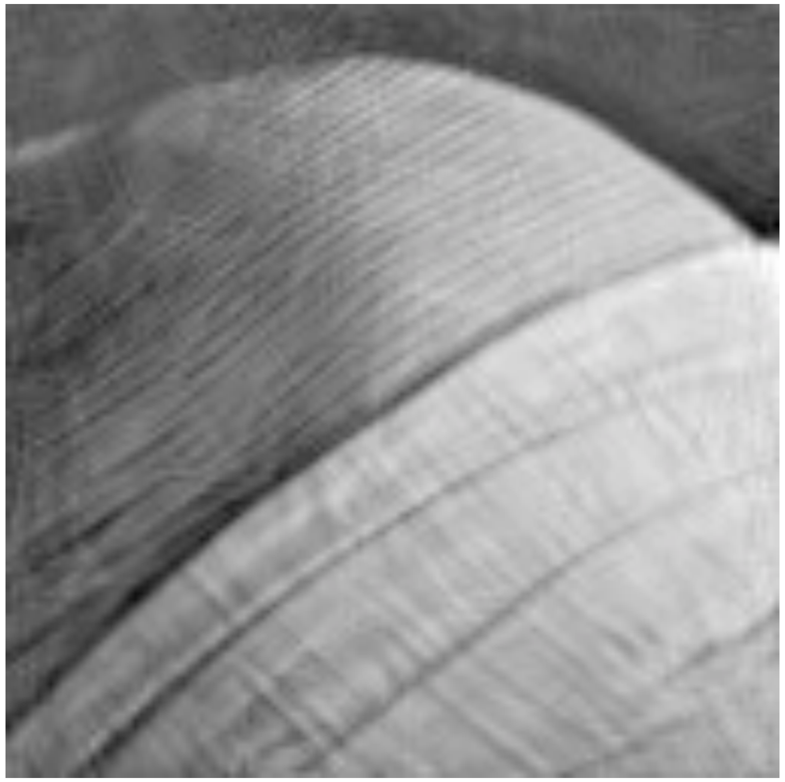

Curvelets

Approximation with Triangulation

Triangulation $(\mathcal{V}, \mathcal{F})$:

$$\left| \begin{array}{l} \text{Vertices } \mathcal{V} = \{v_i\}_{i=1}^{M}. \\ \text{Faces } \mathcal{F} \subset \{1, \ldots, M\}^3. \end{array} \right.$$

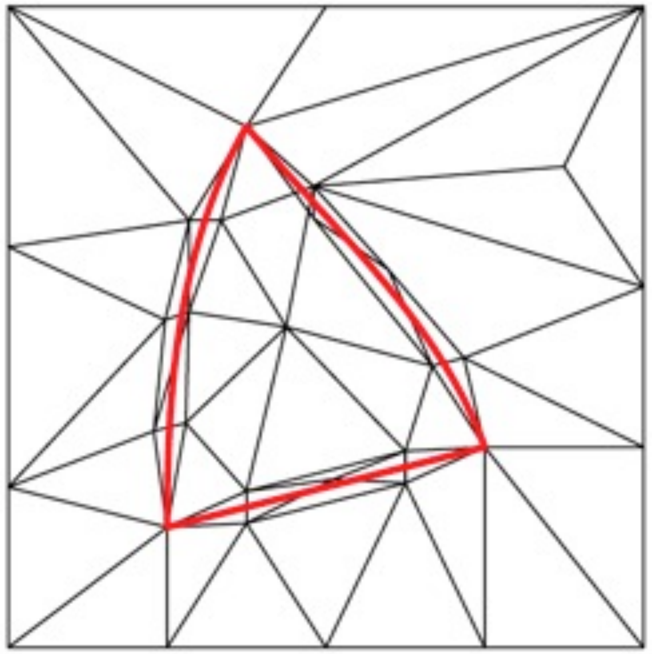

Approximation with Triangulation

Triangulation $(\mathcal{V}, \mathcal{F})$:
$$\left| \begin{array}{l} \text{Vertices } \mathcal{V} = \{v_i\}_{i=1}^M. \\ \text{Faces } \mathcal{F} \subset \{1, \ldots, M\}^3. \end{array} \right.$$

Piecewise linear approximation: $f_M = \sum_{m=1}^{M} \lambda_m \varphi_m$

$$\lambda = \underset{\mu}{\operatorname{argmin}} \, \| f - \sum_m \mu_m \varphi_m \|$$

$\varphi_m(v_i) = \delta_i^m$ is affine on each face of $\mathcal{F}$.

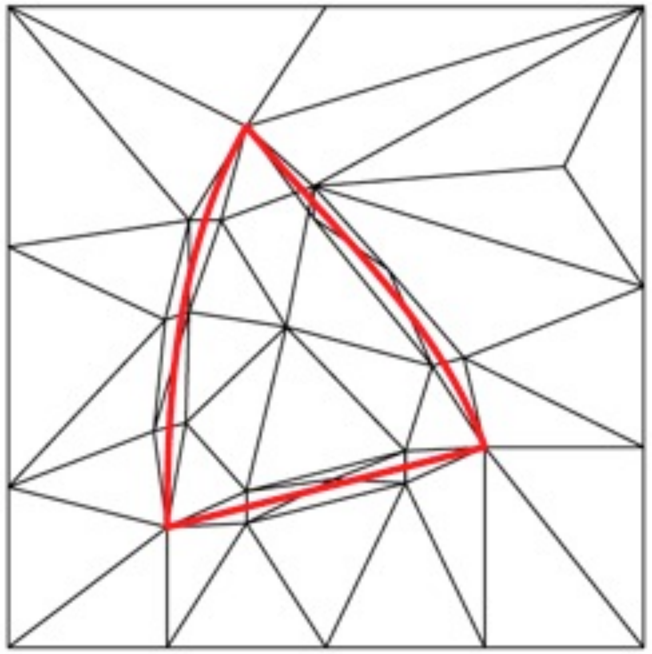

Approximation with Triangulation

Triangulation $(\mathcal{V}, \mathcal{F})$: | Vertices $\mathcal{V} = \{v_i\}_{i=1}^{M}$. Faces $\mathcal{F} \subset \{1, \ldots, M\}^3$.

Piecewise linear approximation: $f_M = \sum_{m=1}^{M} \lambda_m \varphi_m$

$$\lambda = \underset{\mu}{\operatorname{argmin}} \|f - \sum_m \mu_m \varphi_m\|$$

$\varphi_m(v_i) = \delta_i^m$ is affine on each face of $\mathcal{F}$.

> *Theorem:*
>
> There exists $(\mathcal{V}, \mathcal{F})$ such that
>
> $$\|f - f_M\| \leqslant C_f M^{-2}$$

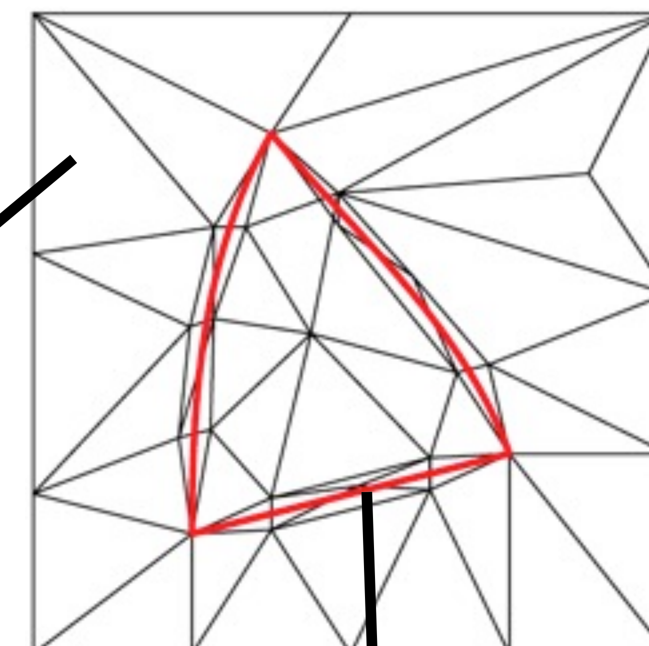

Regular areas:
$\sim M/2$ equilateral triangles.

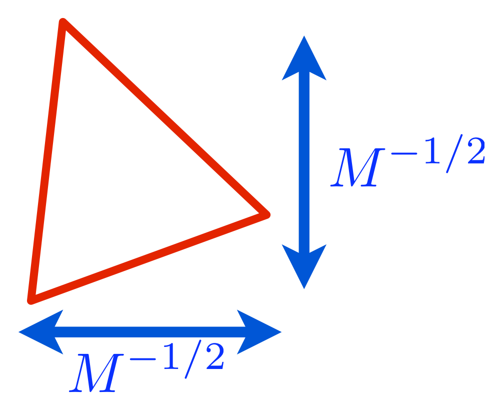

Singular areas:
$\sim M/2$ anisotropic triangles.

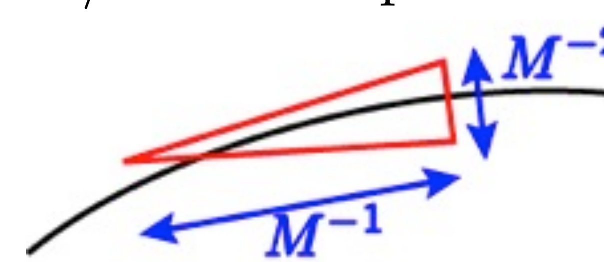

Approximation with Triangulation

Triangulation $(\mathcal{V}, \mathcal{F})$: $\left|\begin{array}{l} \text{Vertices } \mathcal{V} = \{v_i\}_{i=1}^M. \\ \text{Faces } \mathcal{F} \subset \{1, \ldots, M\}^3. \end{array}\right.$

Piecewise linear approximation: $f_M = \sum_{m=1}^{M} \lambda_m \varphi_m$

$$\lambda = \underset{\mu}{\operatorname{argmin}} \; \|f - \sum_m \mu_m \varphi_m\|$$

$\varphi_m(v_i) = \delta_i^m$ is affine on each face of $\mathcal{F}$.

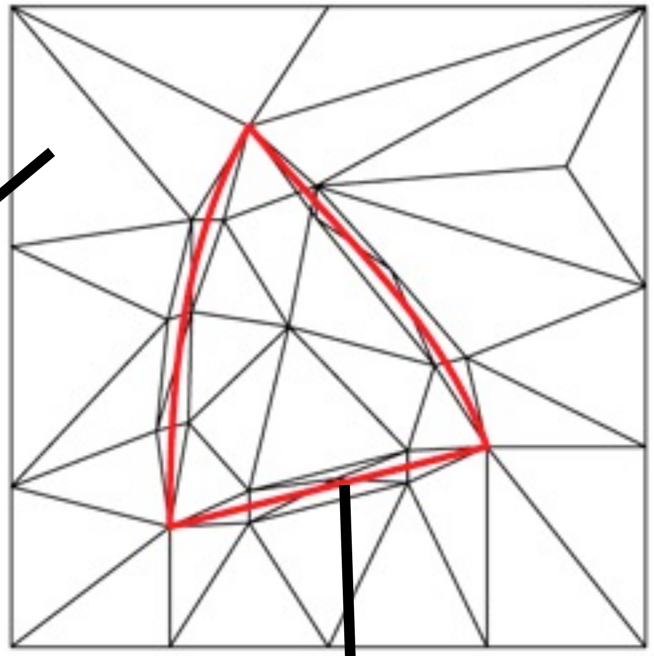

Theorem:

There exists $(\mathcal{V}, \mathcal{F})$ such that

$$\|f - f_M\| \leqslant C_f M^{-2}$$

Regular areas:

$\sim M/2$ equilateral triangles.

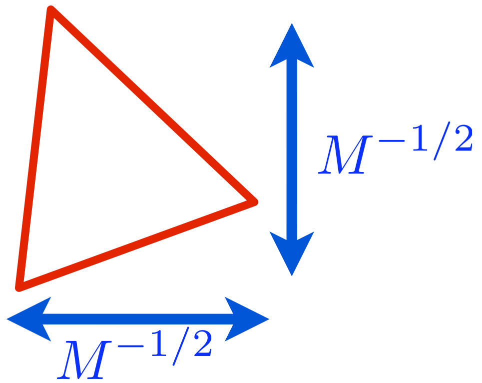

Optimal $(\mathcal{V}, \mathcal{F})$: NP-hard.

Provably good greedy schemes:

[Mirebeau, Cohen, 2009]

Singular areas:

$\sim M/2$ anisotropic triangles.

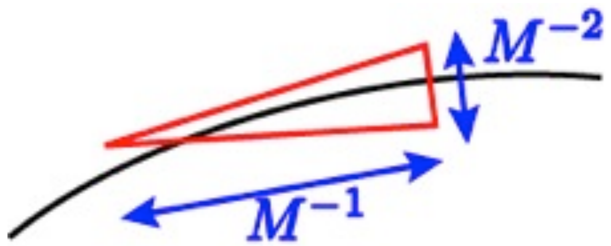

Greedy Triangulation Optimization

Bougleux, Peyré, Cohen, ICCV'09

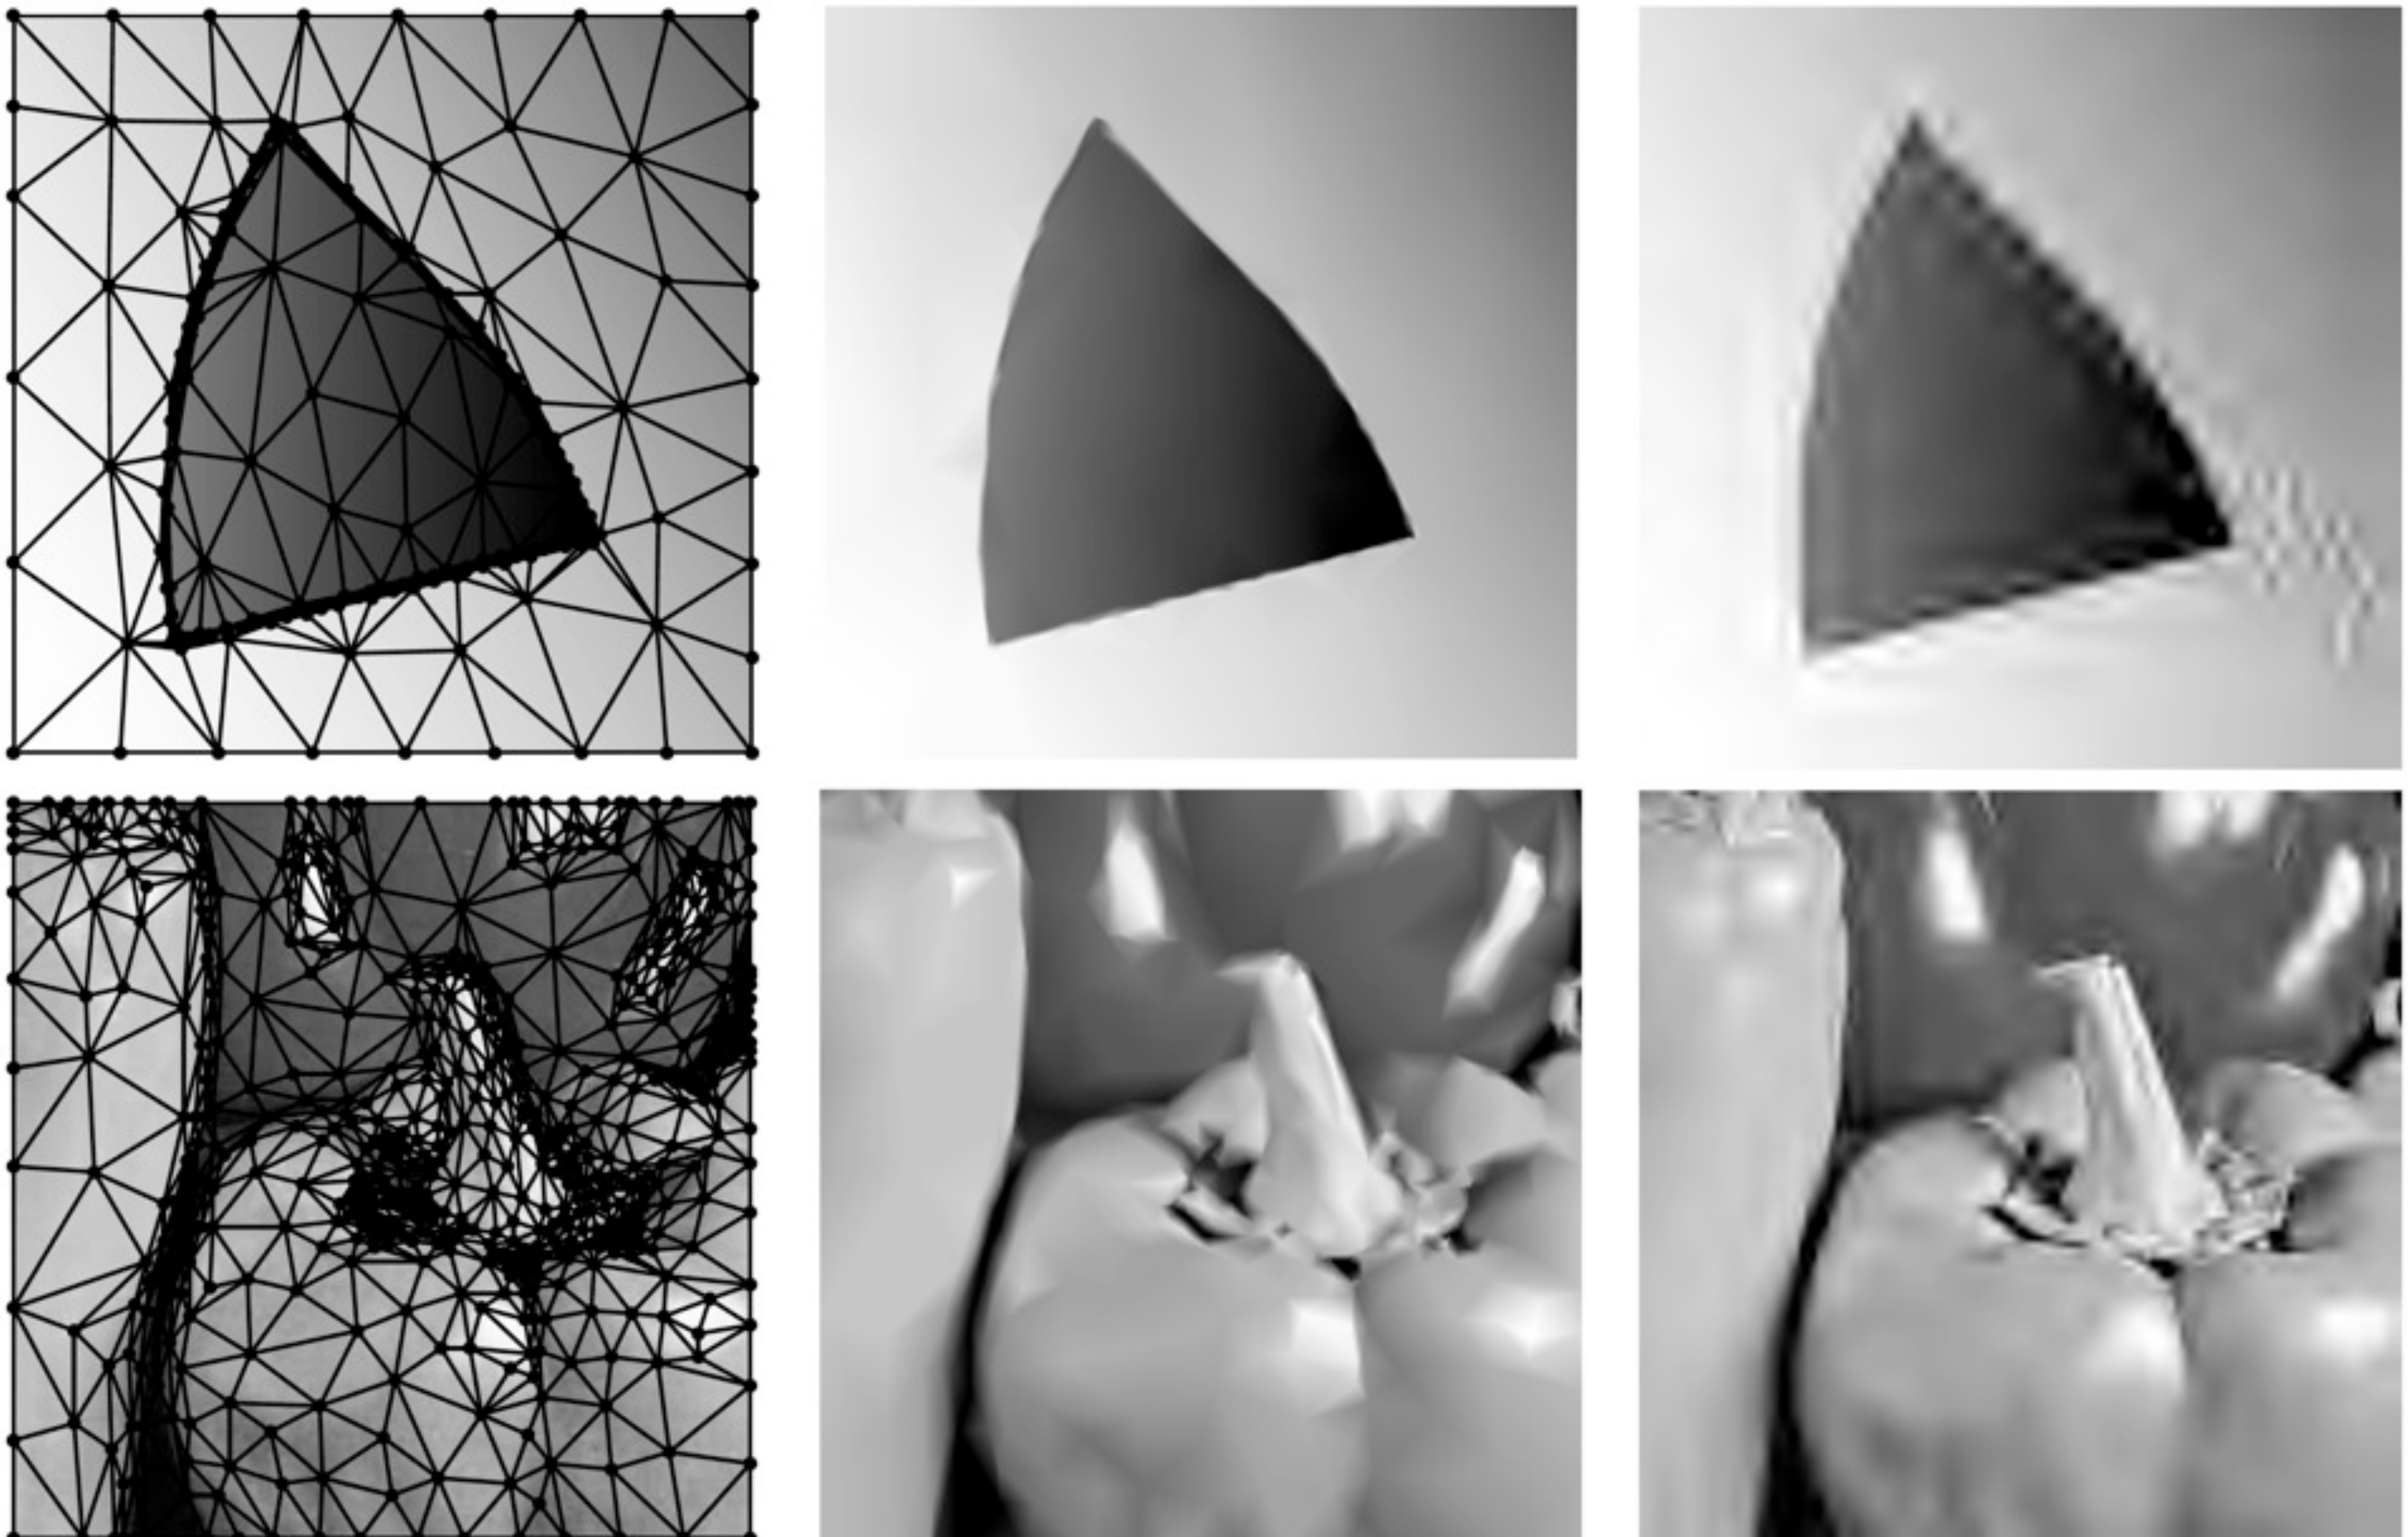

$M = 200$

$M = 600$

Anisotropic triangulation

JPEG2000

Overview

- Sparsity for Approximation
- Sparsity for Processing
- Geometric Images
- Adaptive Geometric Processing
- Adaptive Inverse Problems Regularization
- Geometric Texture Synthesis

Geometric Multiscale Processing

Image $f \in \mathbb{R}^N$

Wavelet transform

Wavelet coefficients

$$f_j^\omega[n] = \langle f,\, \psi_{j,n}^\omega \rangle$$

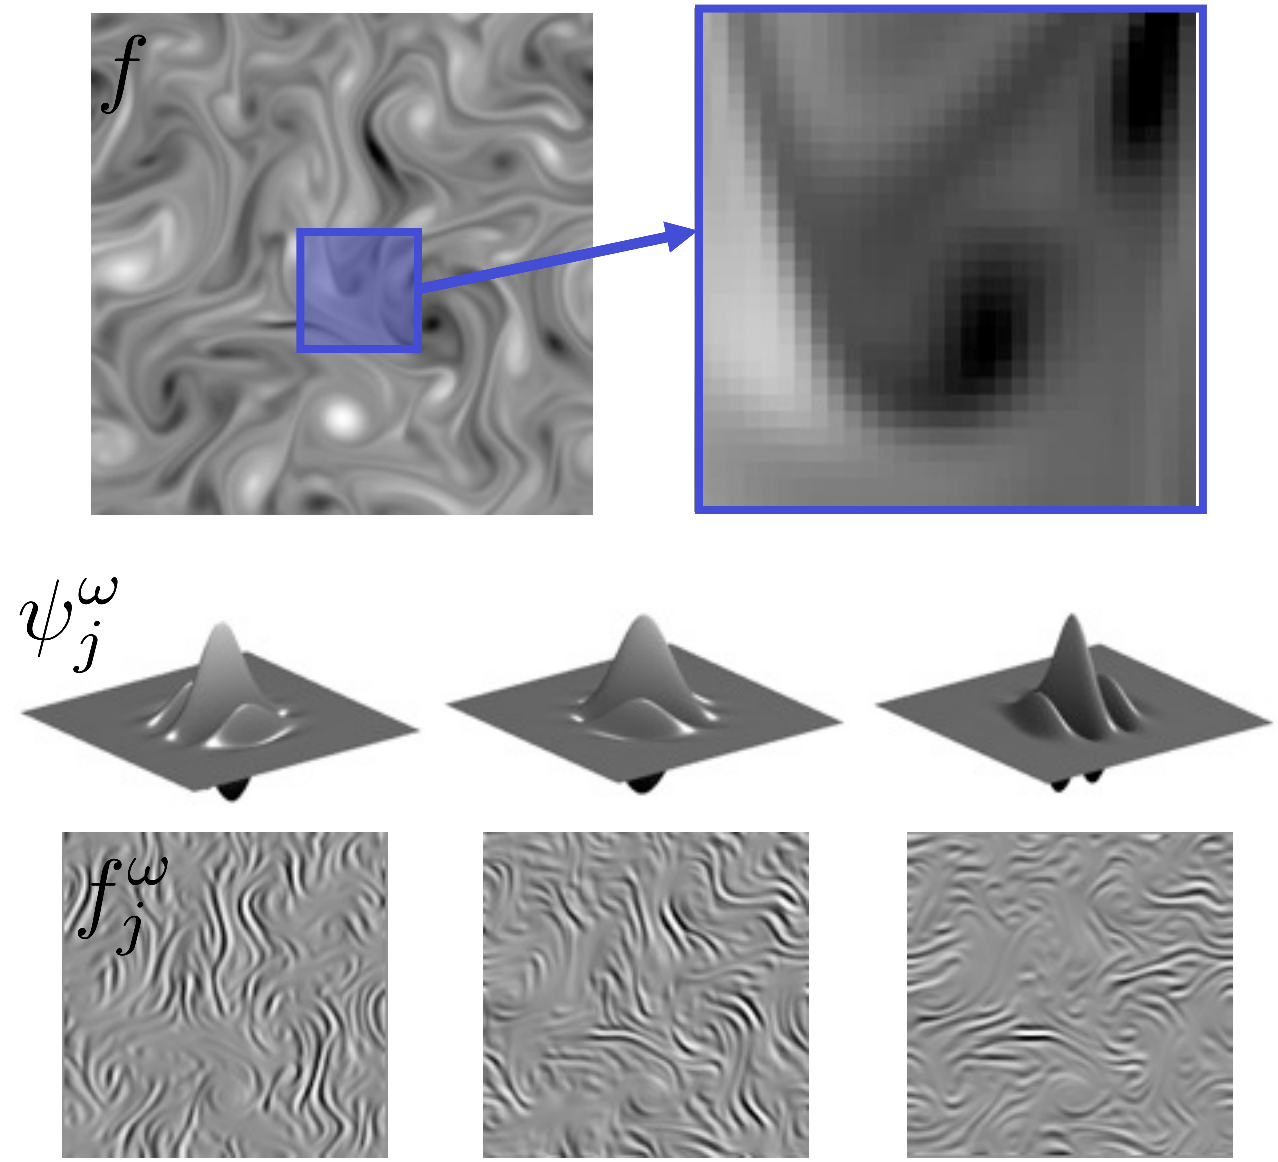

Geometric Multiscale Processing

Image $f \in \mathbb{R}^N$

Wavelet transform

Wavelet coefficients

$$f_j^\omega[n] = \langle f, \psi_{j,n}^\omega \rangle$$

Geometry

λ

Geometric transform

Geometric coefficients

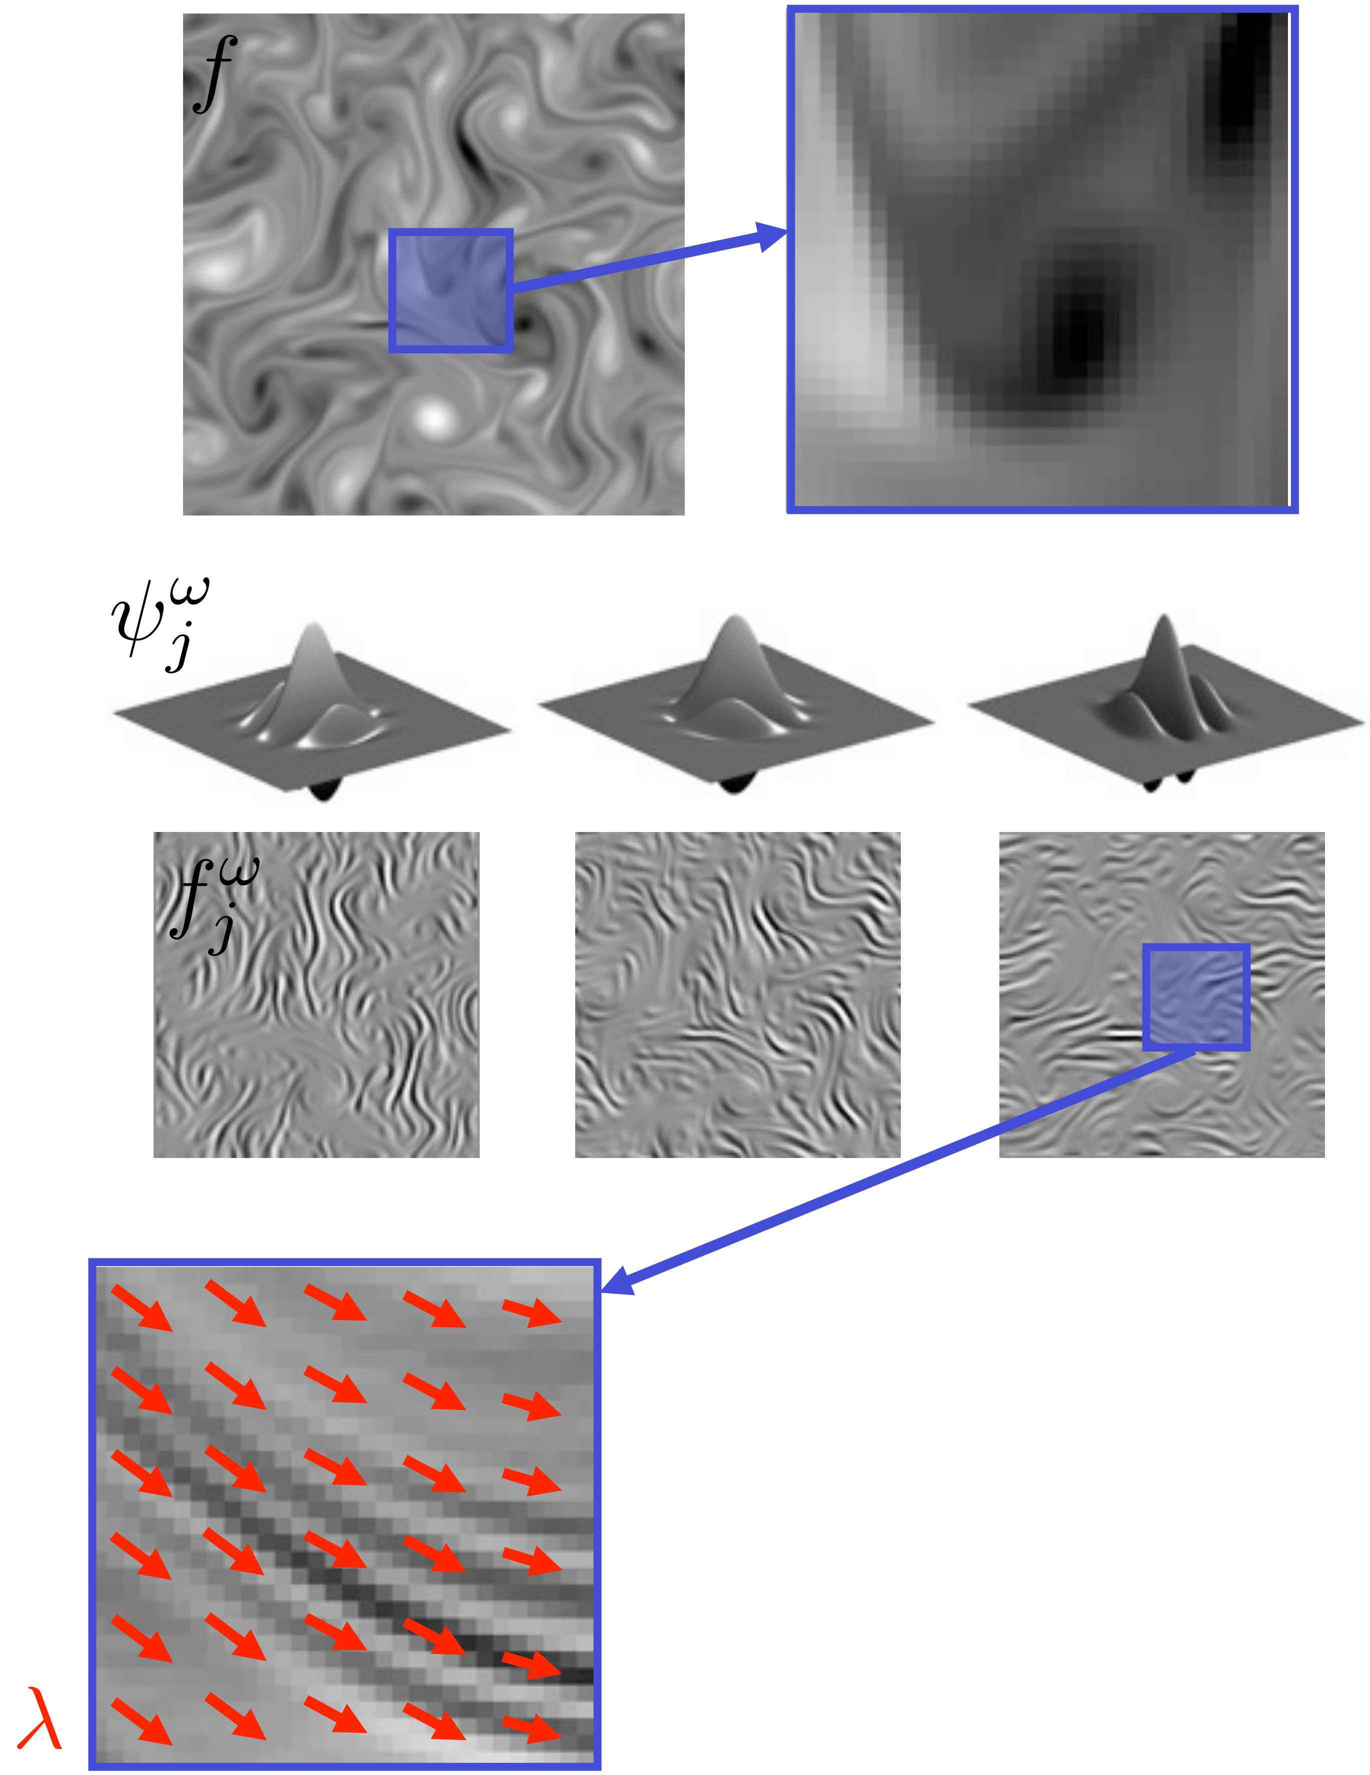

Geometric Multiscale Processing

Image $f \in \mathbb{R}^N$

Wavelet transform

Wavelet coefficients

$$f_j^\omega[n] = \langle f,\, \psi_{j,n}^\omega \rangle$$

Geometry

λ

Geometric transform

Geometric coefficients

$$b[j,\ell,n] = \langle f_j^\omega,\, b_{\ell,n}^\lambda \rangle$$

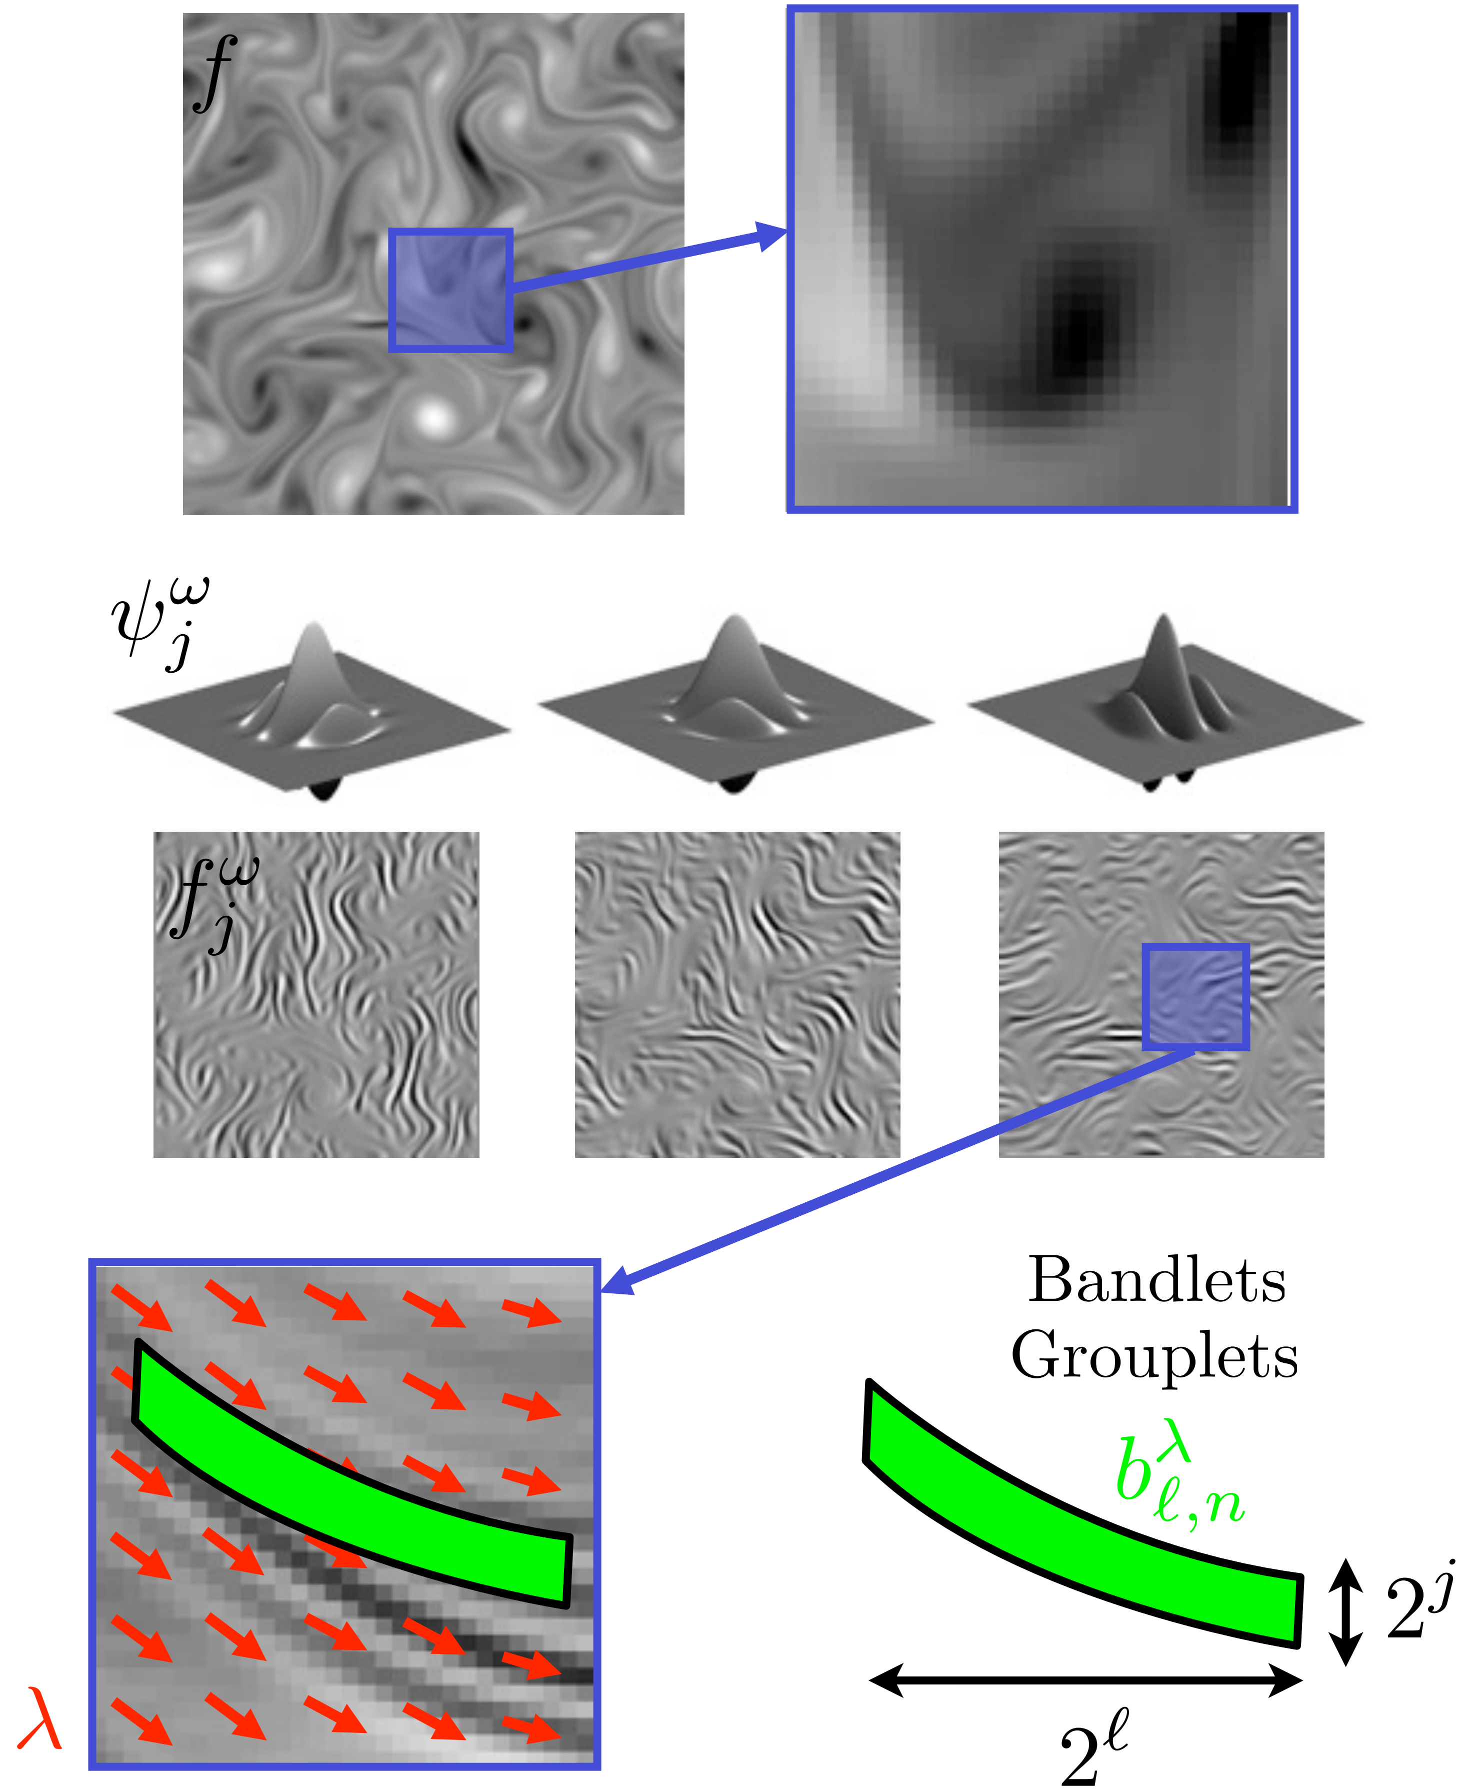

Bandlets vs. Grouplets

Bandlets: [Le Pennec, Mallat, 2005]
[Mallat, Peyré, 2007]

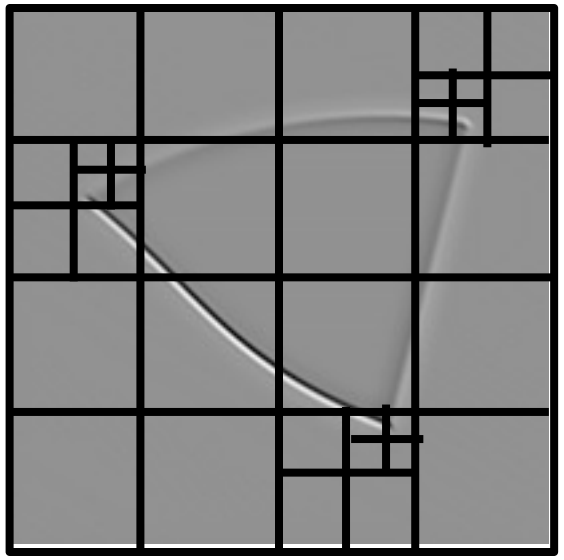

Bandlets vs. Grouplets

Bandlets: [Le Pennec, Mallat, 2005]
[Mallat, Peyré, 2007]

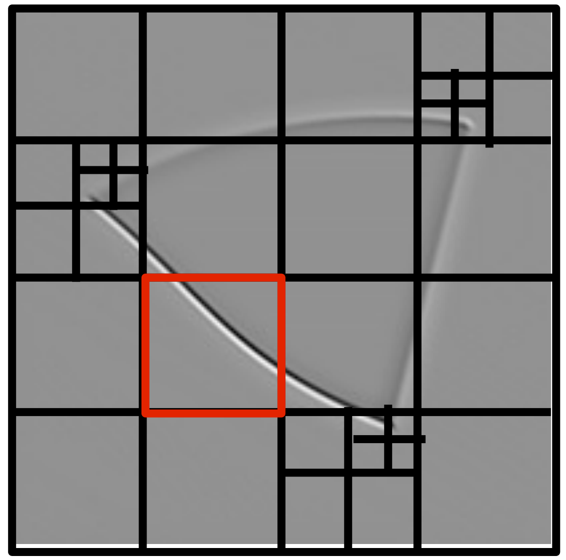

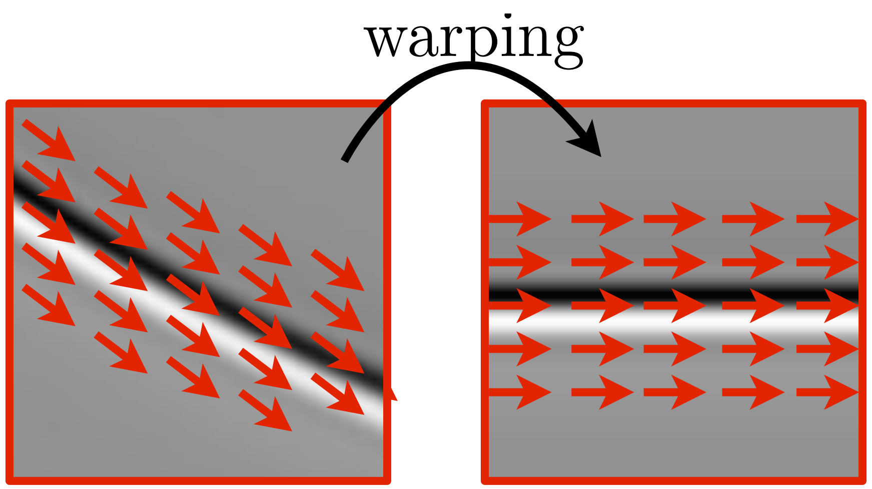

Bandlets vs. Grouplets

Bandlets: [Le Pennec, Mallat, 2005]
[Mallat, Peyré, 2007]

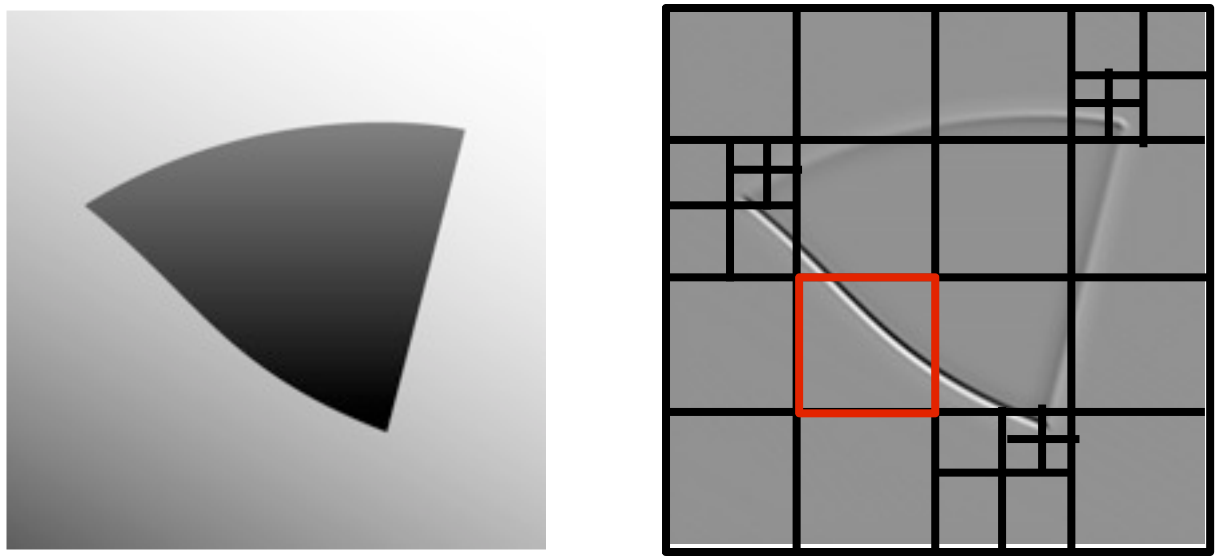

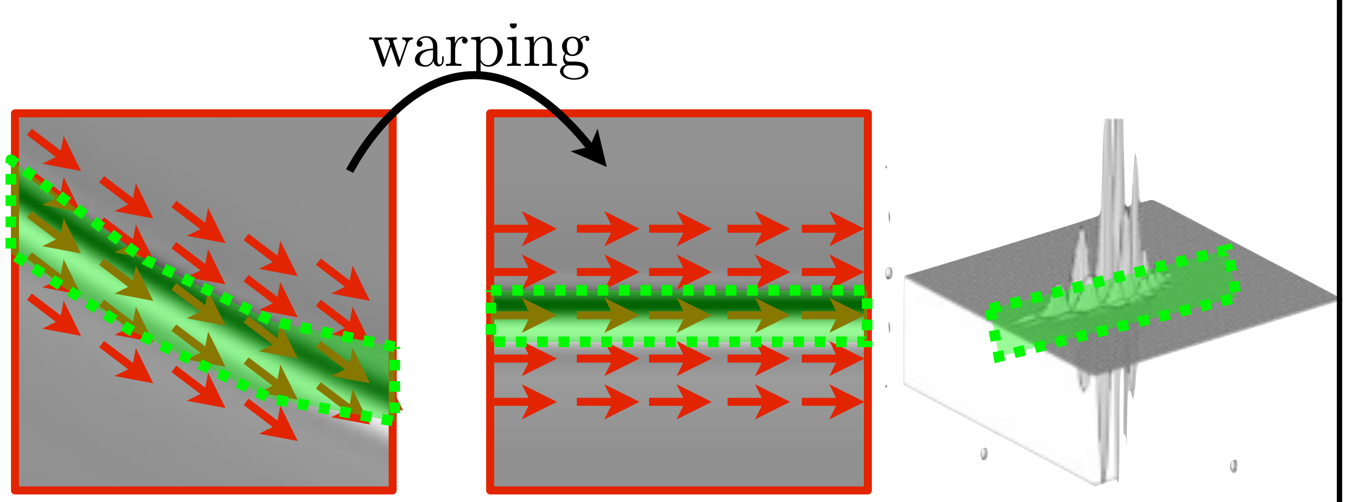

Bandlets vs. Grouplets

Bandlets: [Le Pennec, Mallat, 2005]
[Mallat, Peyré, 2007]

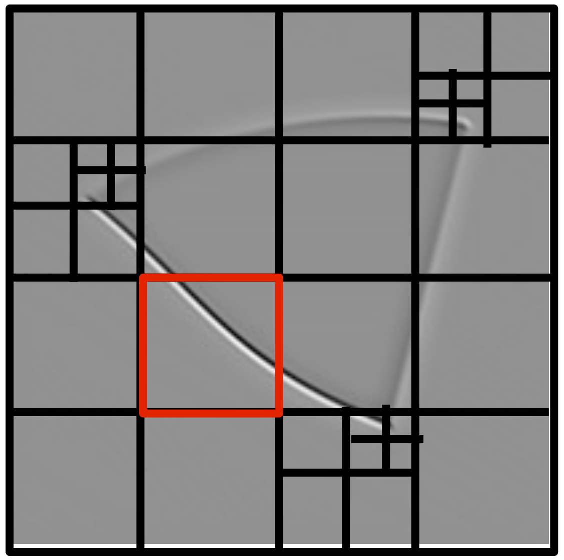

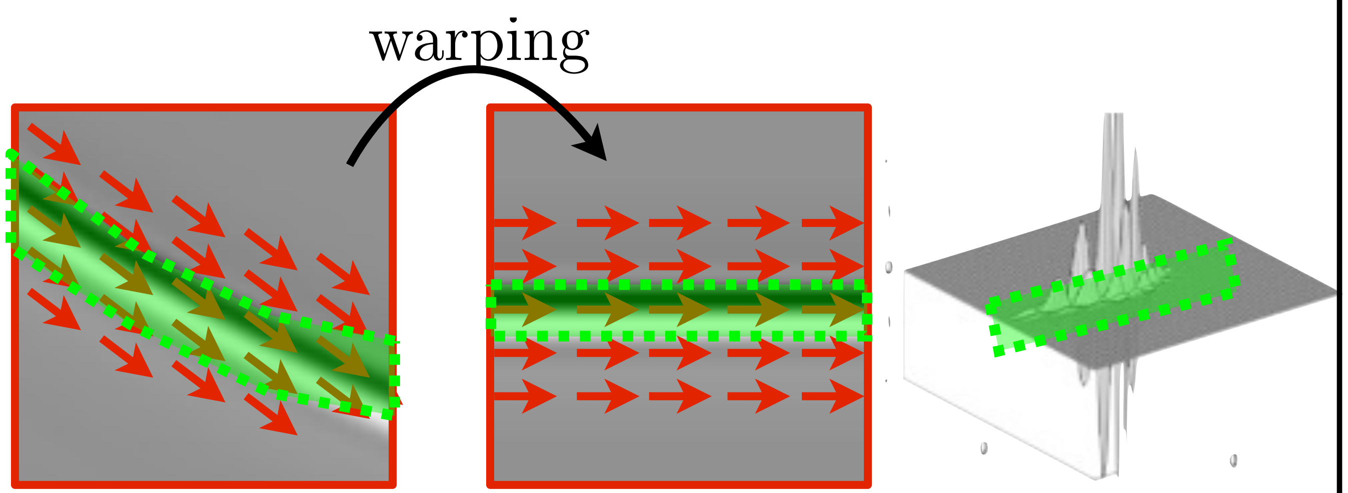

Structured set of bases (quadtrees):
$\rightarrow$ fast best-basis search algorithm.

Bandlets vs. Grouplets

Bandlets: [Le Pennec, Mallat, 2005]
[Mallat, Peyré, 2007]

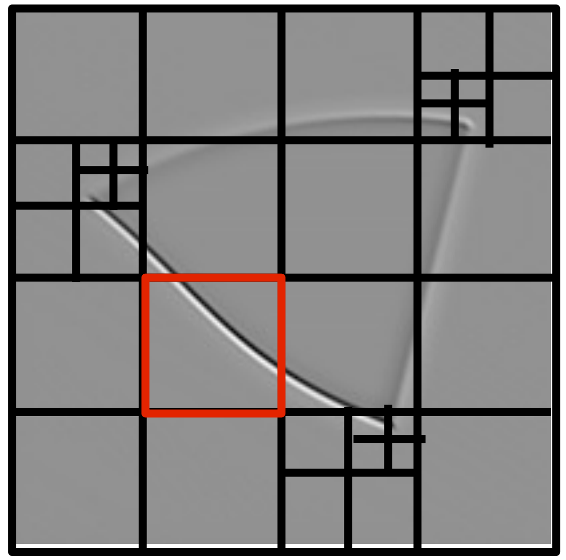

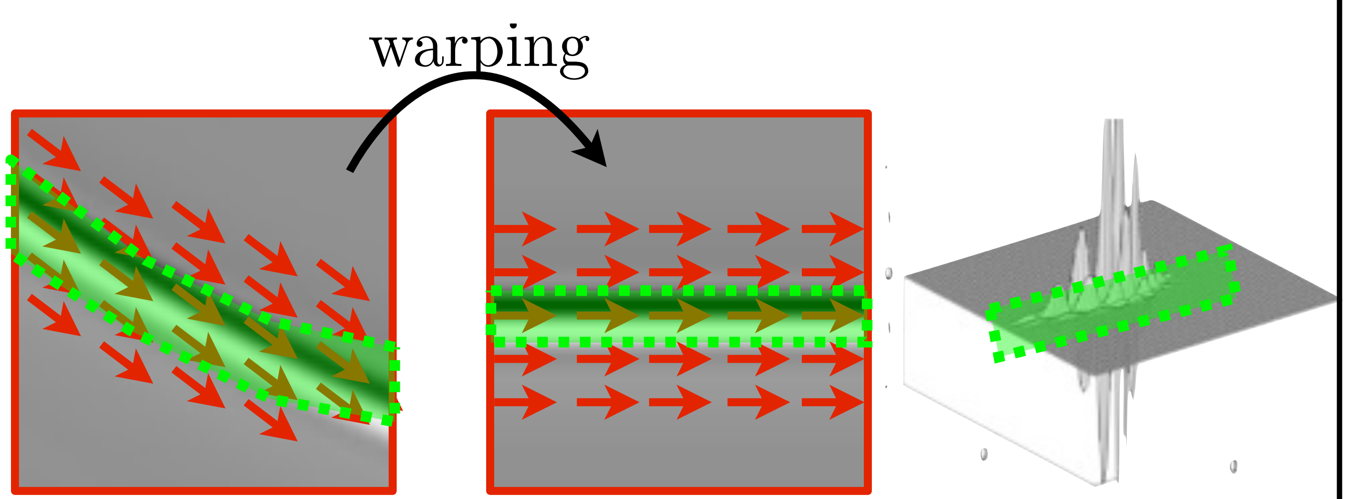

Structured set of bases (quadtrees):
$\longrightarrow$ fast best-basis search algorithm.

Approximation of a $\mathcal{C}^\alpha$ cartoon image:

$$\|f - f_M\|^2 = O(M^{-\alpha}) \qquad M = M_{\text{band}} + M_\lambda$$

Bandlets vs. Grouplets

Bandlets: [Le Pennec, Mallat, 2005] [Mallat, Peyré, 2007]

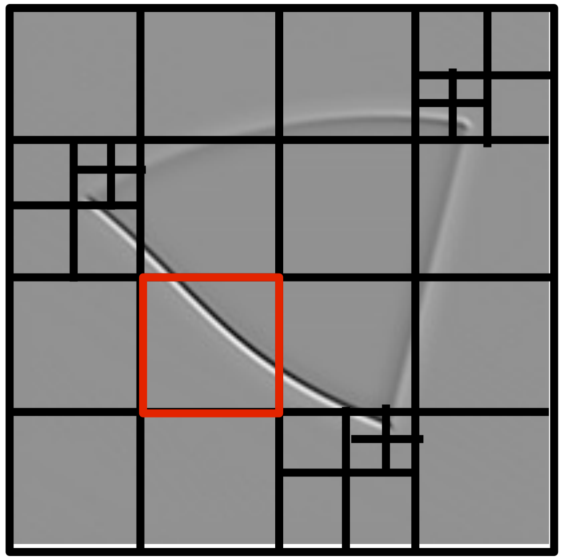

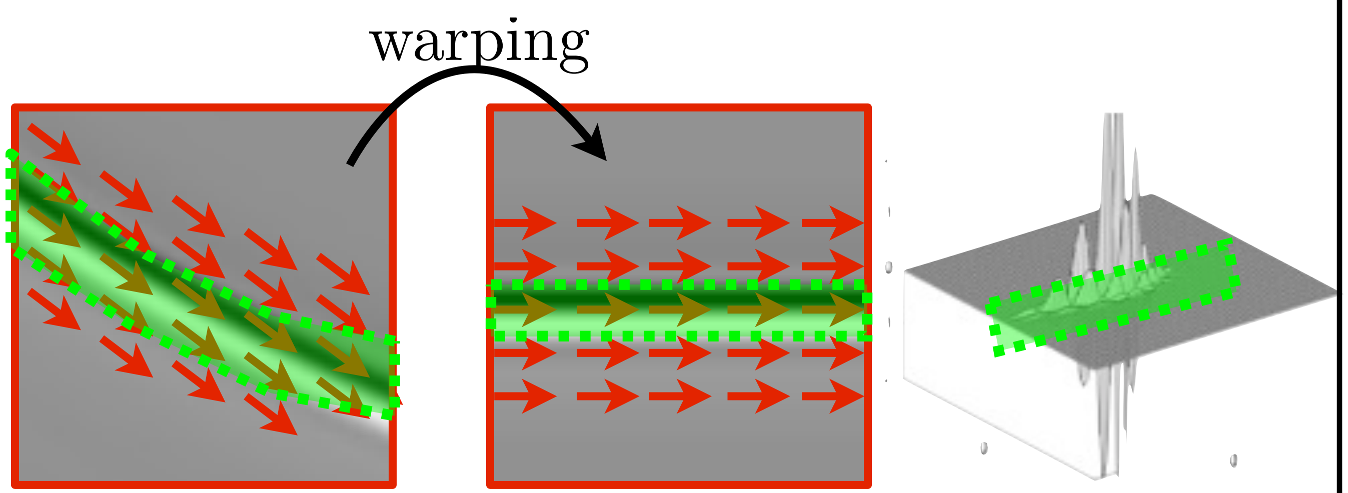

Structured set of bases (quadtrees):
$\longrightarrow$ fast best-basis search algorithm.

Approximation of a $\mathcal{C}^\alpha$ cartoon image:

$$\|f - f_M\|^2 = O(M^{-\alpha}) \qquad M = M_{\text{band}} + M_\lambda$$

Grouplets: [Mallat, 2009] [Peyré, 2010]

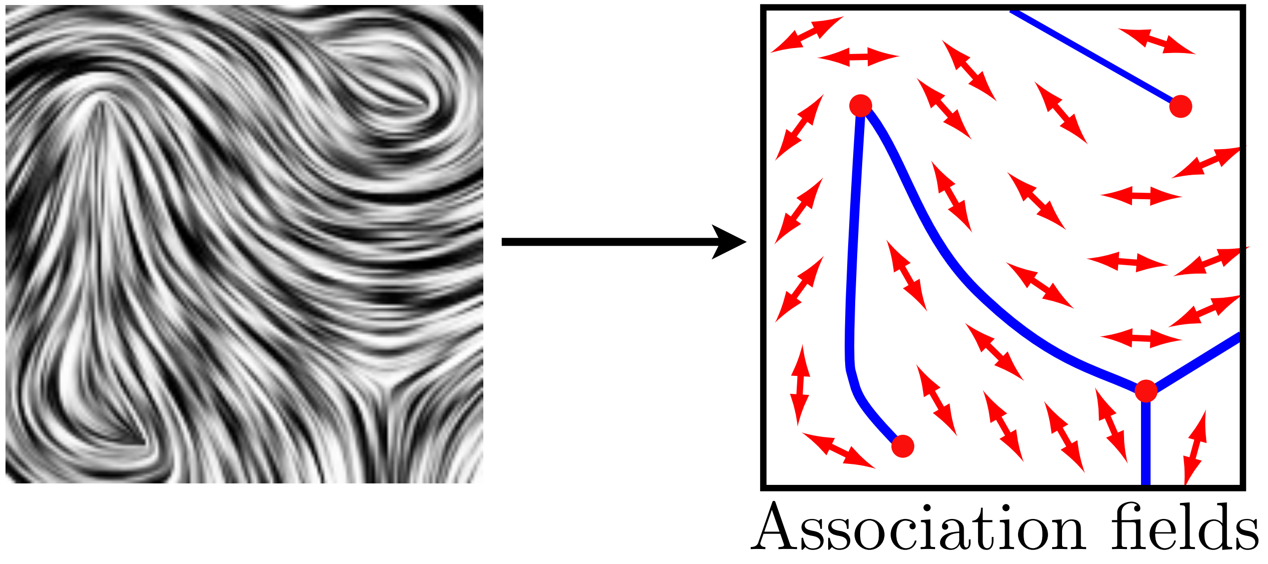

Bandlets vs. Grouplets

Bandlets: [Le Pennec, Mallat, 2005]
[Mallat, Peyré, 2007]

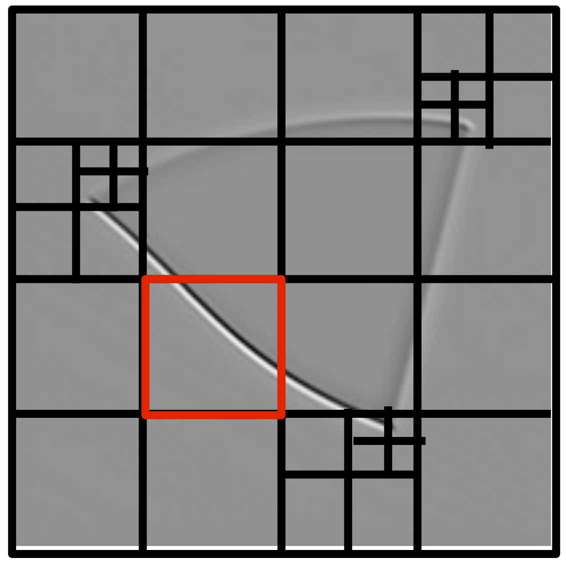

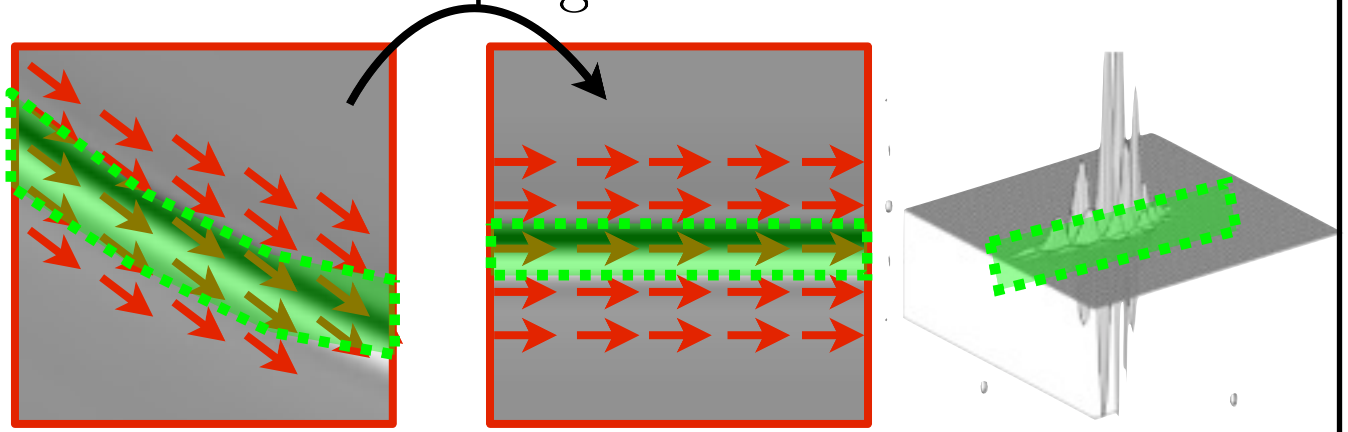

Structured set of bases (quadtrees):
$\longrightarrow$ fast best-basis search algorithm.

Approximation of a $\mathcal{C}^\alpha$ cartoon image:

$$\|f - f_M\|^2 = O(M^{-\alpha}) \qquad M = M_{\text{band}} + M_\lambda$$

Grouplets: [Mallat, 2009] [Peyré, 2010]

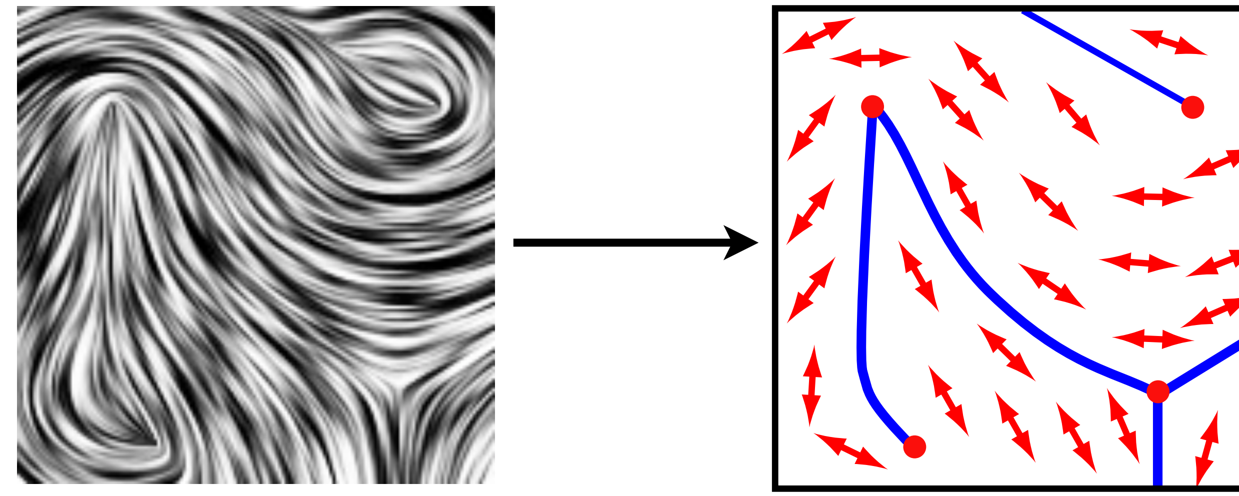

Association fields

Atoms follow the flow λ.

$\ell = 1$ $\quad$ $\ell = 2$ $\quad$ $\ell = 3$ $\quad$ $\ell = 4$

Bandlets vs. GroupLets

Bandlets: [Le Pennec, Mallat, 2005] [Mallat, Peyré, 2007]

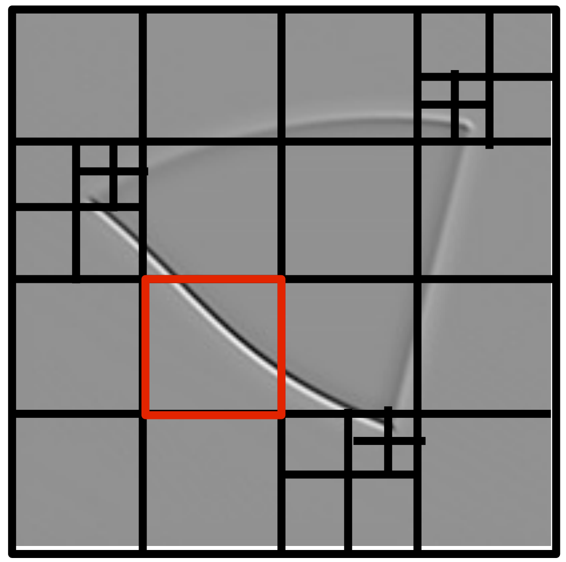

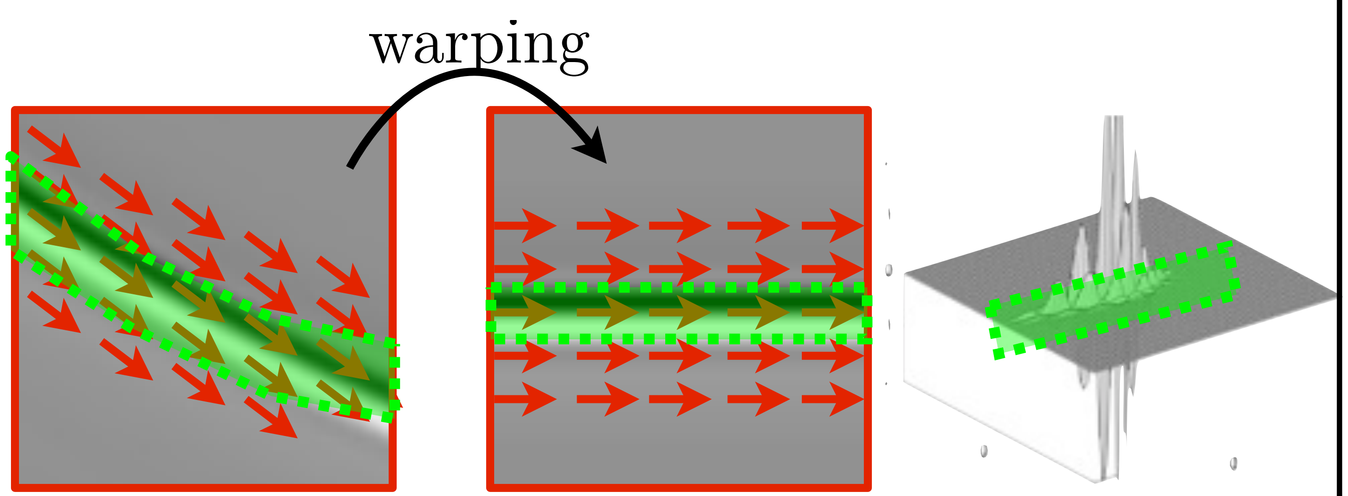

Structured set of bases (quadtrees):
$\rightarrow$ fast best-basis search algorithm.

Approximation of a $\mathcal{C}^\alpha$ cartoon image:

$$\|f - f_M\|^2 = O(M^{-\alpha}) \qquad M = M_{\text{band}} + M_\lambda$$

Grouplets: [Mallat, 2009] [Peyré, 2010]

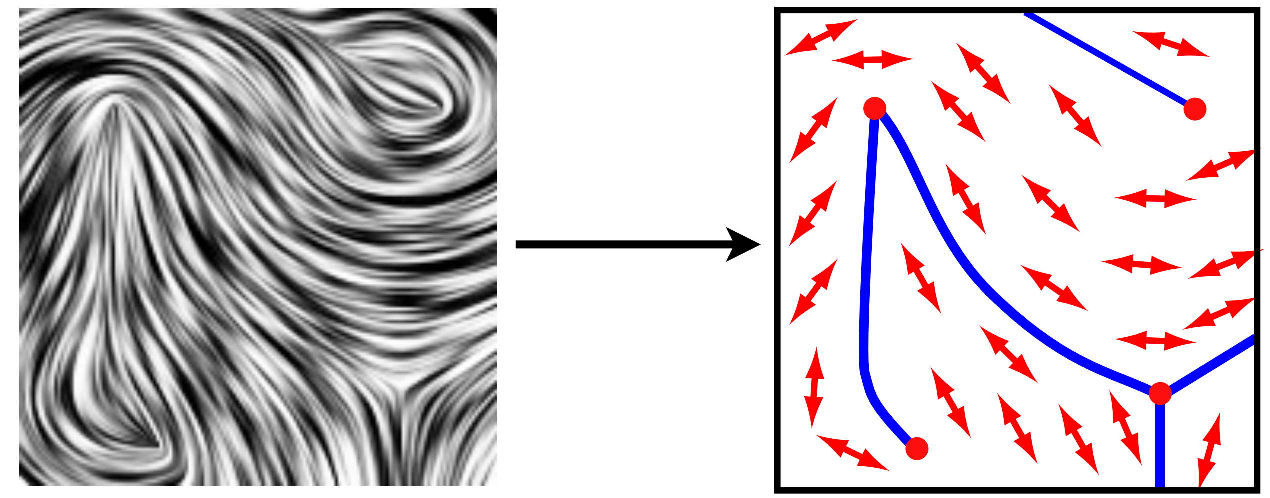

Association fields

Atoms follow the flow λ.

$\ell = 1$ $\ell = 2$ $\ell = 3$ $\ell = 4$

Un-structured set of bases (flows):
$\rightarrow$ sub-optimal optical flow algorithms.

Bandlets vs. Grouplets

Bandlets: [Le Pennec, Mallat, 2005]
[Mallat, Peyré, 2007]

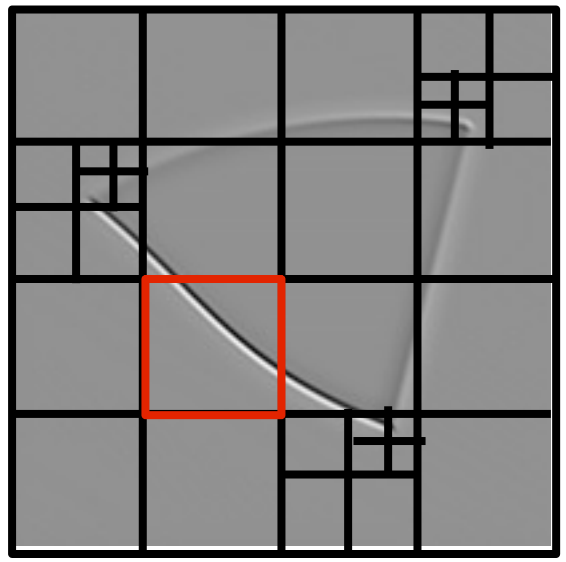

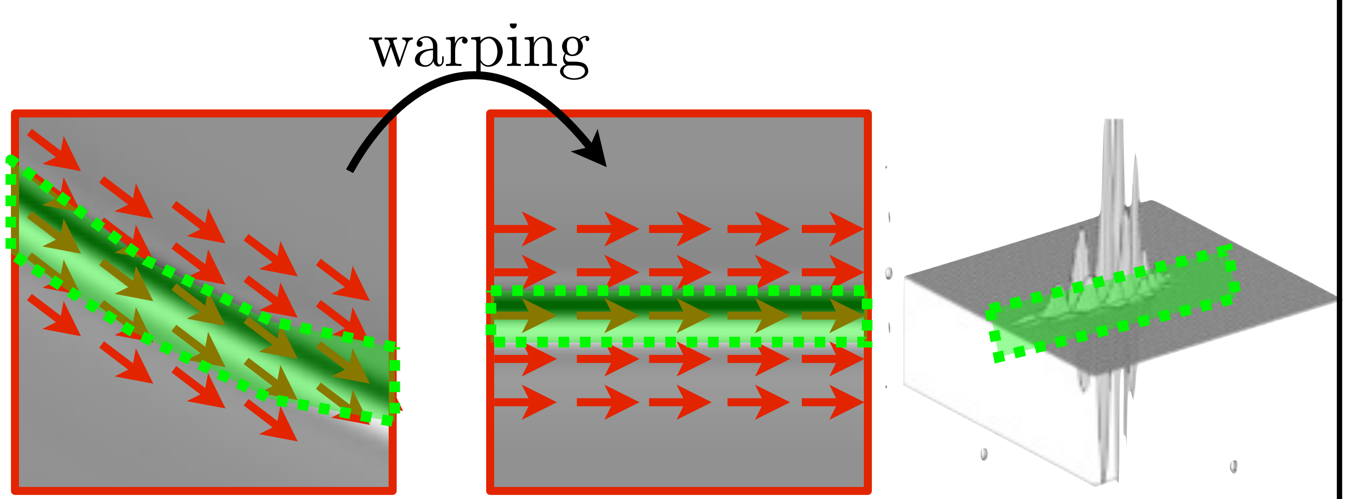

Structured set of bases (quadtrees):
$\longrightarrow$ fast best-basis search algorithm.

Approximation of a $\mathcal{C}^\alpha$ cartoon image:

$$\|f - f_M\|^2 = O(M^{-\alpha}) \qquad M = M_{\text{band}} + M_\lambda$$

Grouplets: [Mallat, 2009] [Peyré, 2010]

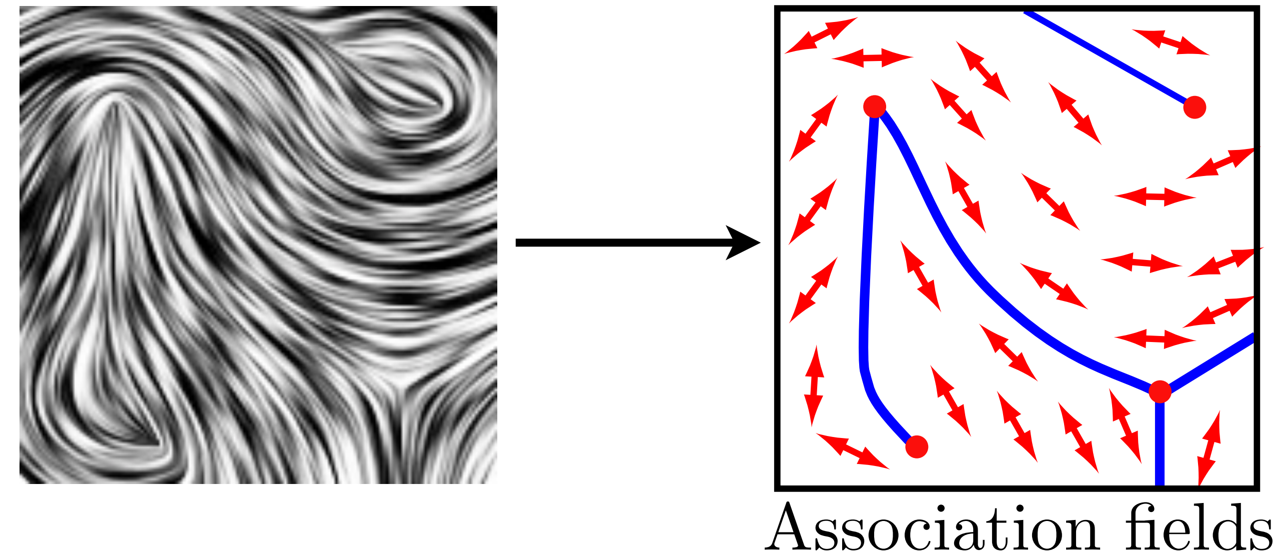

Atoms follow the flow λ.

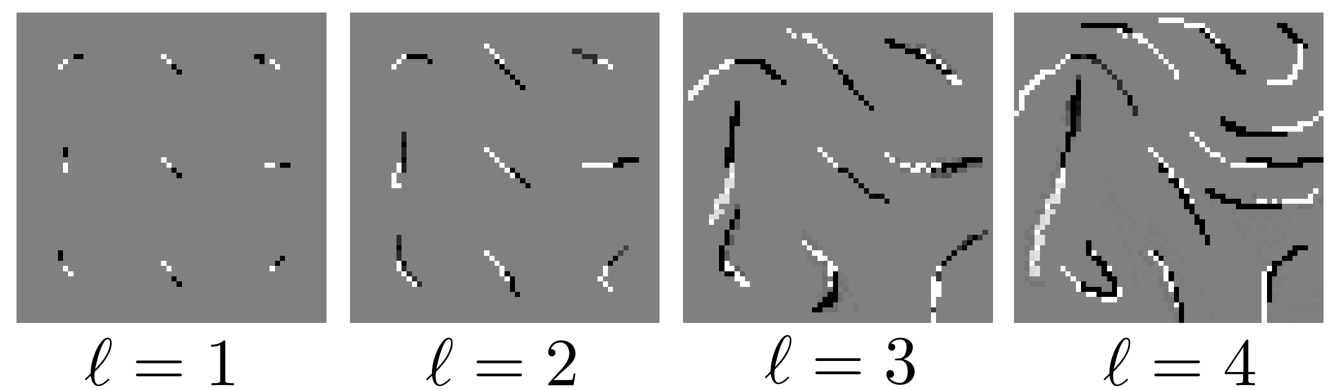

Un-structured set of bases (flows):
$\longrightarrow$ sub-optimal optical flow algorithms.

No optimality results for approximation.

Better processing of textures.

Bandlet Compression and Denoising

Original

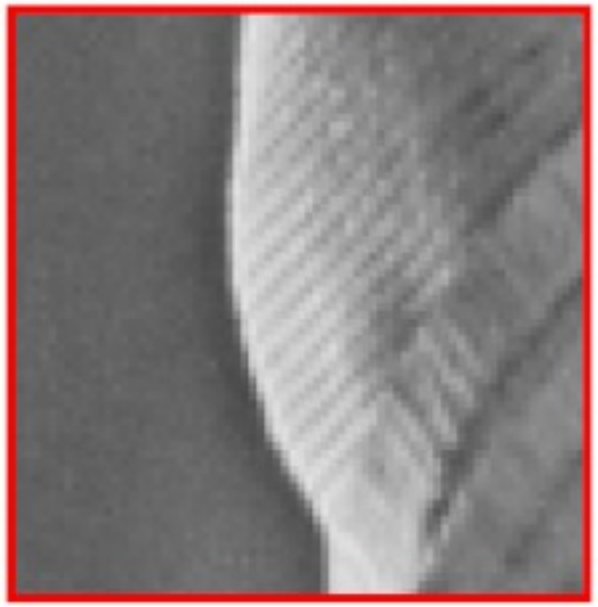

Wavelets

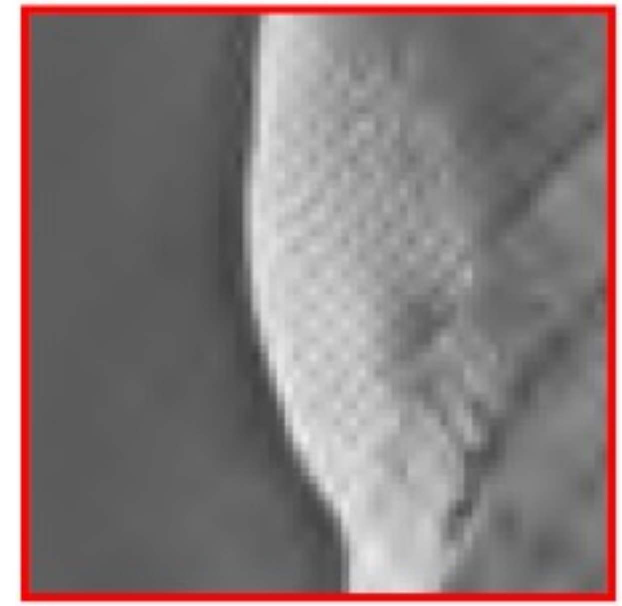

Bandlets

Bandlet Compression and Denoising

Original

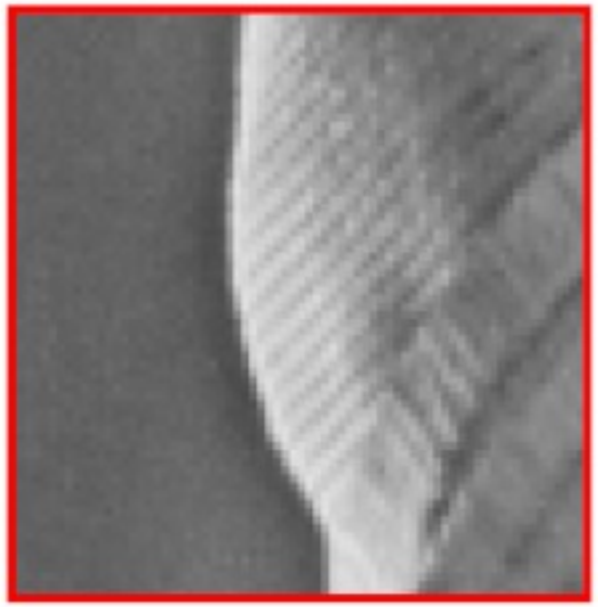

Wavelets

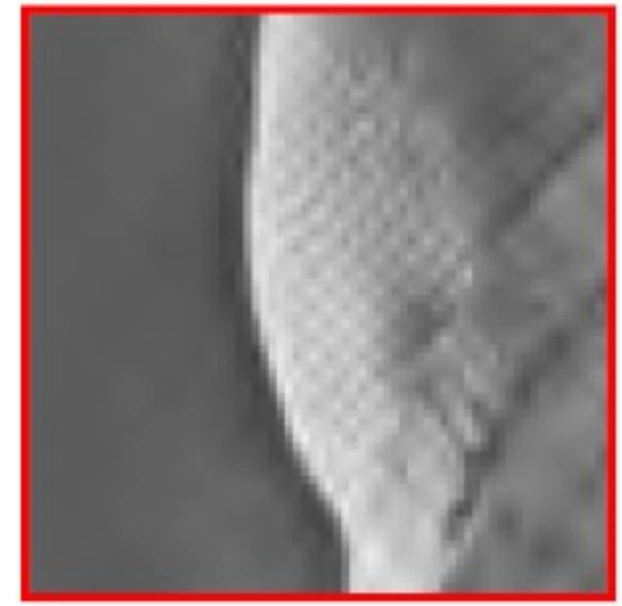

Bandlets

CNES case study (X. Delaunay PhD thesis):
On board low complexity bandlet-like coder.
$\longrightarrow$ Local grouping of wavelet coefficients.

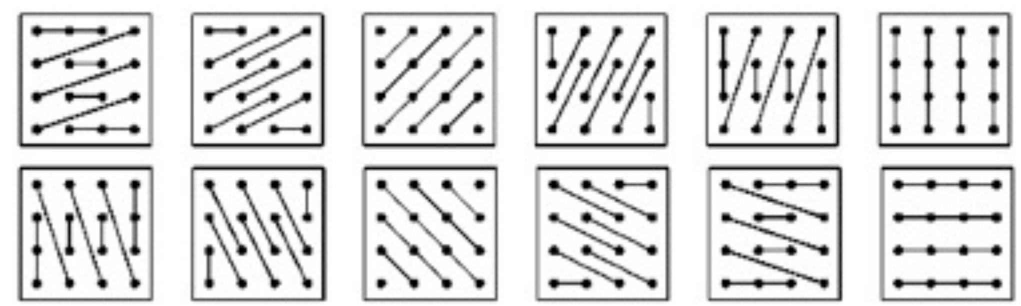

Bandlet Compression and Denoising

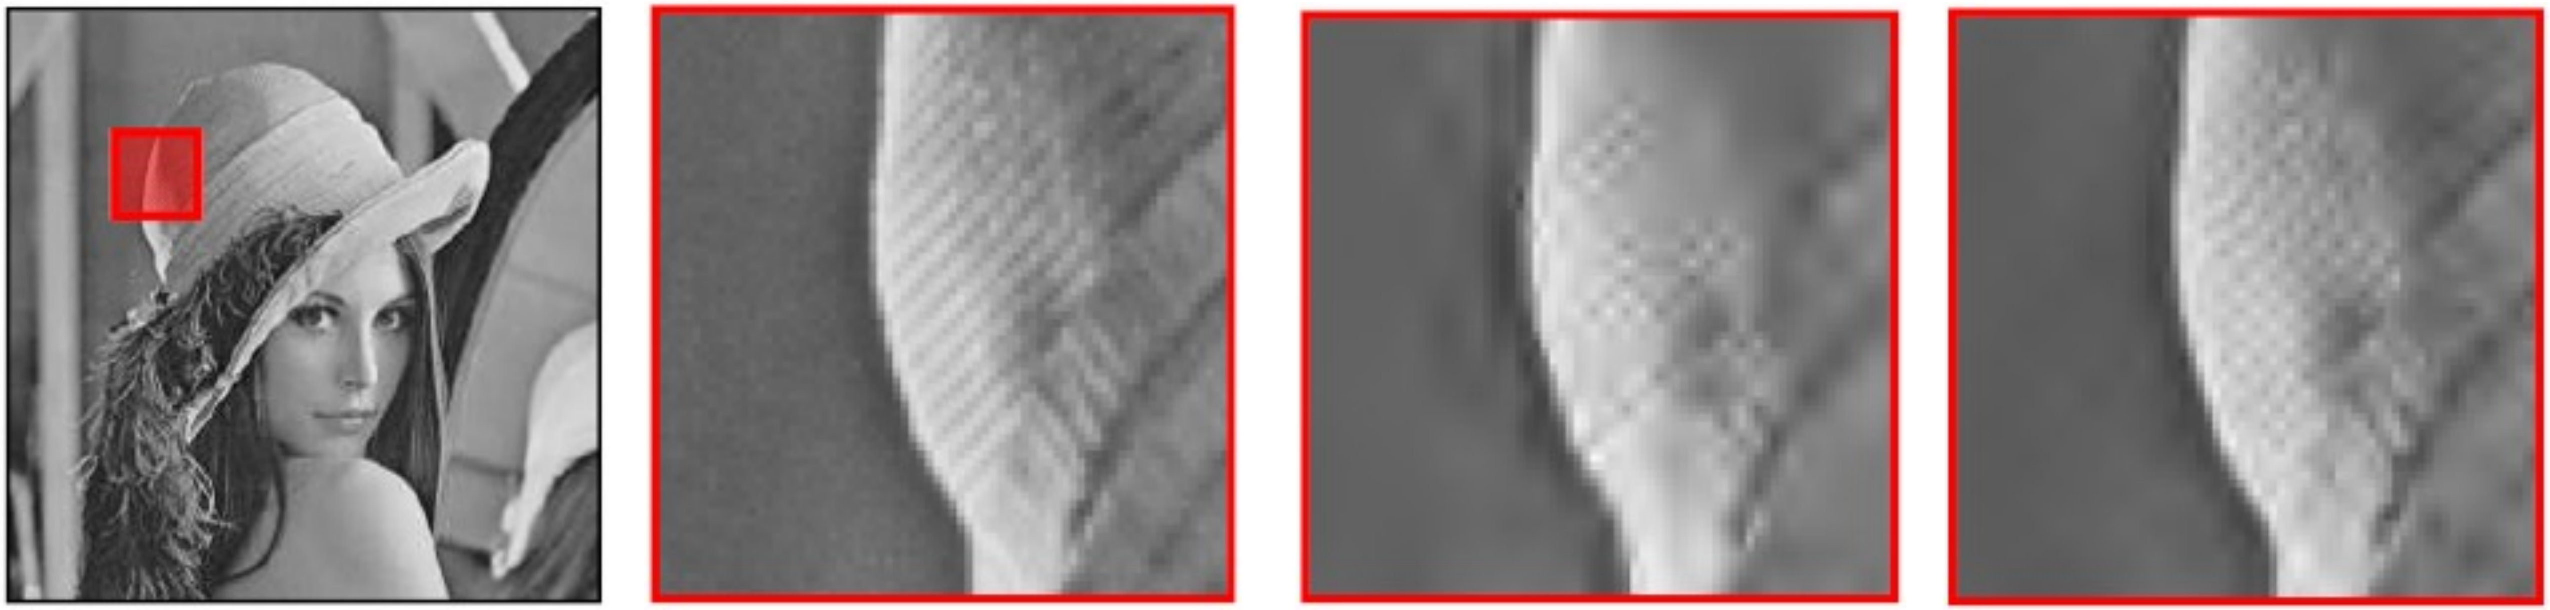

Original Wavelets Bandlets

CNES case study (X. Delaunay PhD thesis):
On board low complexity bandlet-like coder.
$\longrightarrow$ Local grouping of wavelet coefficients.

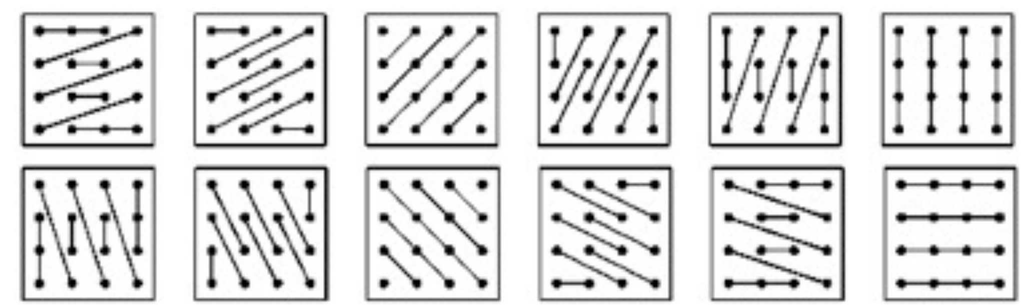

Denoising: λ is estimated from the noisy image $y = f + w$.

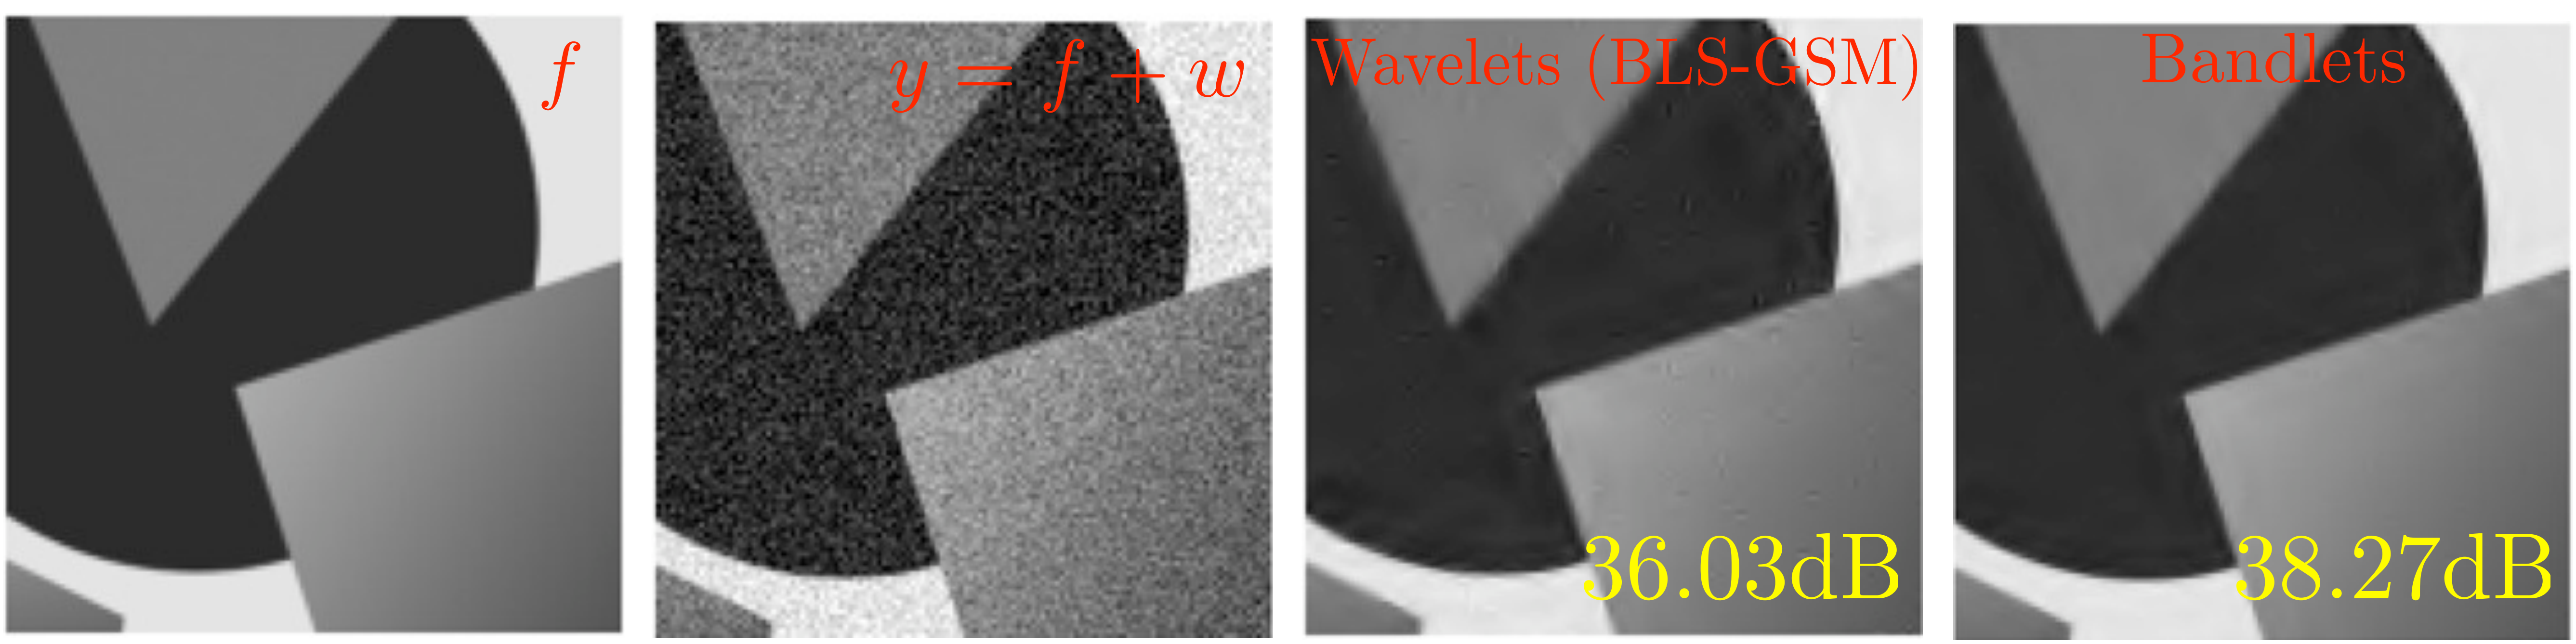

Overview

- Sparsity for Approximation
- Sparsity for Processing
- Geometric Images
- Adaptive Geometric Processing
- Adaptive Inverse Problems Regularization
- Geometric Texture Synthesis

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)

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)

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

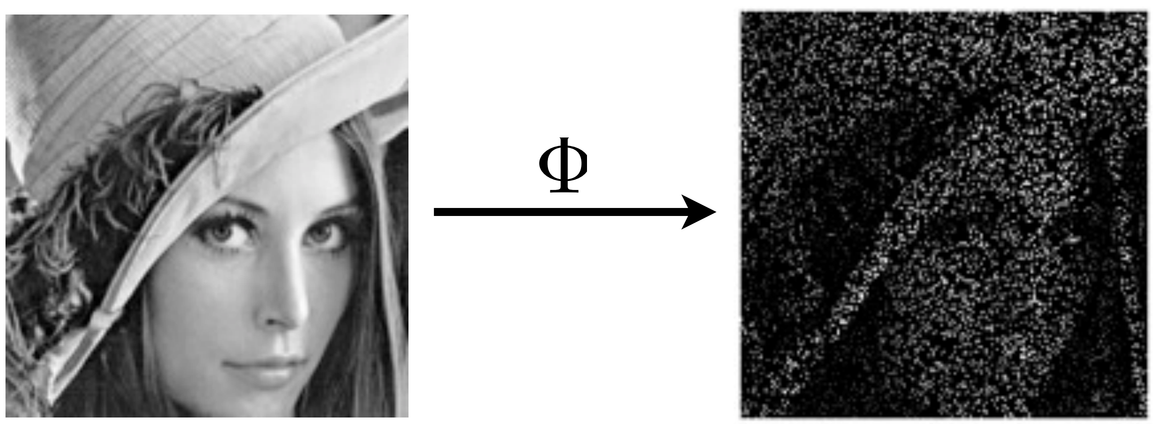

$$(\Phi f)(x) = \begin{cases} 0 & \text{if} \quad x \in \Omega, \\ f(x) & \text{if} \quad x \notin \Omega. \end{cases}$$

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)

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

$\xrightarrow{\Phi}$

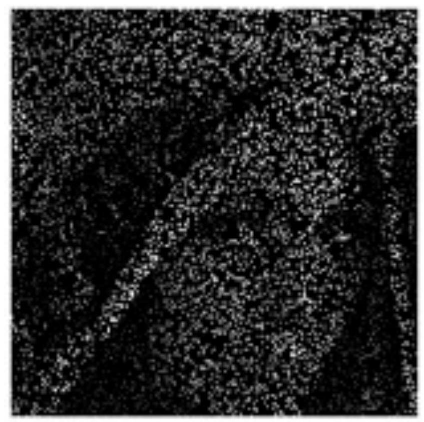

$$(\Phi f)(x) = \begin{cases} 0 & \text{if} \quad x \in \Omega, \\ f(x) & \text{if} \quad x \notin \Omega. \end{cases}$$

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

$\xrightarrow{\Phi}$

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)

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

$\xrightarrow{\Phi}$

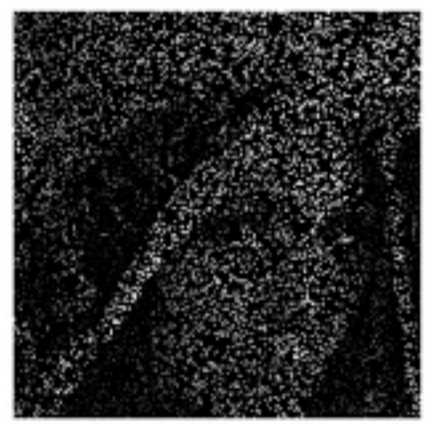

$$(\Phi f)(x) = \begin{cases} 0 & \text{if} \quad x \in \Omega, \\ f(x) & \text{if} \quad x \notin \Omega. \end{cases}$$

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

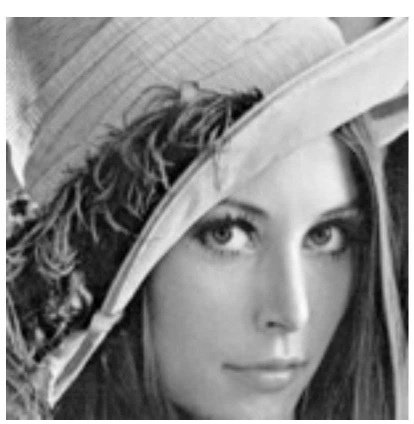

$\xrightarrow{\Phi}$

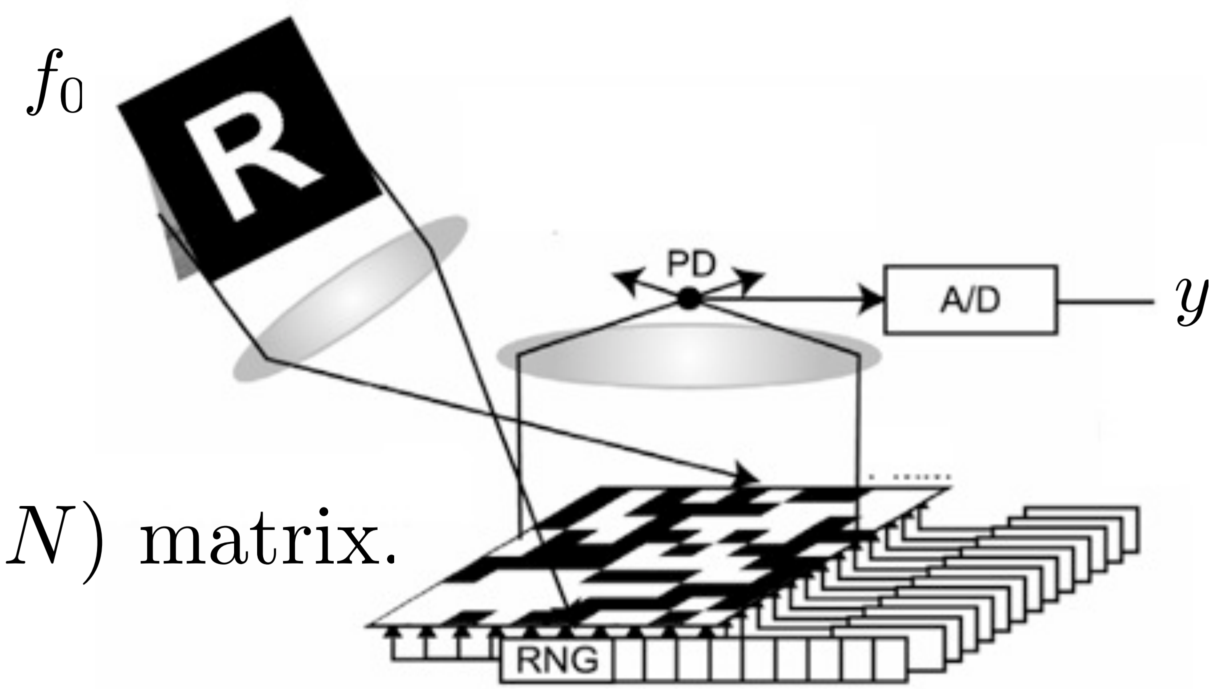

Compressed sensing: $\Phi =$ random (P, N) matrix.

Inverse Problem Regularization

Noisy measurements $y = \Phi f_0 + w \in \mathbb{R}^P, \quad w[n] \sim \mathcal{N}(0, \sigma).$

λ

Inverse Problem Regularization

Noisy measurements $y = \Phi f_0 + w \in \mathbb{R}^P, \quad w[n] \sim \mathcal{N}(0, \sigma)$.

Prior model: $J(f) \in \mathbb{R}$ such that $J(f_0)$ is small for $f_0 \in \Theta$.

Regularized inverse: $$f^\star = \underset{f \in \mathbb{R}^N}{\operatorname{argmin}} \underbrace{\boxed{\frac{1}{2}\|y - \Phi f\|^2}}_{\text{Data fitting}} + \underbrace{\boxed{\lambda J(f)}}_{\text{Regularity}}$$

λ

Inverse Problem Regularization

Noisy measurements $y = \Phi f_0 + w \in \mathbb{R}^P, \quad w[n] \sim \mathcal{N}(0, \sigma).$

Prior model: $J(f) \in \mathbb{R}$ such that $J(f_0)$ is small for $f_0 \in \Theta$.

Regularized inverse: $$f^\star = \underset{f \in \mathbb{R}^N}{\operatorname{argmin}} \underbrace{\boxed{\frac{1}{2}\|y - \Phi f\|^2}}_{\text{Data fitting}} + \underbrace{\boxed{\lambda J(f)}}_{\text{Regularity}}$$

Choice of λ: minimize $\|f^\star - f_0\|$ (oracle)

Trade-off between denoising (λ increases with σ) and regularity of f_0.

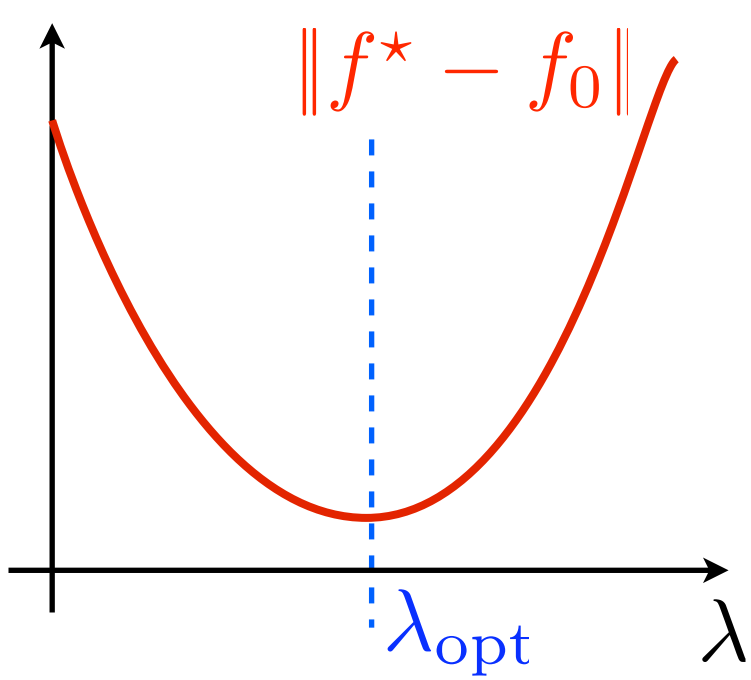

Inverse Problem Regularization

Noisy measurements $y = \Phi f_0 + w \in \mathbb{R}^P$, $\quad w[n] \sim \mathcal{N}(0, \sigma)$.

Prior model: $J(f) \in \mathbb{R}$ such that $J(f_0)$ is small for $f_0 \in \Theta$.

Regularized inverse: $$f^\star = \underset{f \in \mathbb{R}^N}{\operatorname{argmin}} \underbrace{\boxed{\frac{1}{2}\|y - \Phi f\|^2}}_{\text{Data fitting}} + \underbrace{\boxed{\lambda J(f)}}_{\text{Regularity}}$$

Choice of λ: minimize $\|f^\star - f_0\|$ (oracle)

Trade-off between denoising (λ increases with σ) and regularity of f_0.

No noise: $\sigma = 0$, $\lambda \to 0$, minimize

$$f^\star = \underset{f \in \mathbb{R}^N, \Phi f = y}{\operatorname{argmin}} J(f)$$

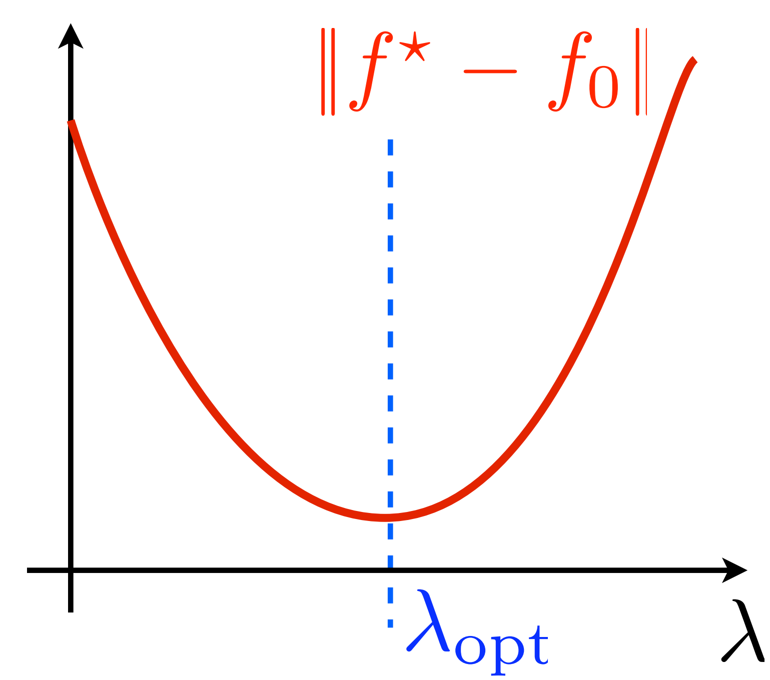

Smooth and Cartoon Priors

Prior model: energy $J(f) \in \mathbb{R}$ low for images of the model $f \in \Theta$.

Sobolev pseudo-norm: $$J(f) = \frac{1}{2}\|f\|^2_{\text{Sob}} = \frac{1}{2}\int \|\nabla_x f\|^2 \mathrm{d}x$$

(norm: $\mu\|f\| + \|f\|_{\text{Sob}}$)

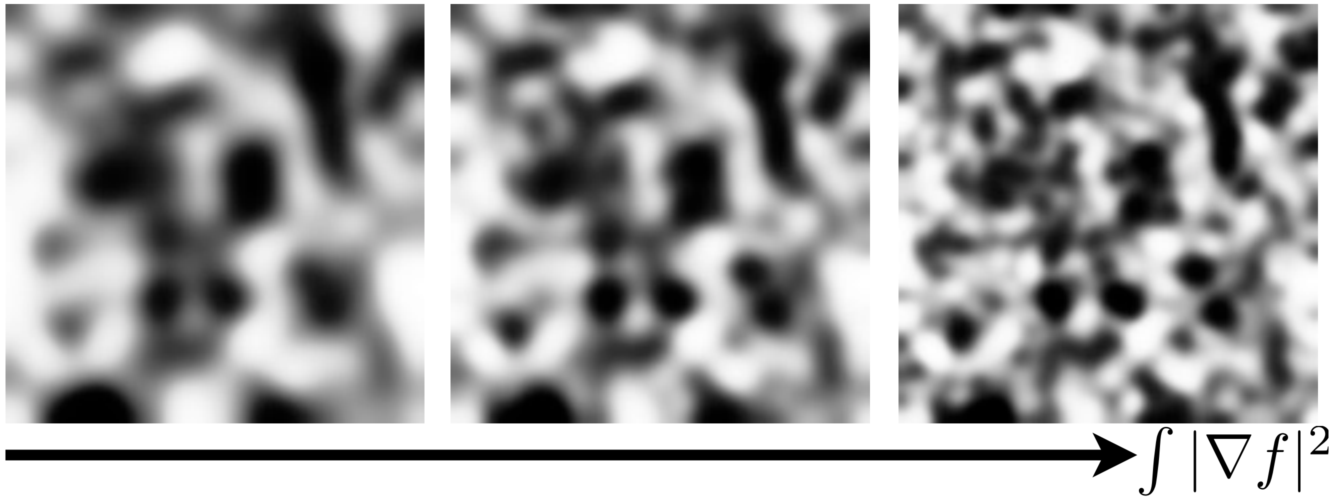

Smooth and Cartoon Priors

Prior model: energy $J(f) \in \mathbb{R}$ low for images of the model $f \in \Theta$.

Sobolev pseudo-norm: $$J(f) = \frac{1}{2}\|f\|^2_{\text{Sob}} = \frac{1}{2}\int \|\nabla_x f\|^2 \mathrm{d}x$$

(norm: $\mu\|f\| + \|f\|_{\text{Sob}}$)

Total variation pseudo-norm: $$J(f) = \|f\|_{\text{TV}} = \int \|\nabla_x f\| \mathrm{d}x$$

$\longrightarrow$ Extension to non-smooth functions $f \in \text{BV}([0,1]^2)$

Co-area formula: $$\|f\|_{\text{TV}} = \int_{\mathbb{R}} \text{length}(\mathcal{C}_t)\mathrm{d}t$$

Level set $\mathcal{C}_t = \{x \setminus f(x) = t\}$

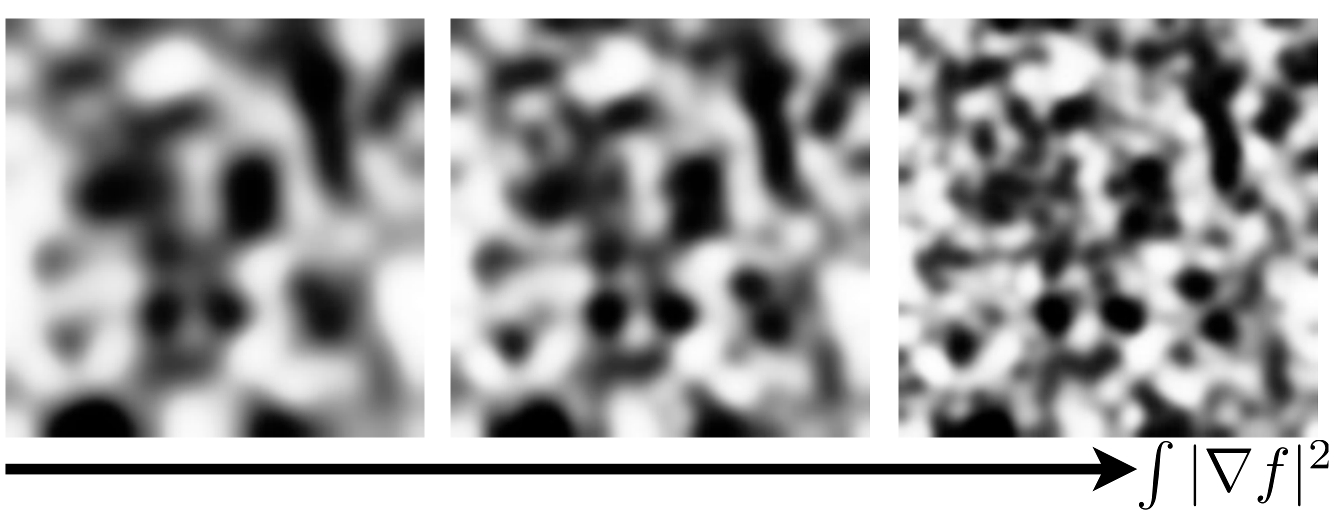

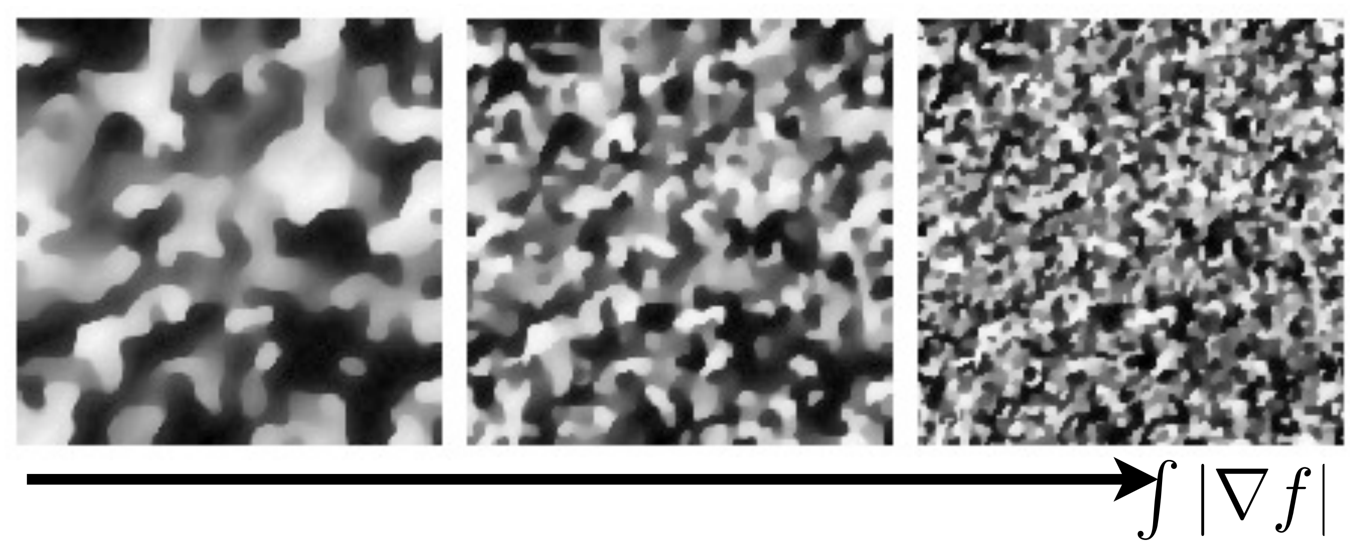

Sparse Priors

Orthogonal basis $\mathcal{B} = \{\psi_m\}_m$ of $\mathbb{R}^N$.

Example: Wavelet basis $\psi_m = \psi_{j,n}$, $m = (j, n)$.

Sparsity: most $\langle f,\, \psi_m \rangle$ are small.

Ideal sparsity: for most m, $\langle f,\, \psi_m \rangle = 0$.

Ideal prior: $J_0(f) = \#\,\{m \setminus \langle f,\, \psi_m \rangle \neq 0\}$

Image f

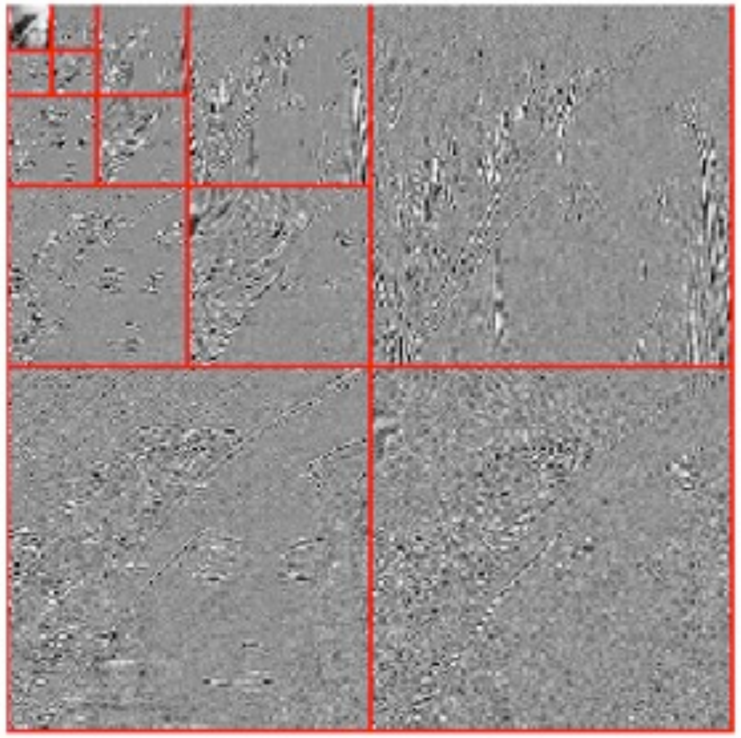

Coefficients $\{\langle f,\, \psi_m \rangle\}_m$

Sparse Priors

Orthogonal basis $\mathcal{B} = \{\psi_m\}_m$ of $\mathbb{R}^N$.

Example: Wavelet basis $\psi_m = \psi_{j,n}$, $m = (j, n)$.

Sparsity: most $\langle f,\, \psi_m \rangle$ are small.

Ideal sparsity: for most m, $\langle f,\, \psi_m \rangle = 0$.

Ideal prior: $J_0(f) = \#\{m \setminus \langle f,\, \psi_m \rangle \neq 0\}$

Image f

Best M-sparse approximation:

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

$$M = J_0(f_M) = \#\{m \setminus |\langle f,\, \psi_m \rangle| > T\}$$

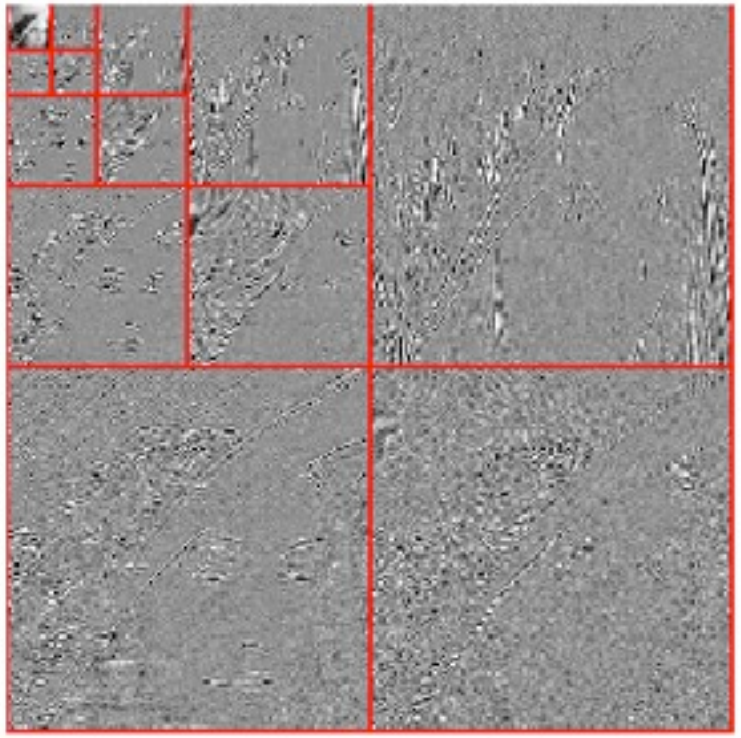

Coefficients $\{\langle f,\, \psi_m \rangle\}_m$

Approximate sparsity: $\|f - f_M\|$ is small.

Convex Relaxation: L1 Prior

"Ideal" sparsity prior: $J_0(f) = \# \{ m \setminus \langle f, \psi_m \rangle \neq 0 \}$

Image with 2 pixels:

$J_0(f) = 0 \longrightarrow$ null image.

$J_0(f) = 1 \longrightarrow$ sparse image.

$J_0(f) = 2 \longrightarrow$ non-sparse image.

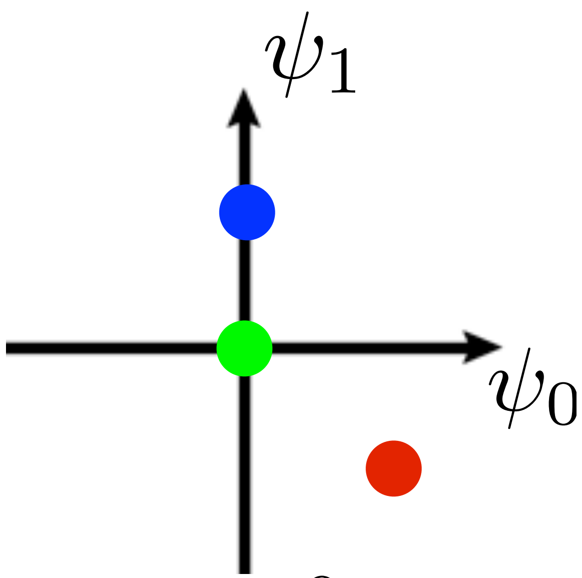

$q = 0$

Convex Relaxation: L1 Prior

"Ideal" sparsity prior: $J_0(f) = \#\{m \setminus \langle f,\, \psi_m \rangle \neq 0\}$

Image with 2 pixels:

$J_0(f) = 0 \longrightarrow$ null image.

$J_0(f) = 1 \longrightarrow$ sparse image.

$J_0(f) = 2 \longrightarrow$ non-sparse image.

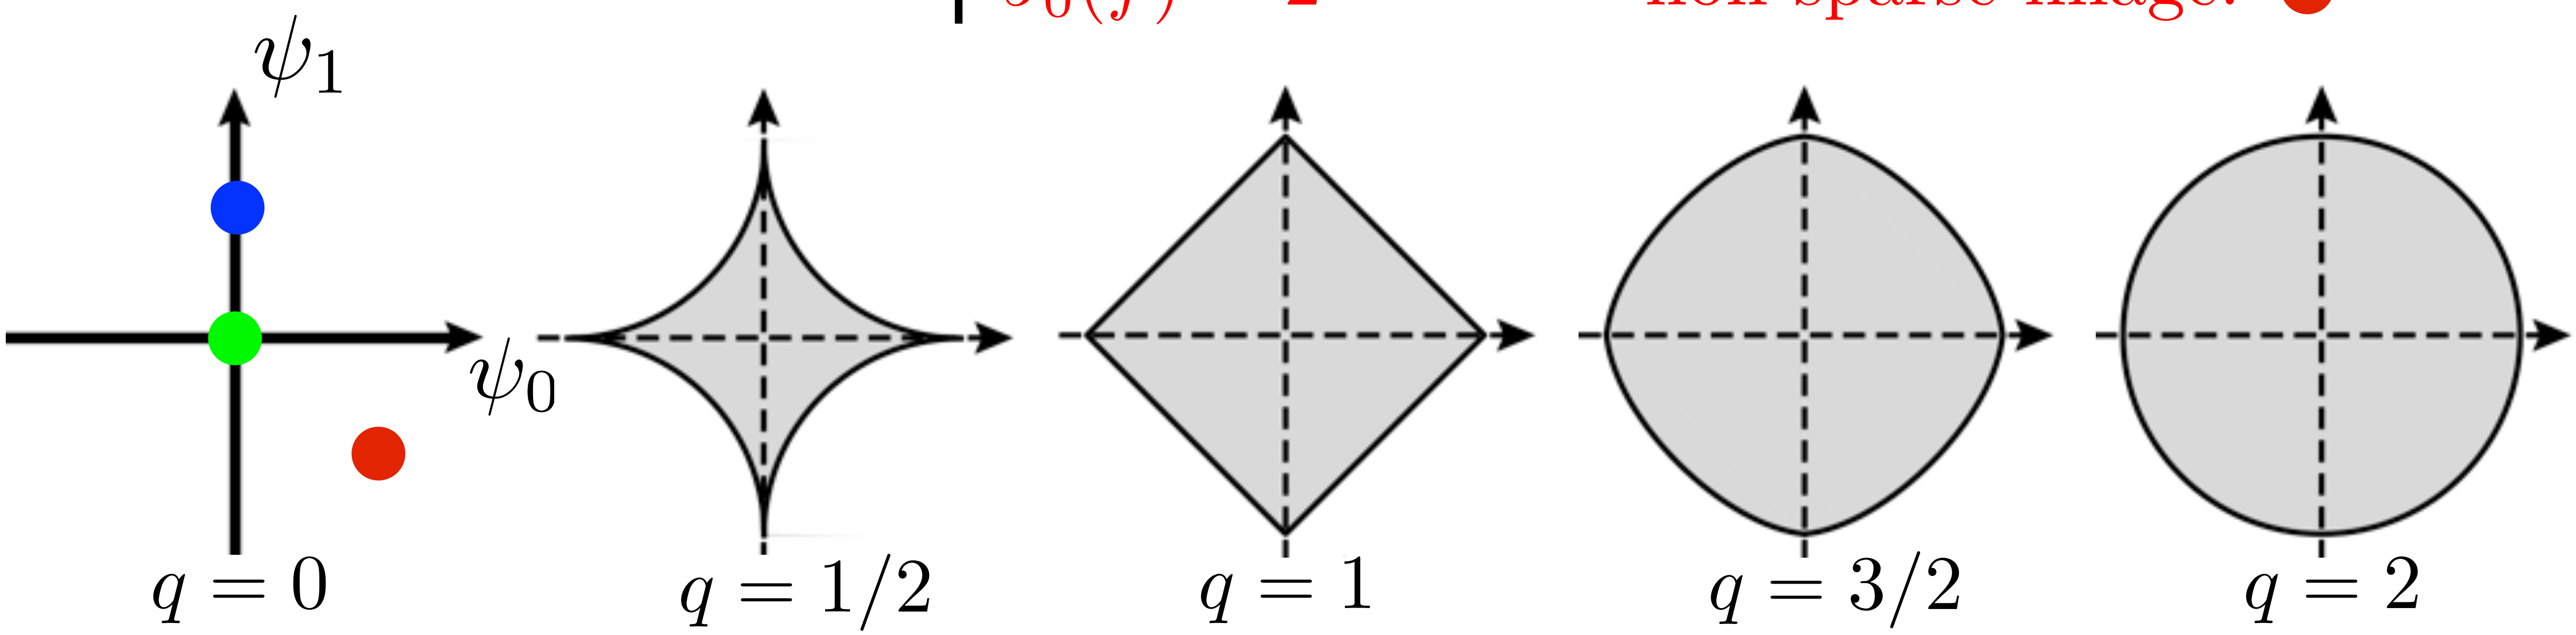

ℓ^q priors: $J_q(f) = \sum_m |\langle f,\, \psi_m \rangle|^q$ (convex for $q \geqslant 1$)

Convex Relaxation: L1 Prior

"Ideal" sparsity prior: $J_0(f) = \#\{m \backslash \langle f,\, \psi_m \rangle \neq 0\}$

Image with 2 pixels:

$J_0(f) = 0 \longrightarrow$ null image.

$J_0(f) = 1 \longrightarrow$ sparse image.

$J_0(f) = 2 \longrightarrow$ non-sparse image.

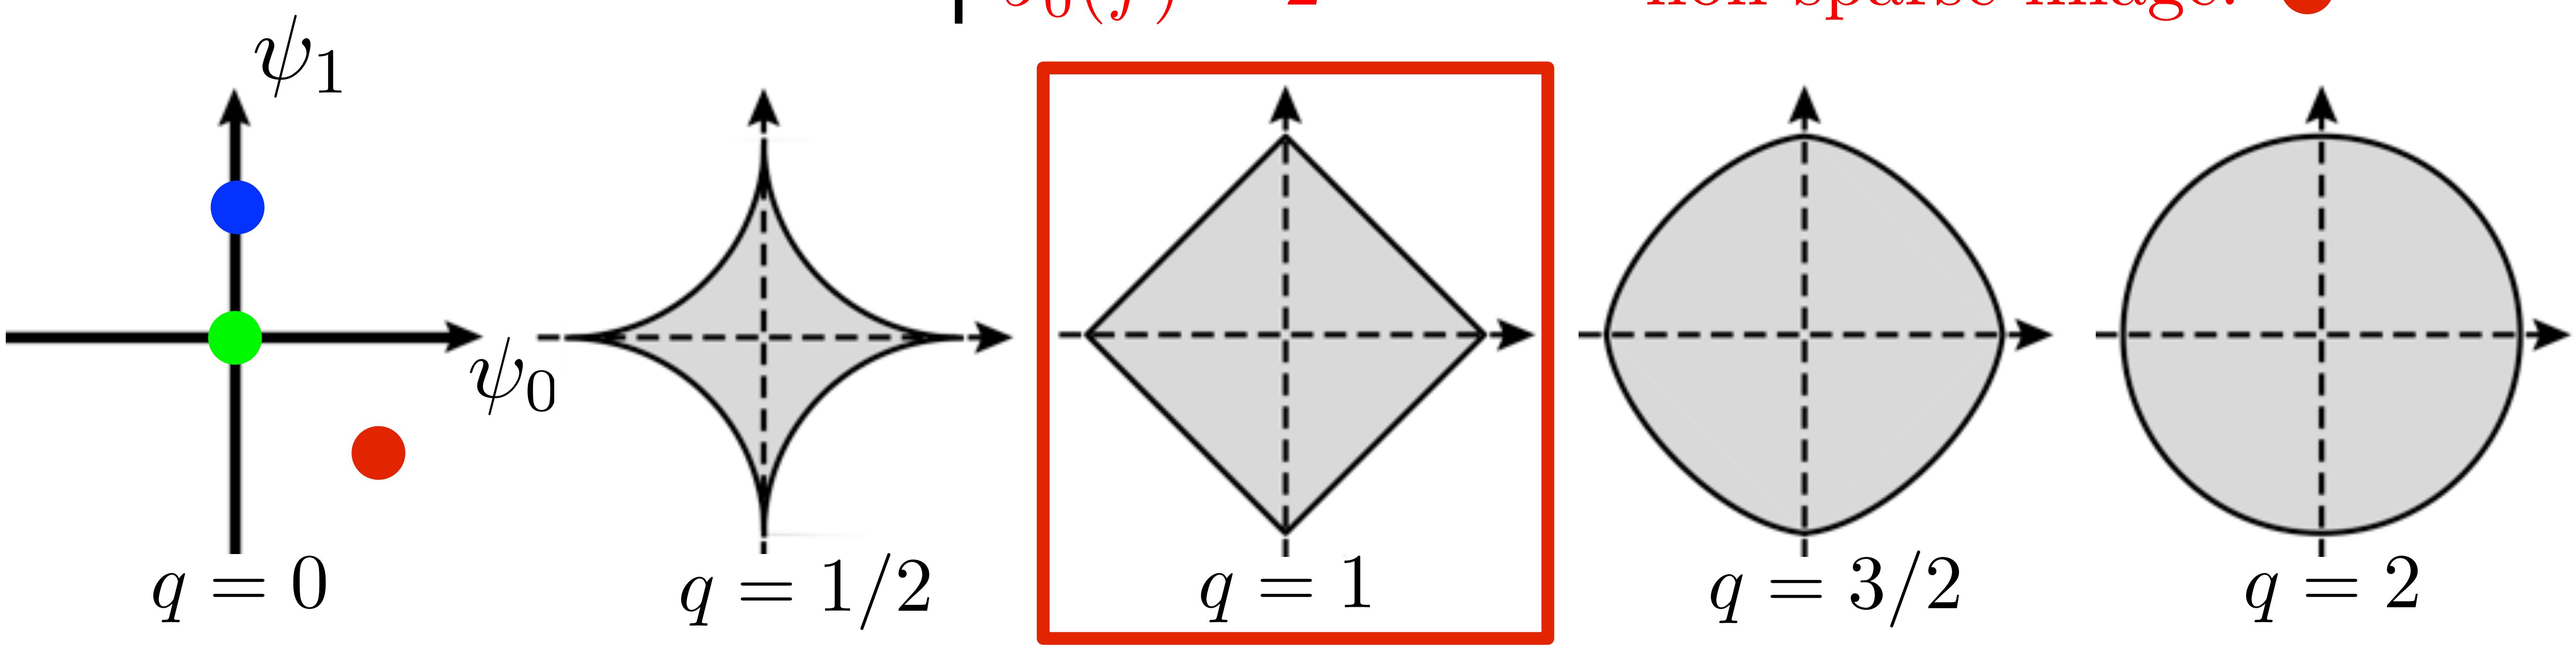

ℓ^q priors: $J_q(f) = \sum_m |\langle f,\, \psi_m \rangle|^q$ (convex for $q \geqslant 1$)

ℓ^1 norm: ℓ^q norm the "closest" to the ℓ^0 ideal sparsity.

Sparse ℓ^1 prior: $\boxed{J_1(f) = \sum_m |\langle f,\, \psi_m \rangle|}$

Noiseless Sparse Regularization

Orthogonal basis $\{\psi_m\}_m$ of $\mathbb{R}^N$. Sparse prior: $J(f) = \sum_m |\langle f,\, \psi_m \rangle|$

Noiseless measurements $y = \Phi f_0$: $$f^\star = \underset{\Phi f = y}{\text{argmin}} \sum_m |\langle f,\, \psi_m \rangle|$$

Convex linear program.

→ Interior points, cf. [Chen, Donoho, Saunders] "basis pursuit".

→ Douglas-Rachford splitting, see [Combettes, Pesquet].

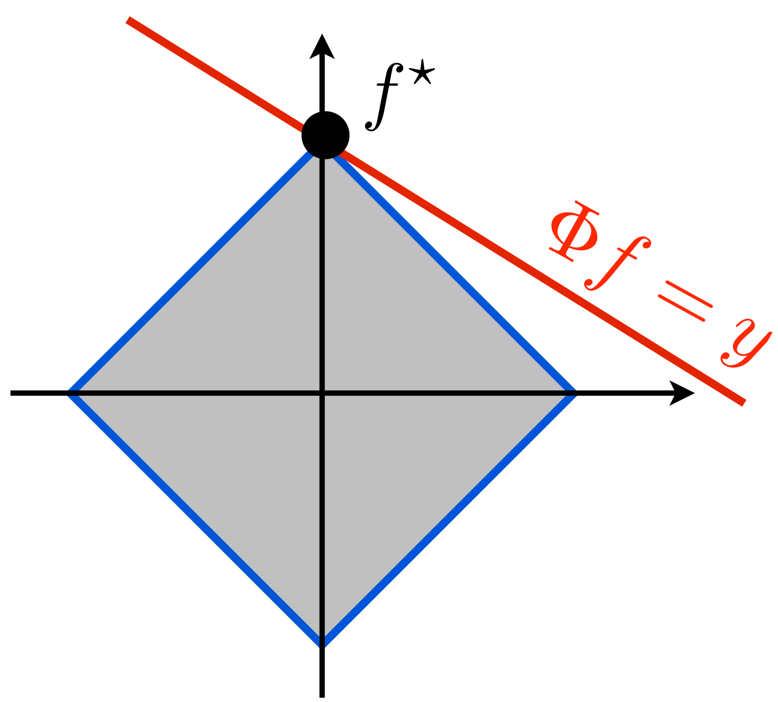

$$f^\star = \underset{\Phi f = y}{\text{argmin}} \sum_m |\langle f,\, \psi_m \rangle|$$

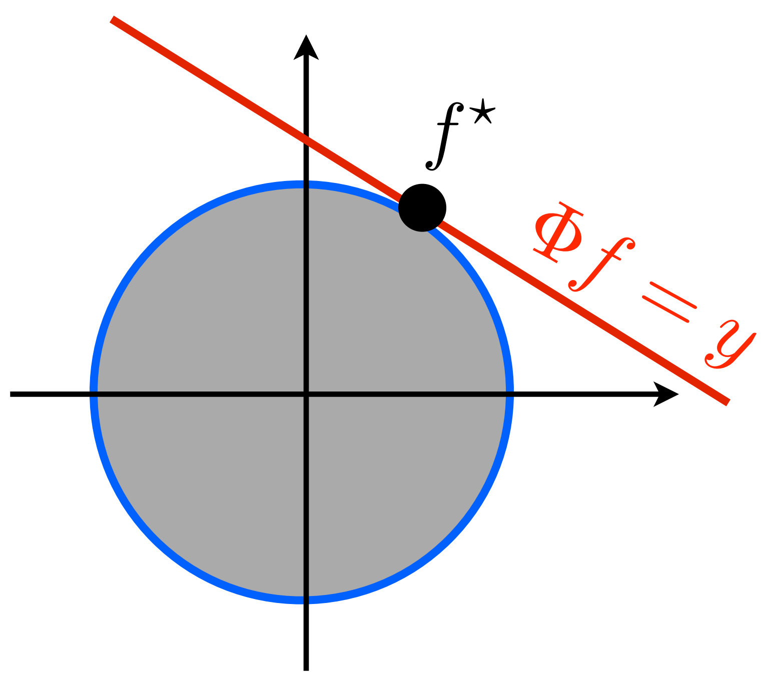

$$f^\star = \underset{\Phi f = y}{\text{argmin}} \sum_m |\langle f,\, \psi_m \rangle|^2 = \|f\|^2$$

Noisy Sparse Regularization

Noisy measurements: $y = \Phi f_0 + w$: $$f^\star = \underset{\|\Phi f - y\| \leqslant \varepsilon}{\operatorname{argmin}} \sum_m |\langle f,\, \psi_m \rangle|.$$

Convex program, can be solved with Lagrangian relaxation.

$$f^\star = \underset{f \in \mathbb{R}^N}{\operatorname{argmin}} \underbrace{\boxed{\frac{1}{2}\|\Phi f - y\|^2}}_{\text{Data fitting}} + \underbrace{\boxed{\lambda \sum_m |\langle f,\, \psi_m \rangle|.}}_{\text{Sparsity}}$$

→ Iterative thresholding, see [Daubechies et al], [Pesquet et al], etc.

→ Nesterov multi-steps schemes.

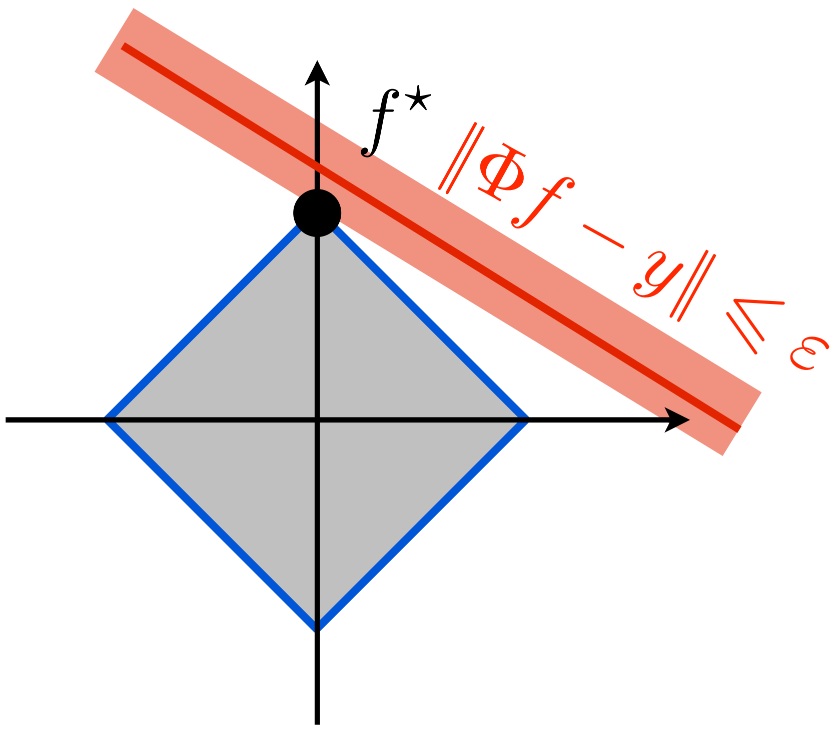

Inpainting Results

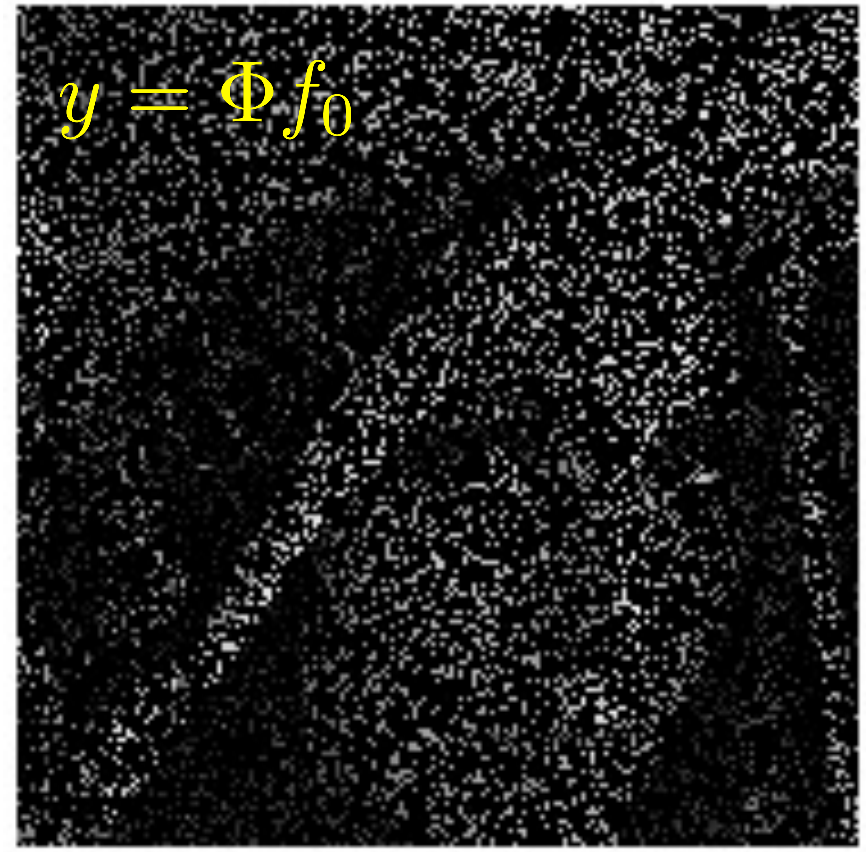

Sobolev, 20.8dB

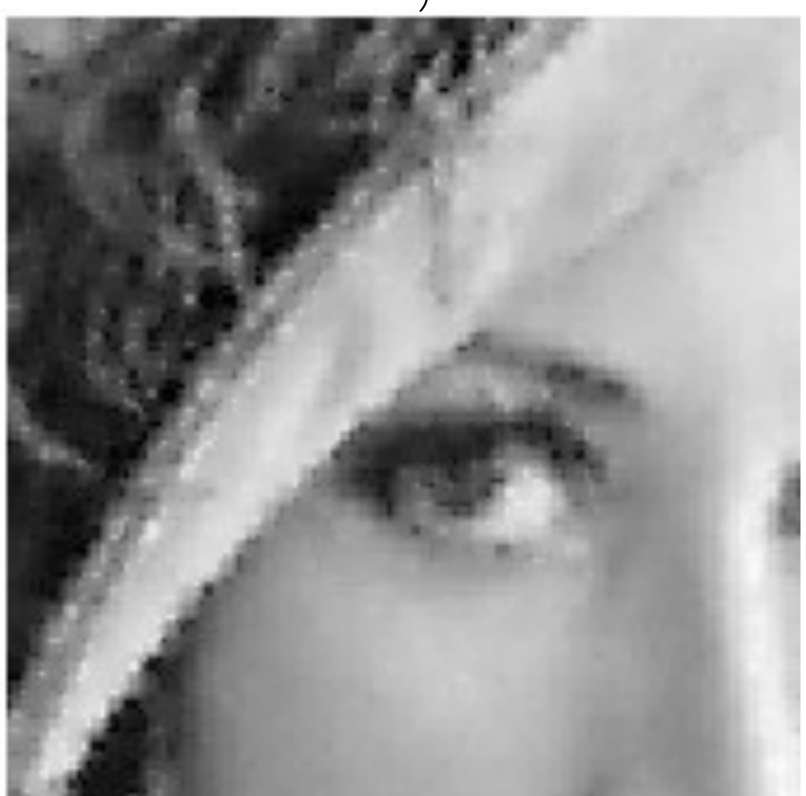

Wavelets orth, 16.6dB

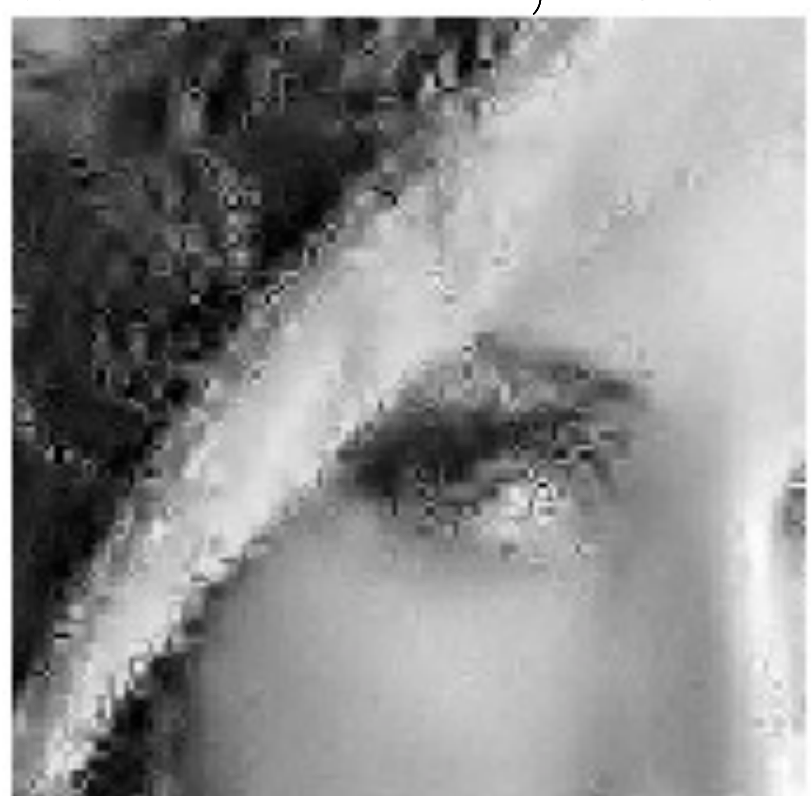

Wavelets TI, 23.6dB

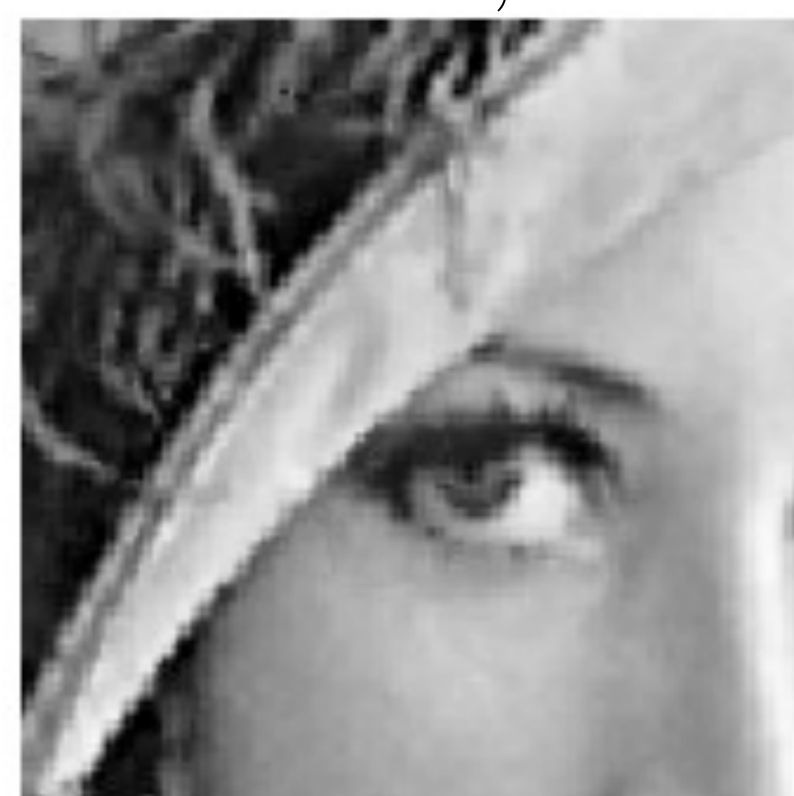

Best-Basis Inverse Problems

Dictionary of ortho-bases $\mathcal{B}^\lambda = \{\psi_m^\lambda\}_m$.

Measurements $y = \Phi f_0 + w$, regularization:

$$\min_{f,\lambda\in\Lambda} E(\lambda, f) = \frac{1}{2}\|y - \Phi f\|^2 + \gamma \sum_m |\langle f,\, \psi_m^\lambda\rangle|$$

$\longrightarrow$ estimate both the image f and its geometry λ.

Best-Basis Inverse Problems

Dictionary of ortho-bases $\mathcal{B}^\lambda = \{\psi_m^\lambda\}_m$.

Measurements $y = \Phi f_0 + w$, regularization:

$$\min_{f,\lambda \in \Lambda} E(\lambda, f) = \frac{1}{2}\|y - \Phi f\|^2 + \gamma \sum_m |\langle f, \psi_m^\lambda \rangle|$$

$\longrightarrow$ estimate both the image f and its geometry λ.

$\longrightarrow$ non-convex problem, use a descent method on (λ, f).

Initialization: $f^{(0)} = 0$.

Gradient step: $\tilde{f}^{(k)} = f^{(k)} - \tau \Phi^* (\Phi f^{(k)} - y)$

Best-Basis Inverse Problems

Dictionary of ortho-bases $\mathcal{B}^\lambda = \{\psi_m^\lambda\}_m$.

Measurements $y = \Phi f_0 + w$, regularization:

$$\min_{f, \lambda \in \Lambda} E(\lambda, f) = \frac{1}{2}\|y - \Phi f\|^2 + \gamma \sum_m |\langle f,\, \psi_m^\lambda \rangle|$$

$\longrightarrow$ estimate both the image f and its geometry λ.

$\longrightarrow$ non-convex problem, use a descent method on (λ, f).

Initialization: $f^{(0)} = 0$.

Gradient step: $\tilde{f}^{(k)} = f^{(k)} - \tau \Phi^* (\Phi f^{(k)} - y)$

Best basis: $\lambda^{(k+1)} = \underset{\lambda}{\text{argmin}} \sum_m \Psi(\langle \tilde{f}^{(k)},\, \psi_m^\lambda \rangle)$ $\quad \Psi(x) = \begin{cases} x^2 & \text{if} \quad |x| \leqslant \gamma, \\ 2\gamma|x| - \gamma^2. \end{cases}$

$\longrightarrow$ Lagrangian minimization, fast best-basis search.

Best-Basis Inverse Problems

Dictionary of ortho-bases $\mathcal{B}^\lambda = \{\psi_m^\lambda\}_m$.

Measurements $y = \Phi f_0 + w$, regularization:

$$\min_{f,\lambda\in\Lambda} E(\lambda, f) = \frac{1}{2}\|y - \Phi f\|^2 + \gamma \sum_m |\langle f,\, \psi_m^\lambda \rangle|$$

$\longrightarrow$ estimate both the image f and its geometry λ.

$\longrightarrow$ non-convex problem, use a descent method on (λ, f).

Initialization: $f^{(0)} = 0$.

Gradient step: $\tilde{f}^{(k)} = f^{(k)} - \tau \Phi^* (\Phi f^{(k)} - y)$

Best basis: $\lambda^{(k+1)} = \underset{\lambda}{\text{argmin}} \sum_m \Psi(\langle \tilde{f}^{(k)},\, \psi_m^\lambda \rangle)$ $\quad \Psi(x) = \begin{cases} x^2 & \text{if} \quad |x| \leqslant \gamma, \\ 2\gamma|x| - \gamma^2. \end{cases}$

$\longrightarrow$ Lagrangian minimization, fast best-basis search.

Thresholding: $f^{(k+1)} = \sum_m s_\gamma(\langle \tilde{f}^{(k)},\, \psi_m^{\lambda^{(k+1)}} \rangle) \psi_m^{\lambda^{(k+1)}}$ $\quad s_\gamma(x) = \max(0, 1 - \frac{\gamma}{|x|}) x$

Best-Basis Inverse Problems

Dictionary of ortho-bases $\mathcal{B}^\lambda = \{\psi_m^\lambda\}_m$.

Measurements $y = \Phi f_0 + w$, regularization:

$$\min_{f, \lambda \in \Lambda} E(\lambda, f) = \frac{1}{2}\|y - \Phi f\|^2 + \gamma \sum_m |\langle f, \psi_m^\lambda \rangle|$$

$\longrightarrow$ estimate both the image f and its geometry λ.

$\longrightarrow$ non-convex problem, use a descent method on (λ, f).

Initialization: $f^{(0)} = 0$.

Gradient step: $\tilde{f}^{(k)} = f^{(k)} - \tau \Phi^* (\Phi f^{(k)} - y)$

Best basis: $\lambda^{(k+1)} = \underset{\lambda}{\operatorname{argmin}} \sum_m \Psi(\langle \tilde{f}^{(k)}, \psi_m^\lambda \rangle)$ $\quad \Psi(x) = \begin{cases} x^2 & \text{if} \quad |x| \leqslant \gamma, \\ 2\gamma|x| - \gamma^2. \end{cases}$

$\longrightarrow$ Lagrangian minimization, fast best-basis search.

Thresholding: $f^{(k+1)} = \sum_m s_\gamma(\langle \tilde{f}^{(k)}, \psi_m^{\lambda^{(k+1)}} \rangle) \psi_m^{\lambda^{(k+1)}}$ $\quad s_\gamma(x) = \max(0, 1 - \frac{\gamma}{|x|}) x$

One has: $\quad E(\lambda^{(k+1)}, f^{(k+1)}) \leqslant E(\lambda^{(k)}, f^{(k)})$

Bandlets Compressed Sensing

Compressed sensing acquisition: $\Phi f = \{\langle f, \varphi_k \rangle\}_{k=0}^{P-1}$ $\quad P = N/6$

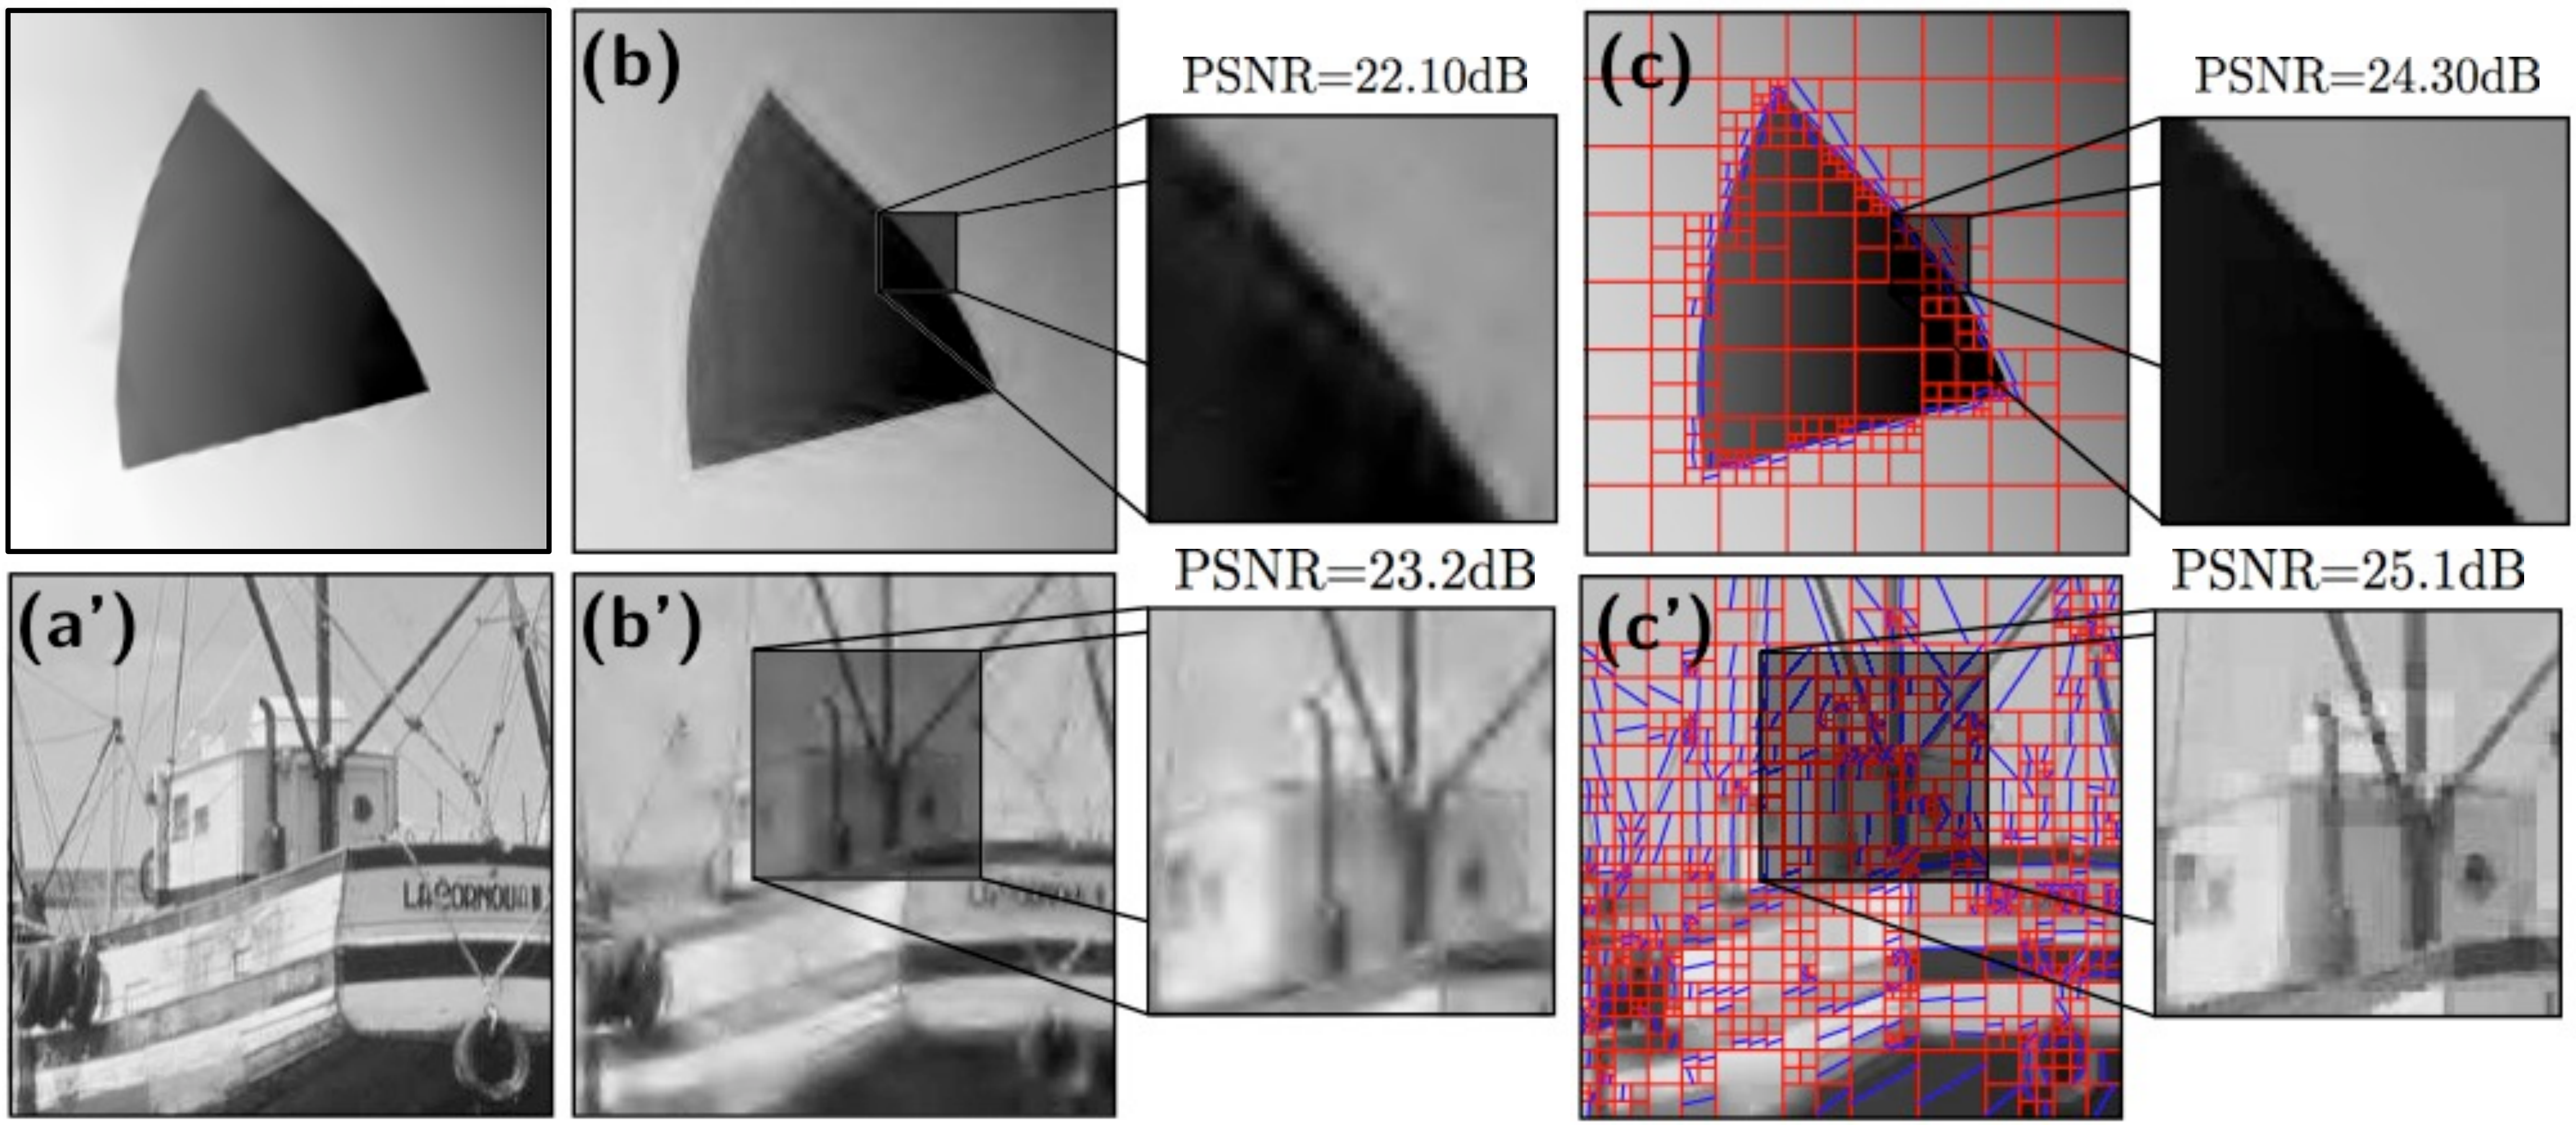

Original image Wavelets (TI) Bandlets

Grouplets Texture Inpainting

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

Grouplet inpainting: iteratively estimate the geometry λ in Ω. [Peyré, 2010]

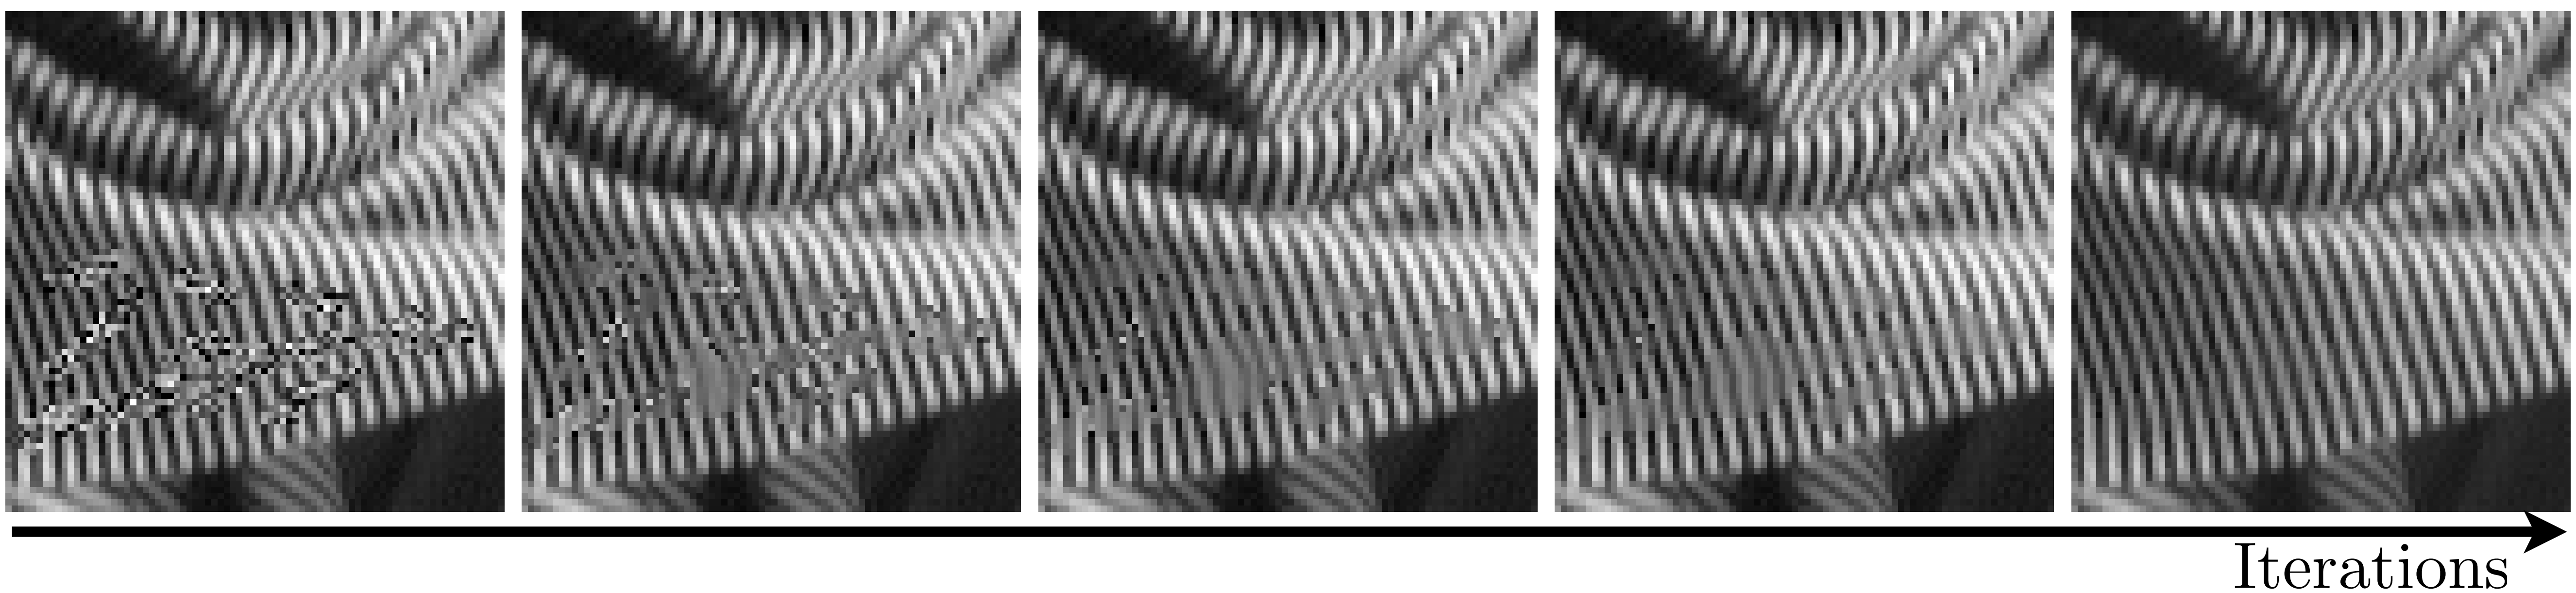

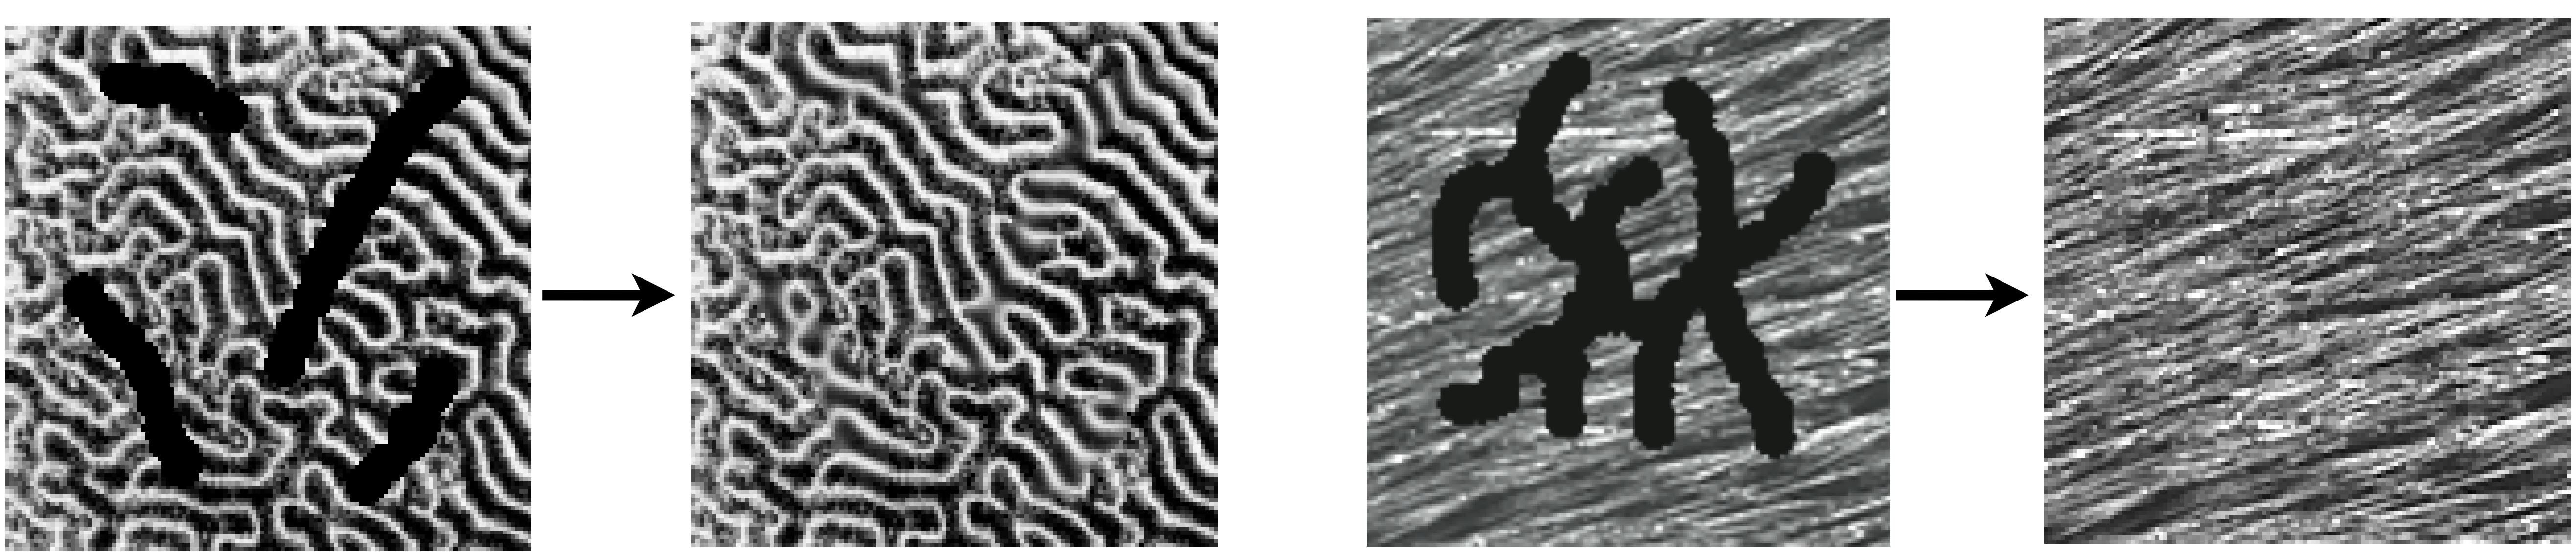

Overview

- Sparsity for Approximation
- Sparsity for Processing
- Geometric Images
- Adaptive Geometric Processing
- Adaptive Inverse Problems Regularization
- Geometric Texture Synthesis

Texture Synthesis

Texture synthesis: generate $f^\star$ perceptually similar to some input f

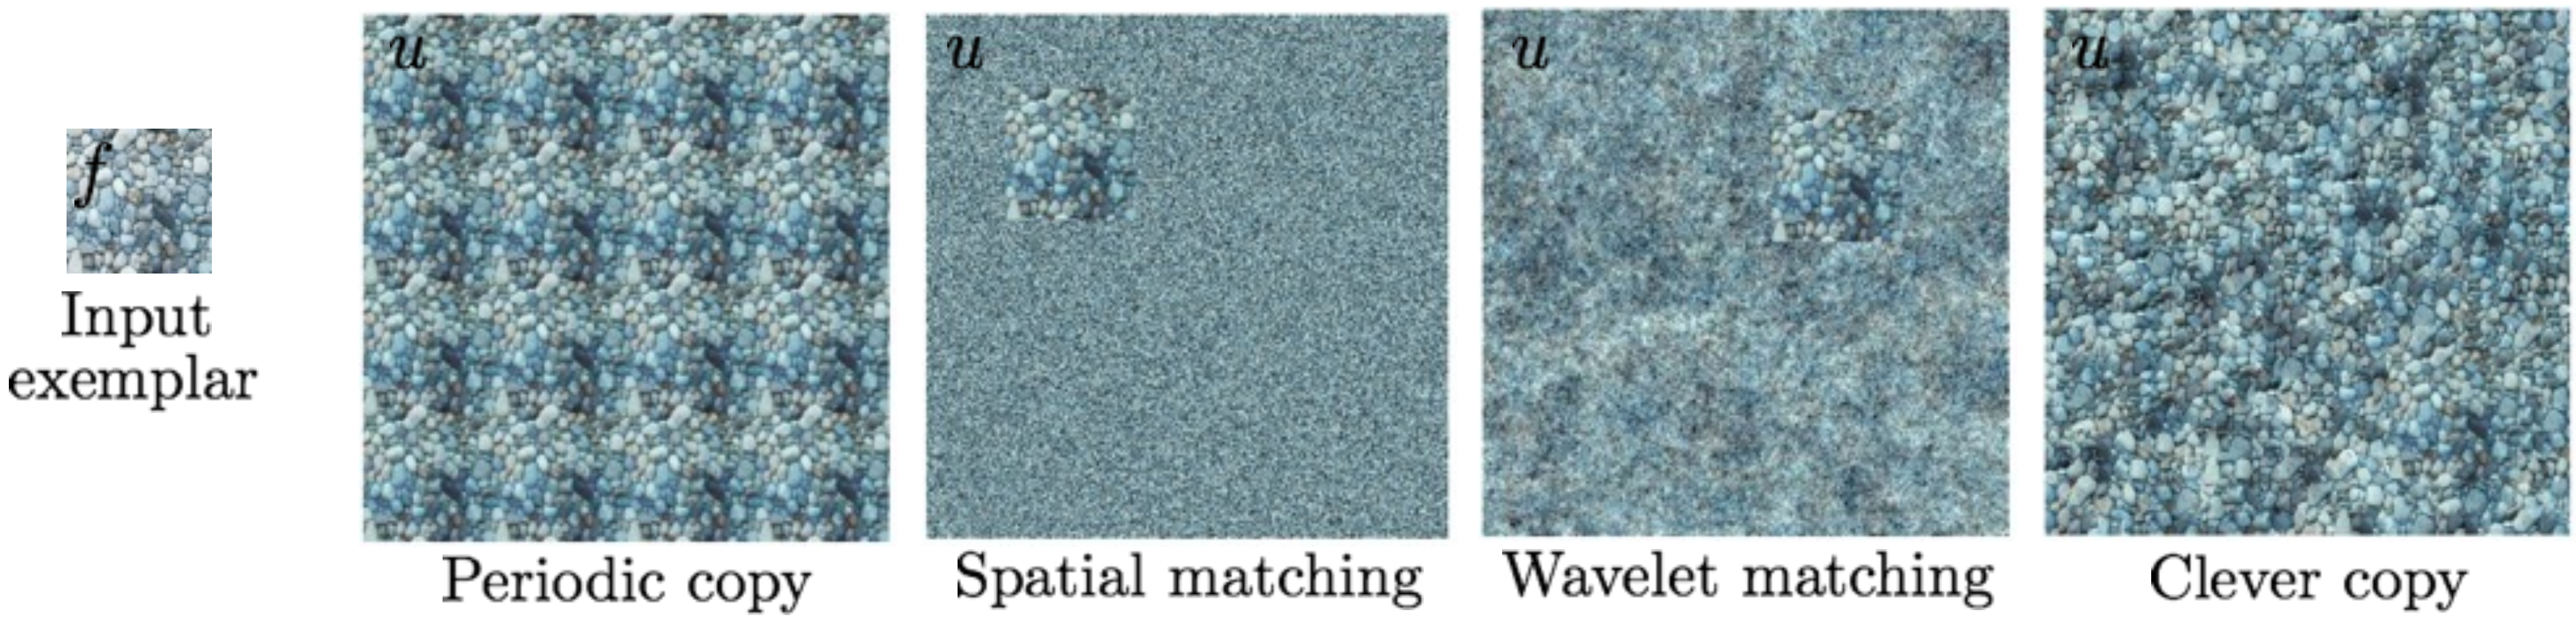

Texture Synthesis

Texture synthesis: generate $f^\star$ perceptually similar to some input f

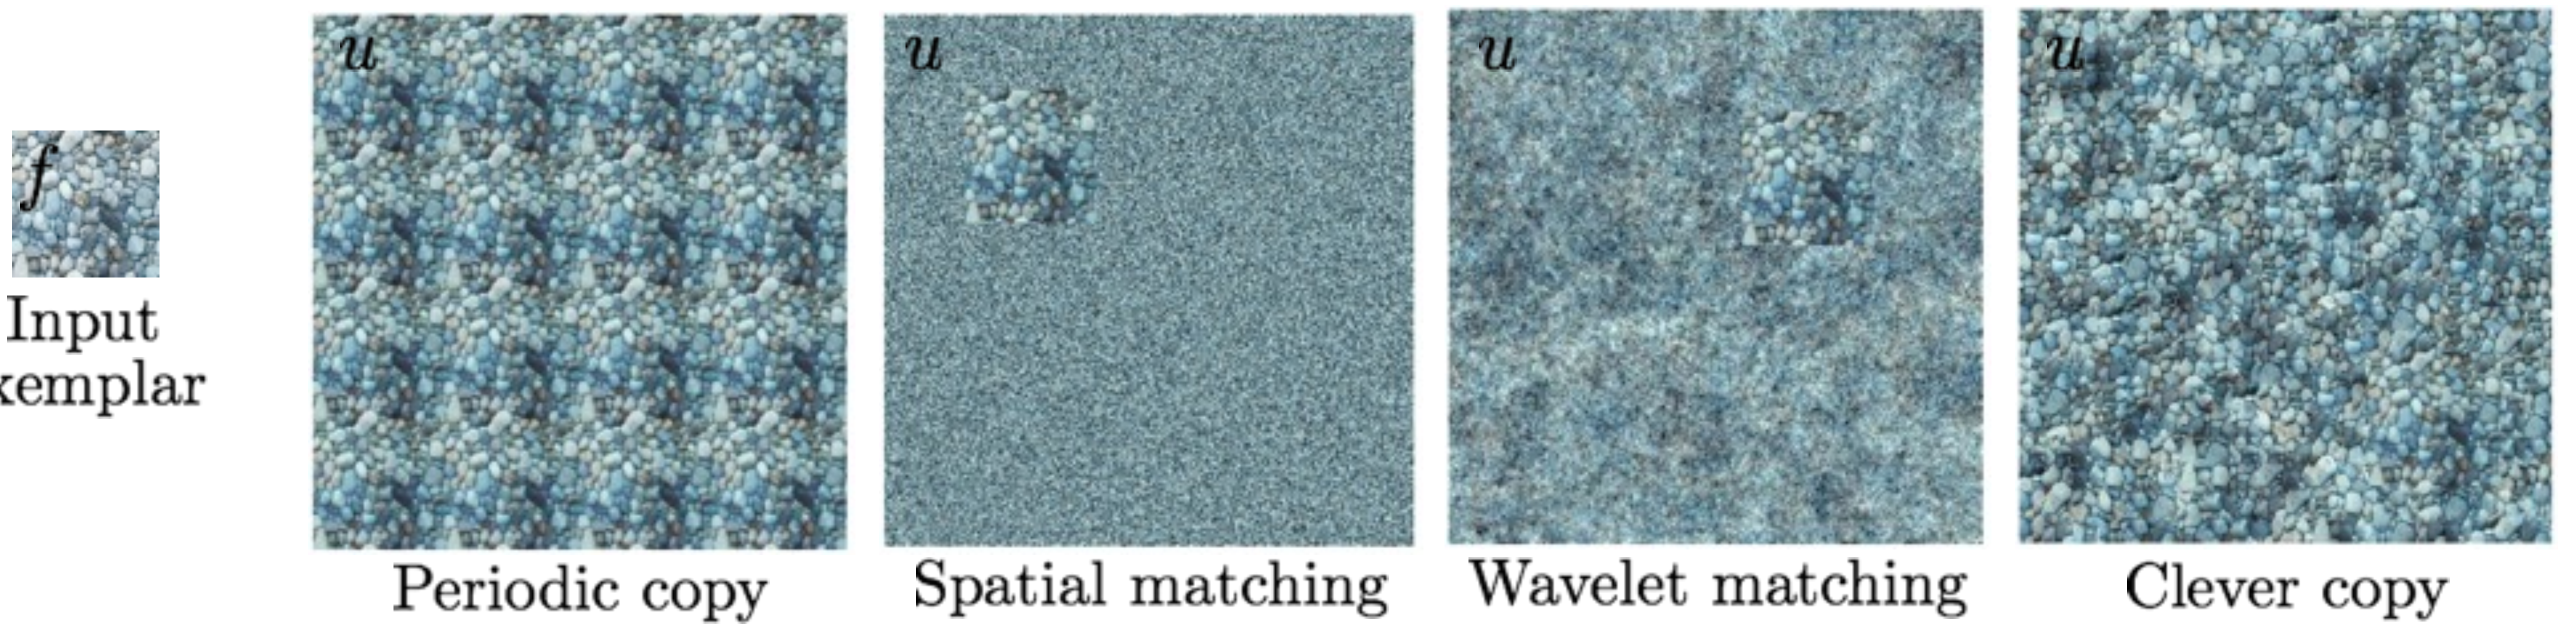

Mathematical model: $f \in \Theta_\Gamma$ where Γ is a geometry.

Analysis of f using a basis $\mathcal{B}(\Gamma) = \{\psi_m^\Gamma\}_m$ (bandlets, grouplets, etc).

Θ_Γ encodes statistical constraints on $\{\langle f, \psi_m^\Gamma \rangle\}_m$.

Synthesis: draw $f^\star \in \Theta_\Gamma$ uniformly at random.

Possible to modify the geometry $\Gamma \mapsto \Gamma^\star$ and draw $f^\star \in \Theta_{\Gamma^\star}$

Wavelet Texture Synthesis

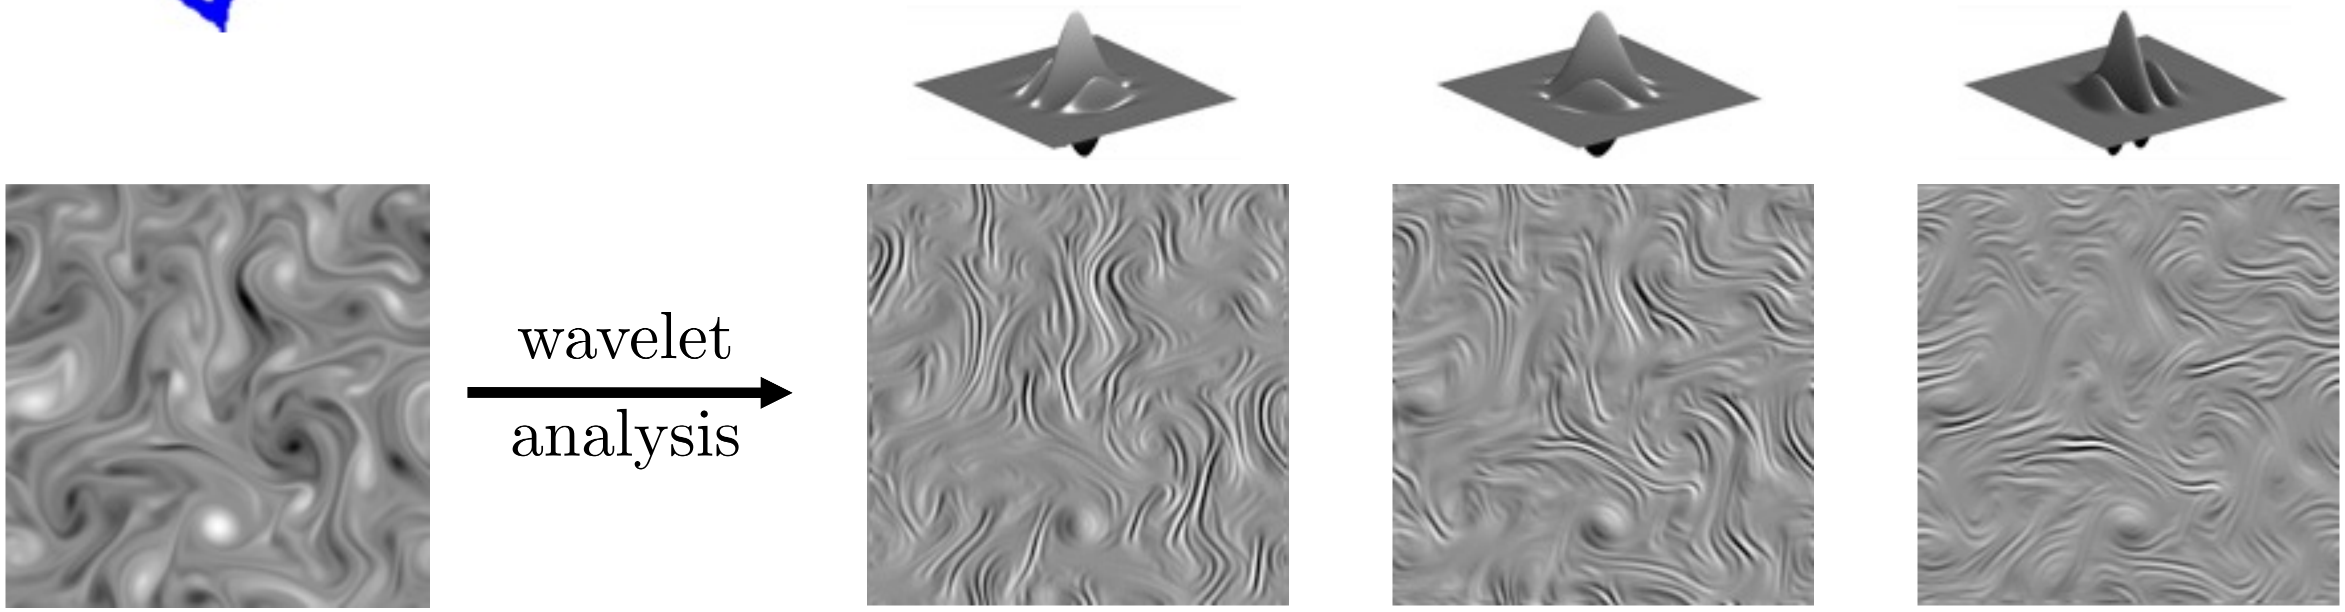

Wavelet Texture Synthesis

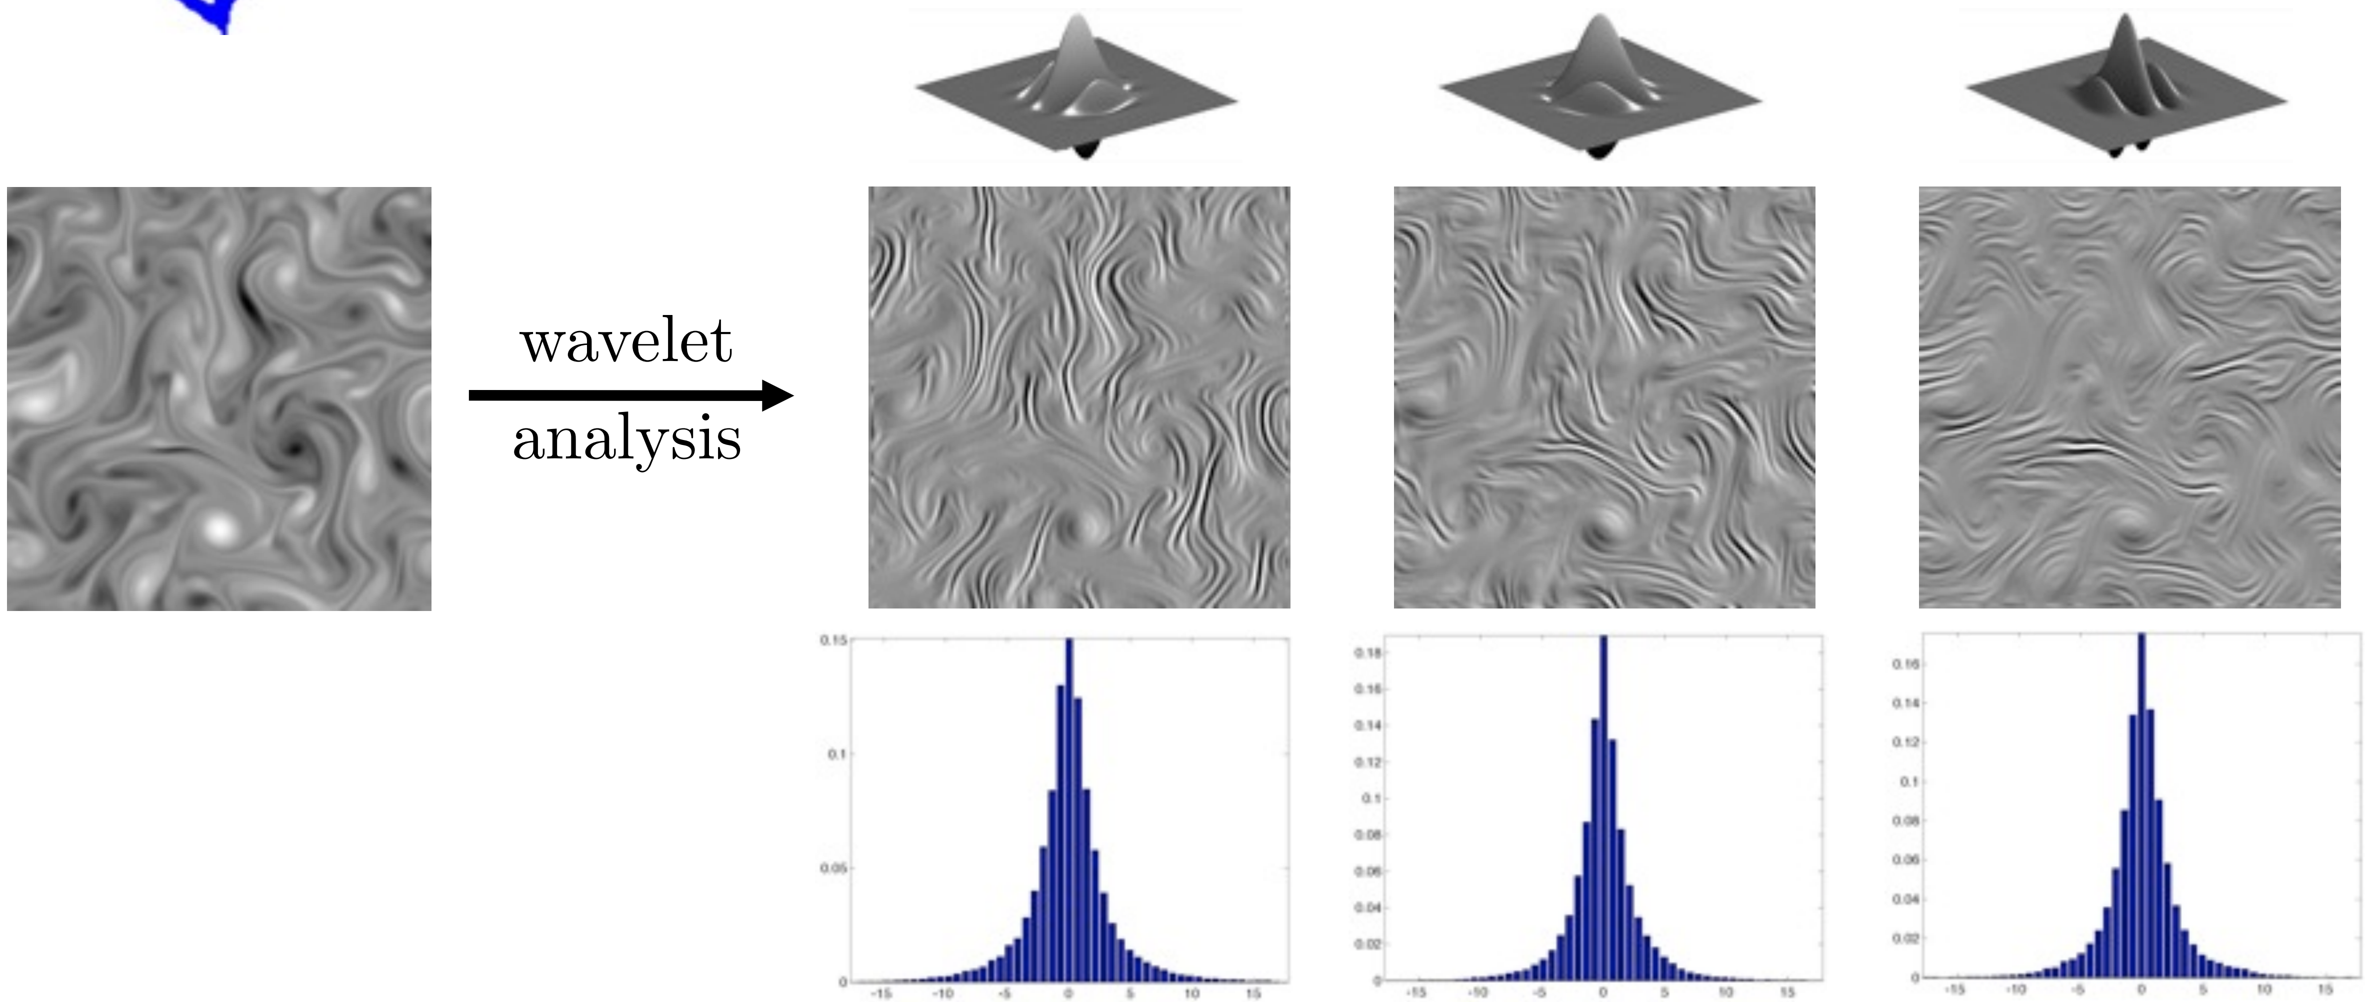

Wavelet Texture Synthesis

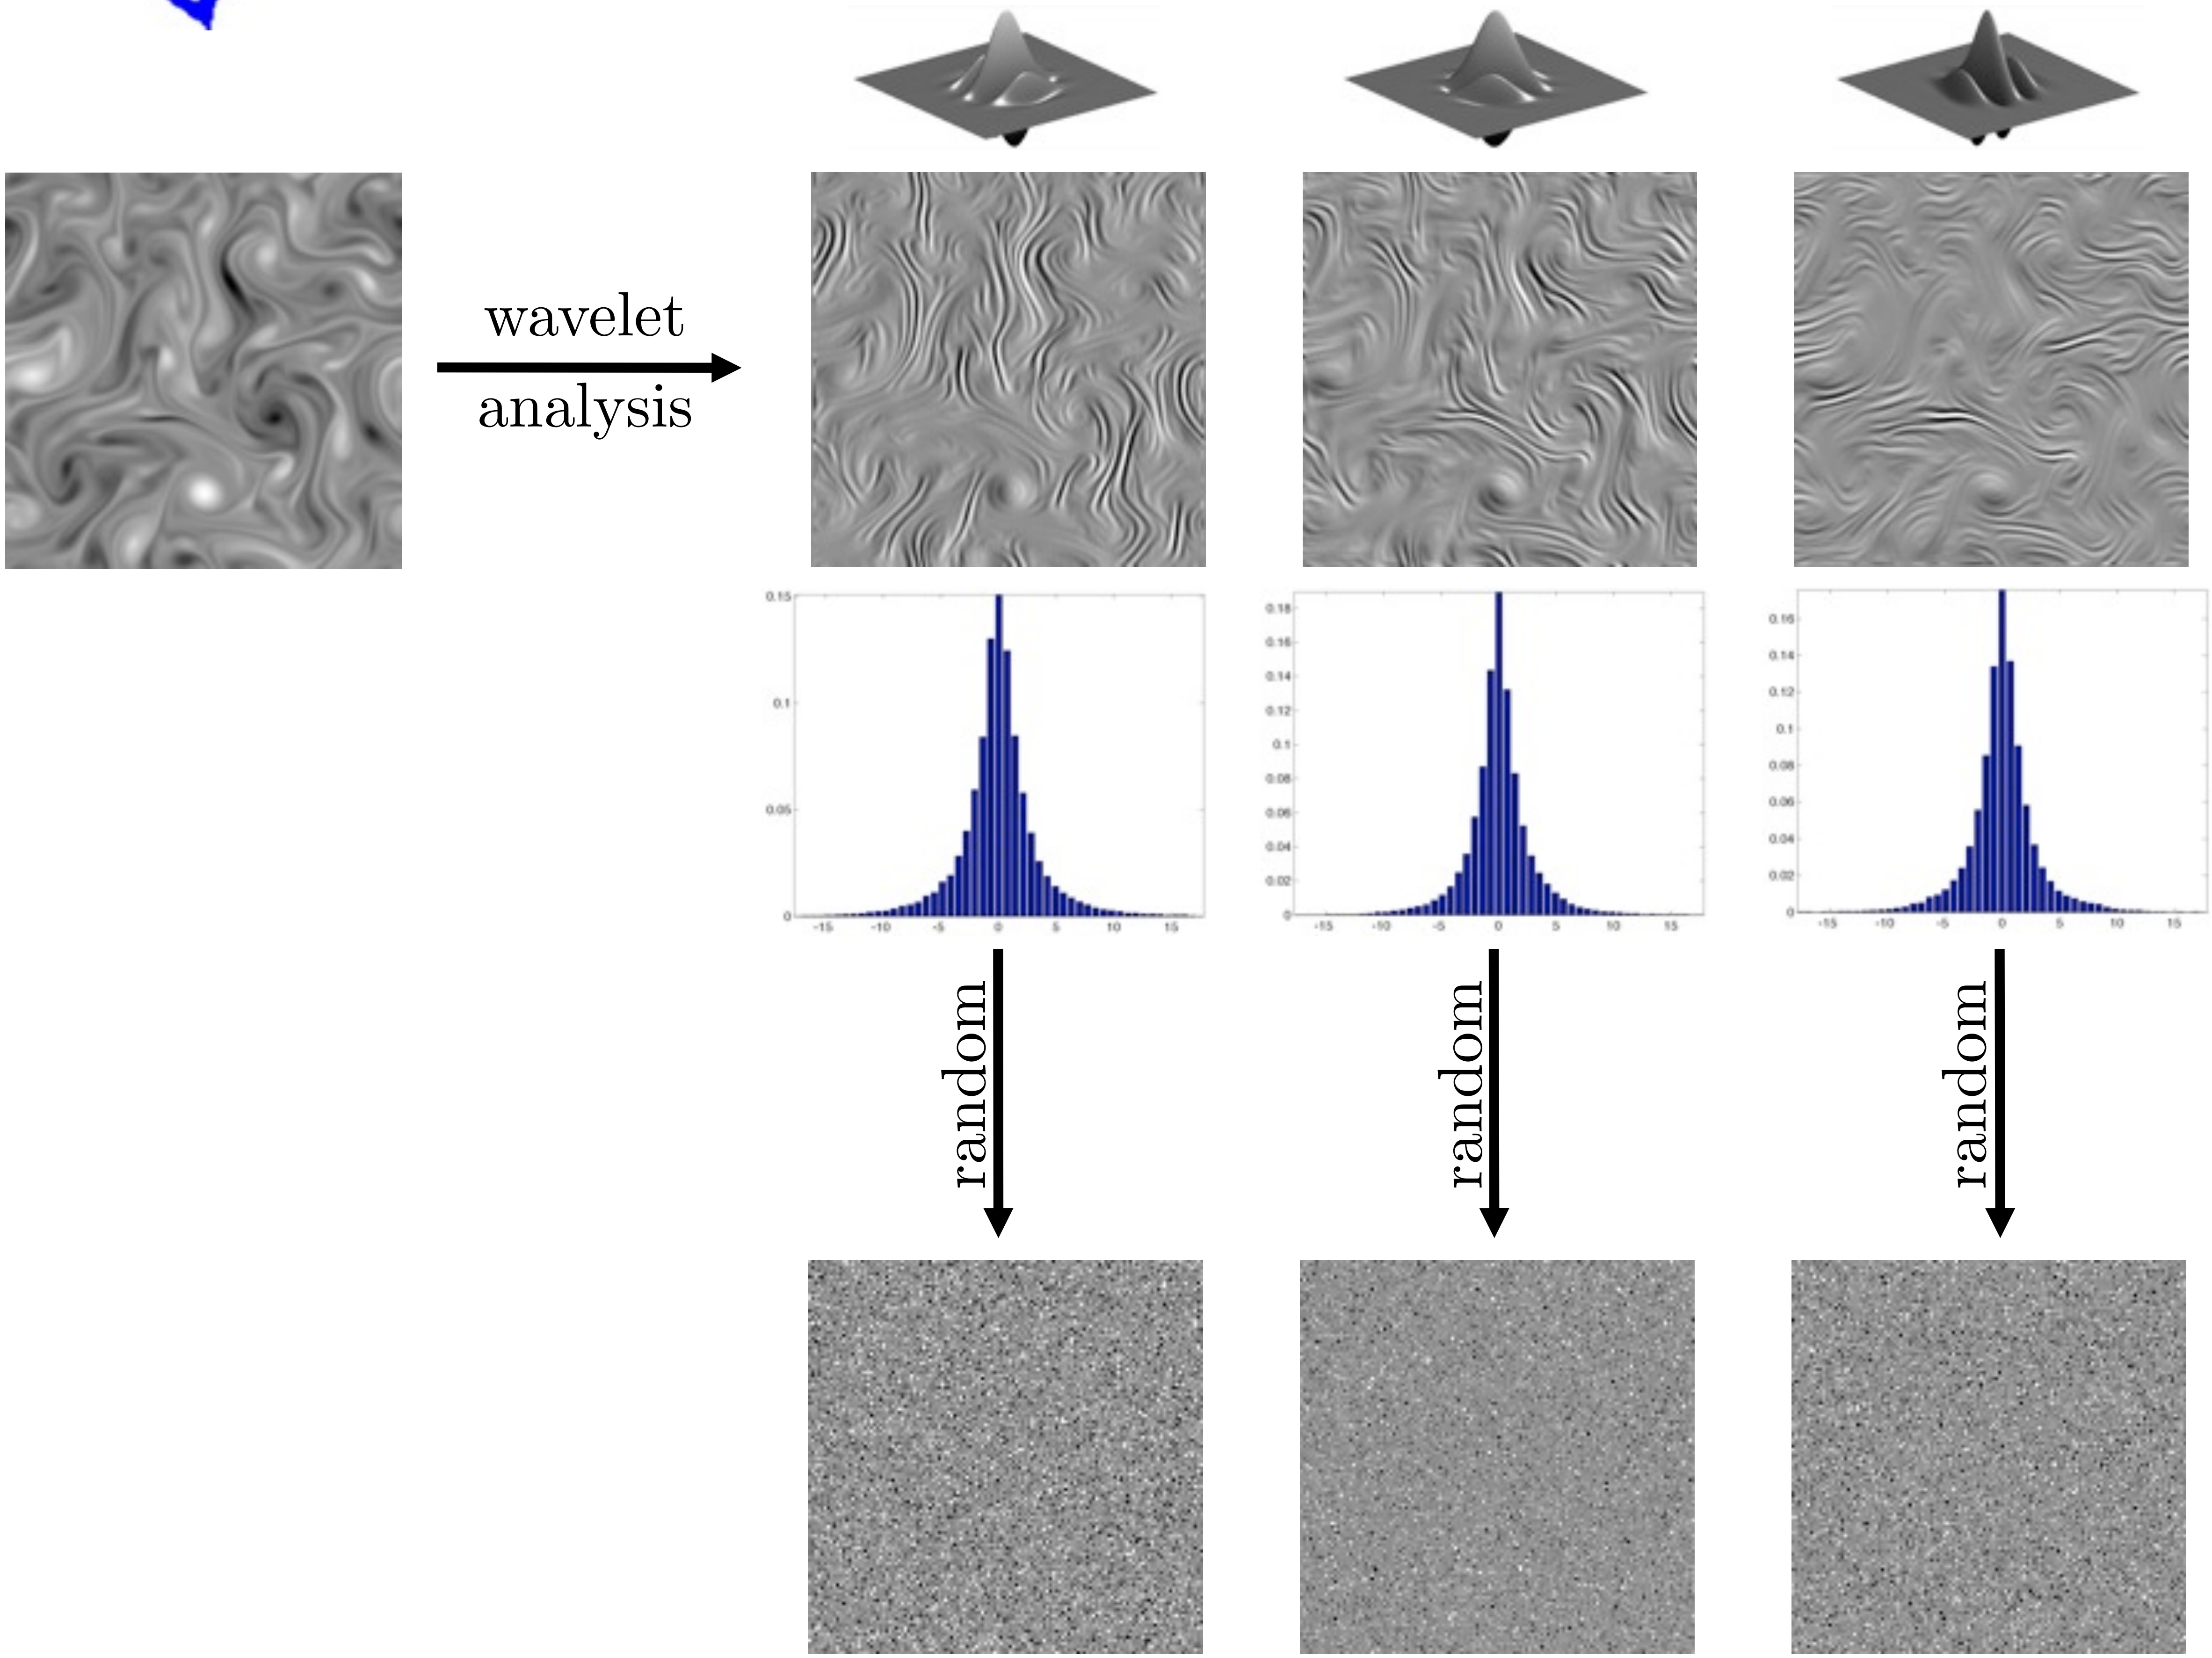

Wavelet Texture Synthesis

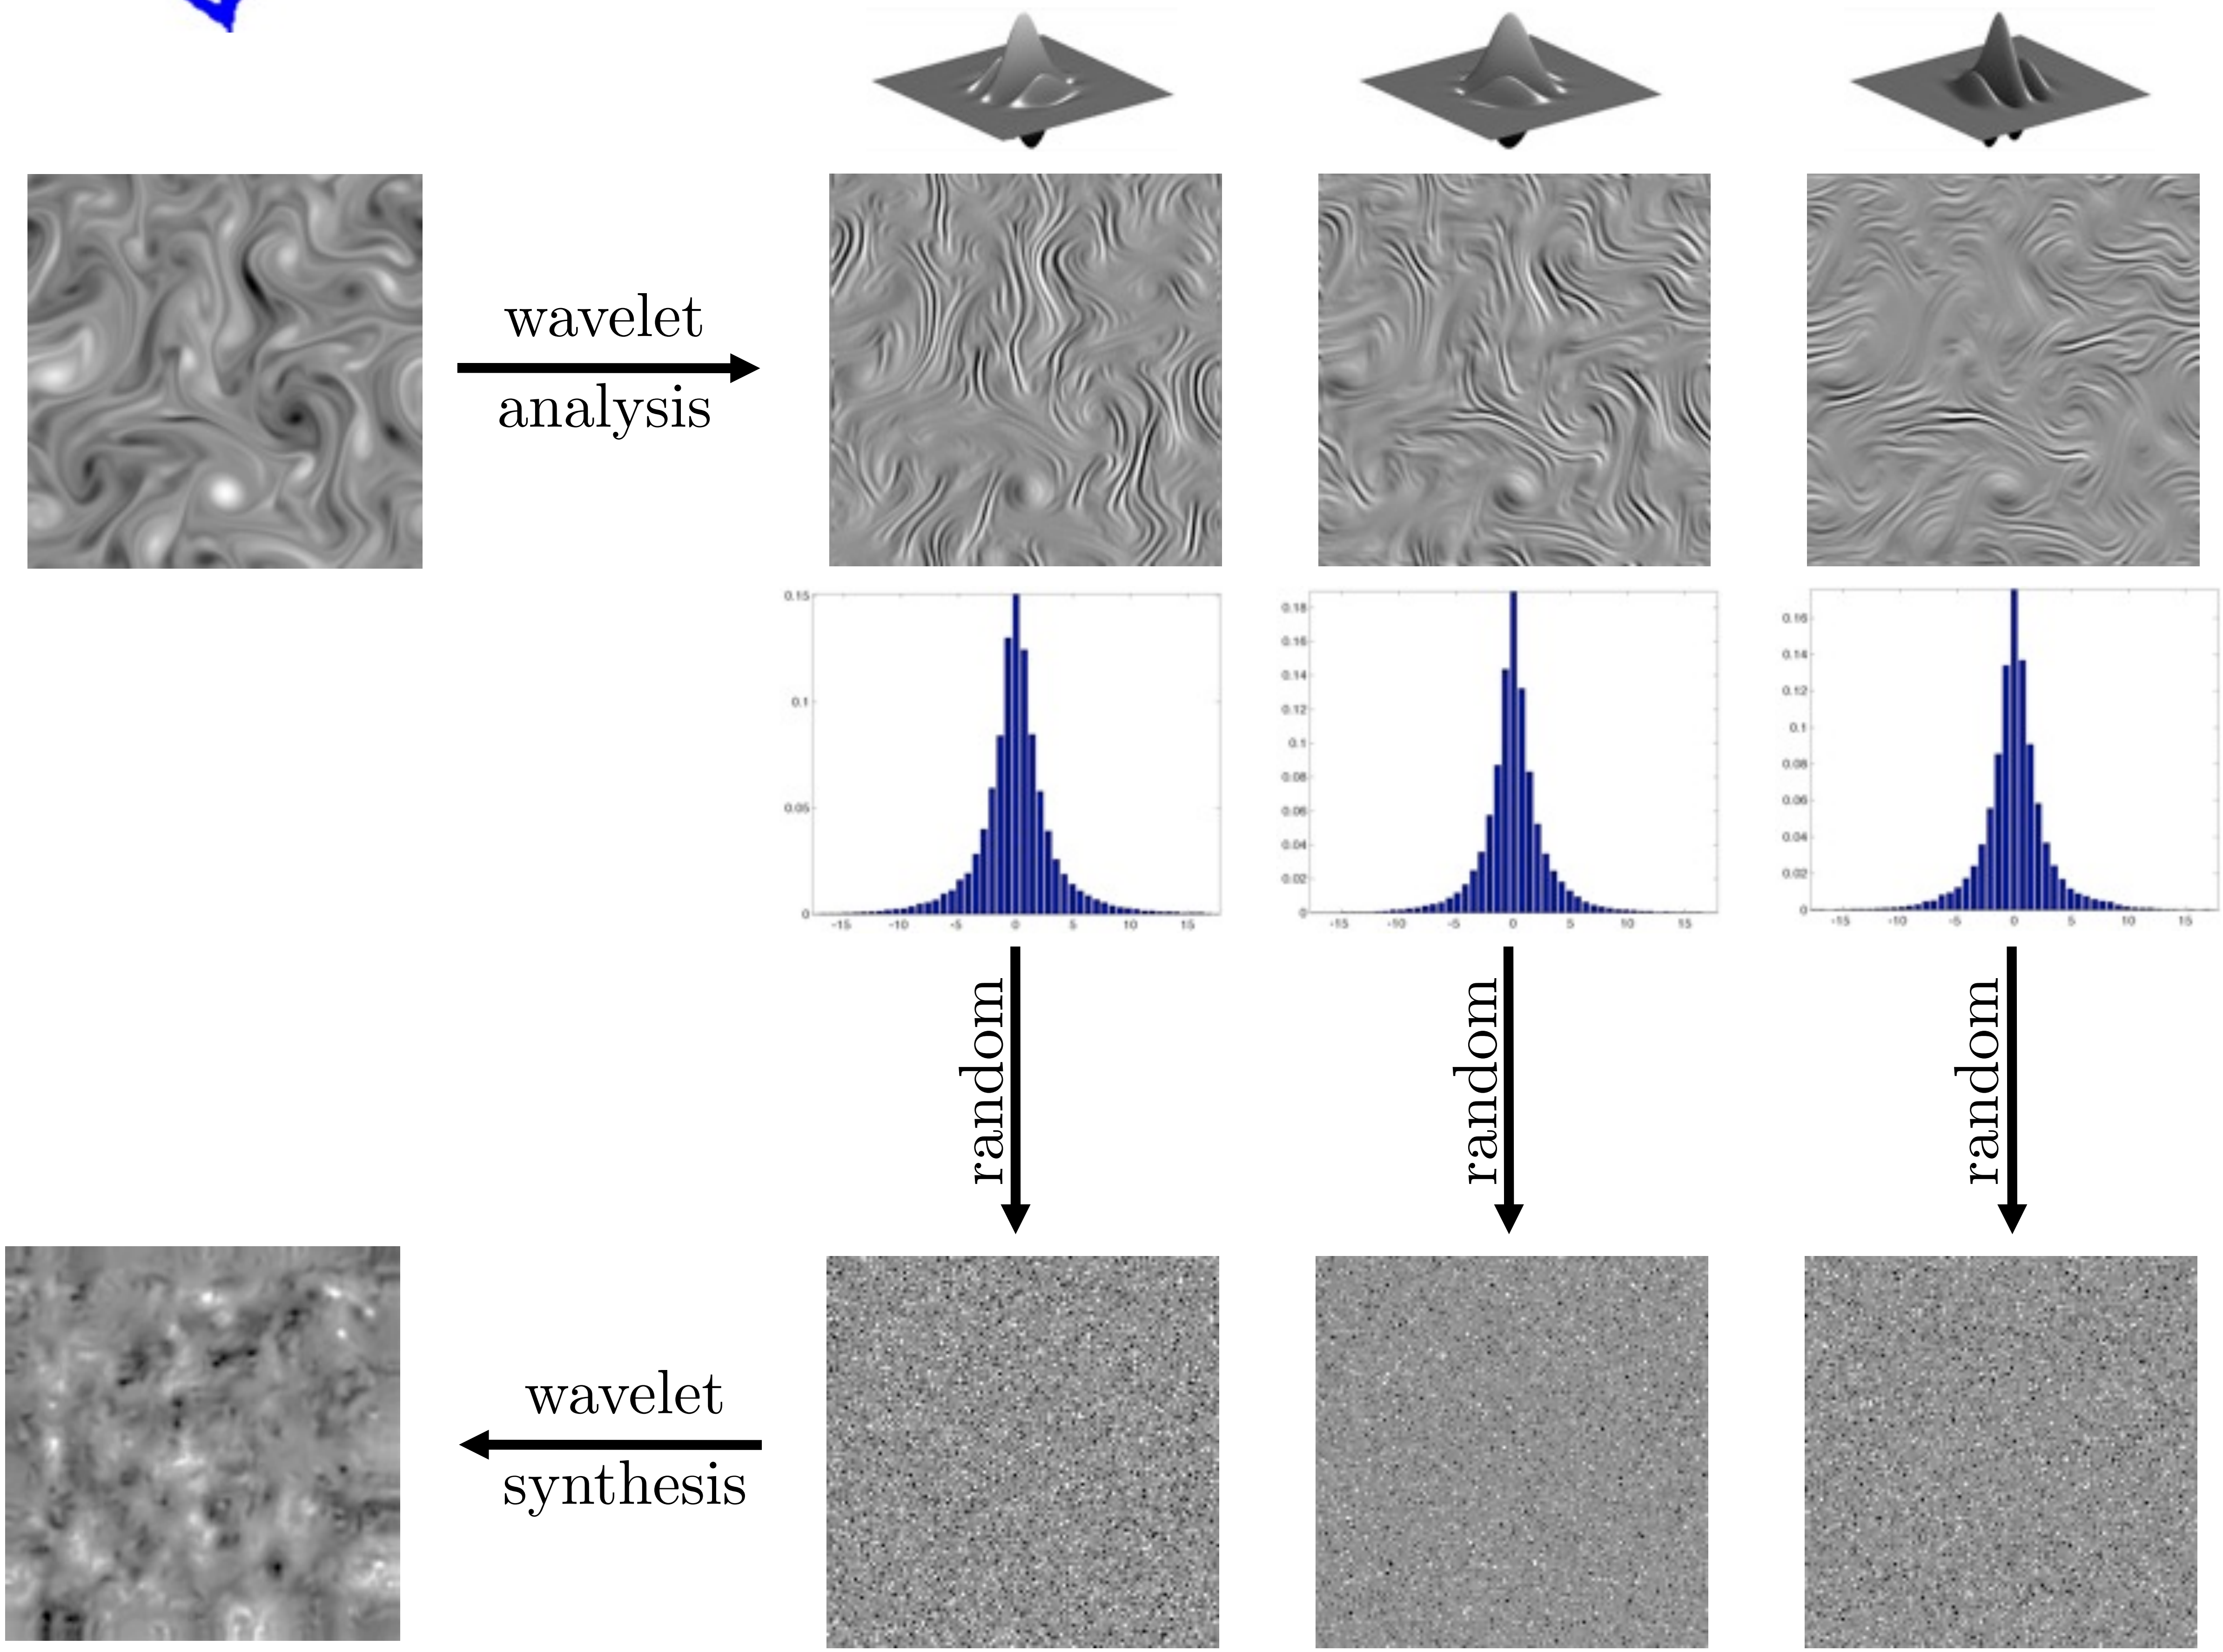

Grouplets Texture Synthesis

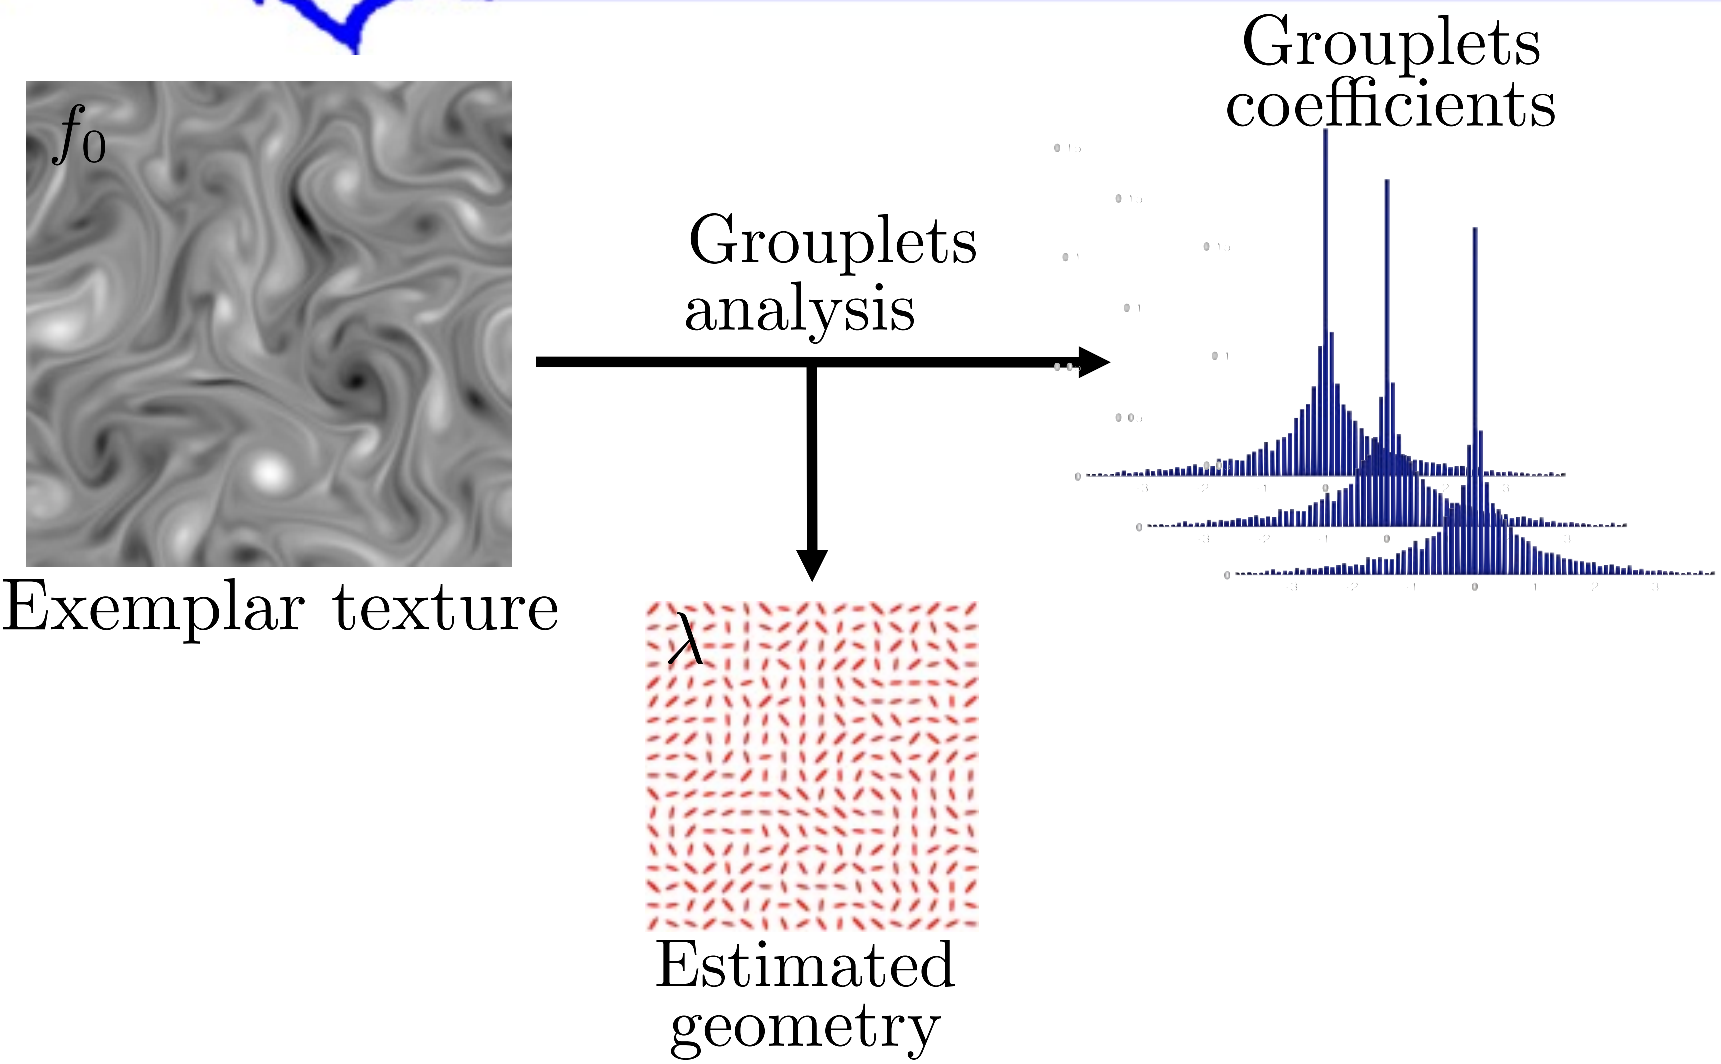

Analysis of the geometry: estimate λ

[Peyré, 2010]

Grouplets Texture Synthesis

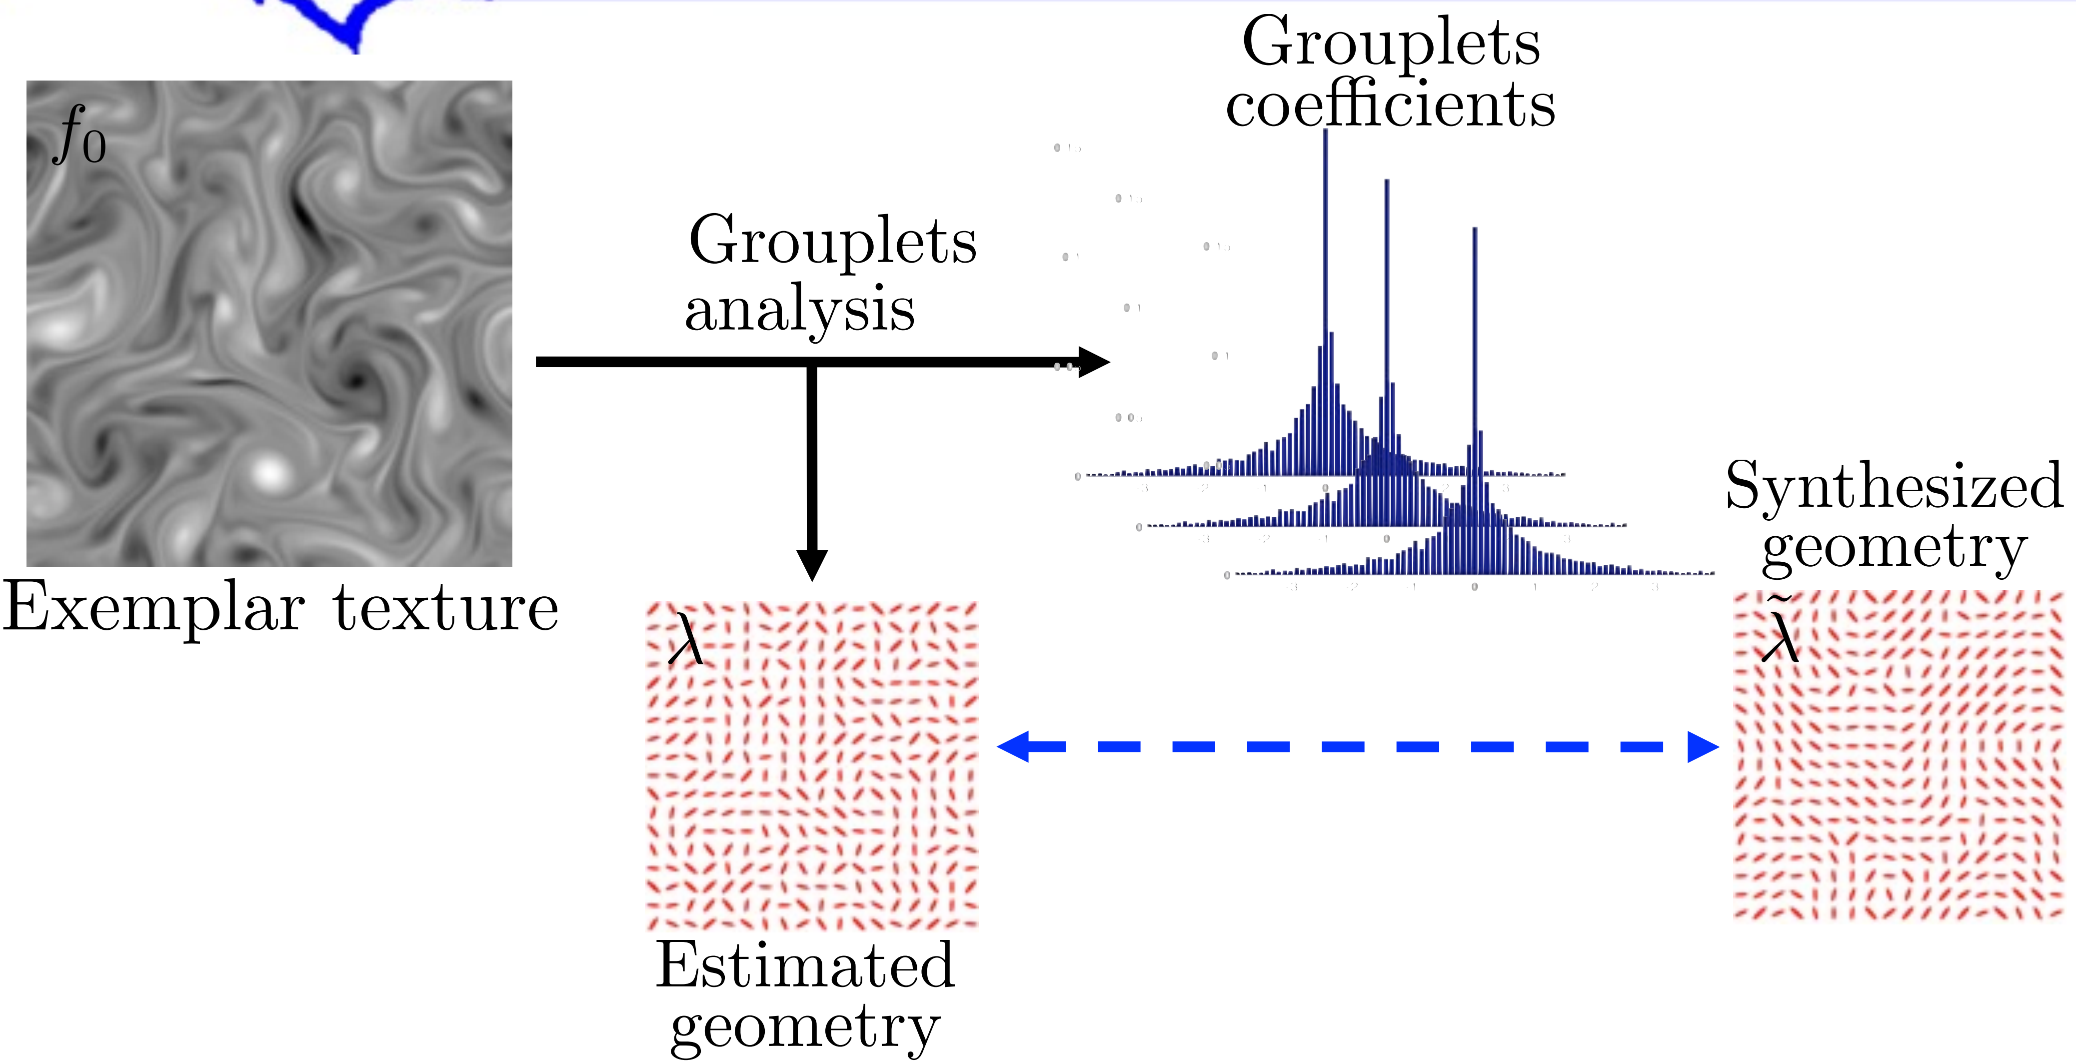

Analysis of the geometry: estimate λ
[Peyré, 2010]

Synthesis of the geometry: $\lambda \longleftrightarrow \tilde{\lambda}$

– Given by the user.

– Statistical model of (multiscale) association fields.

Grouplets Texture Synthesis

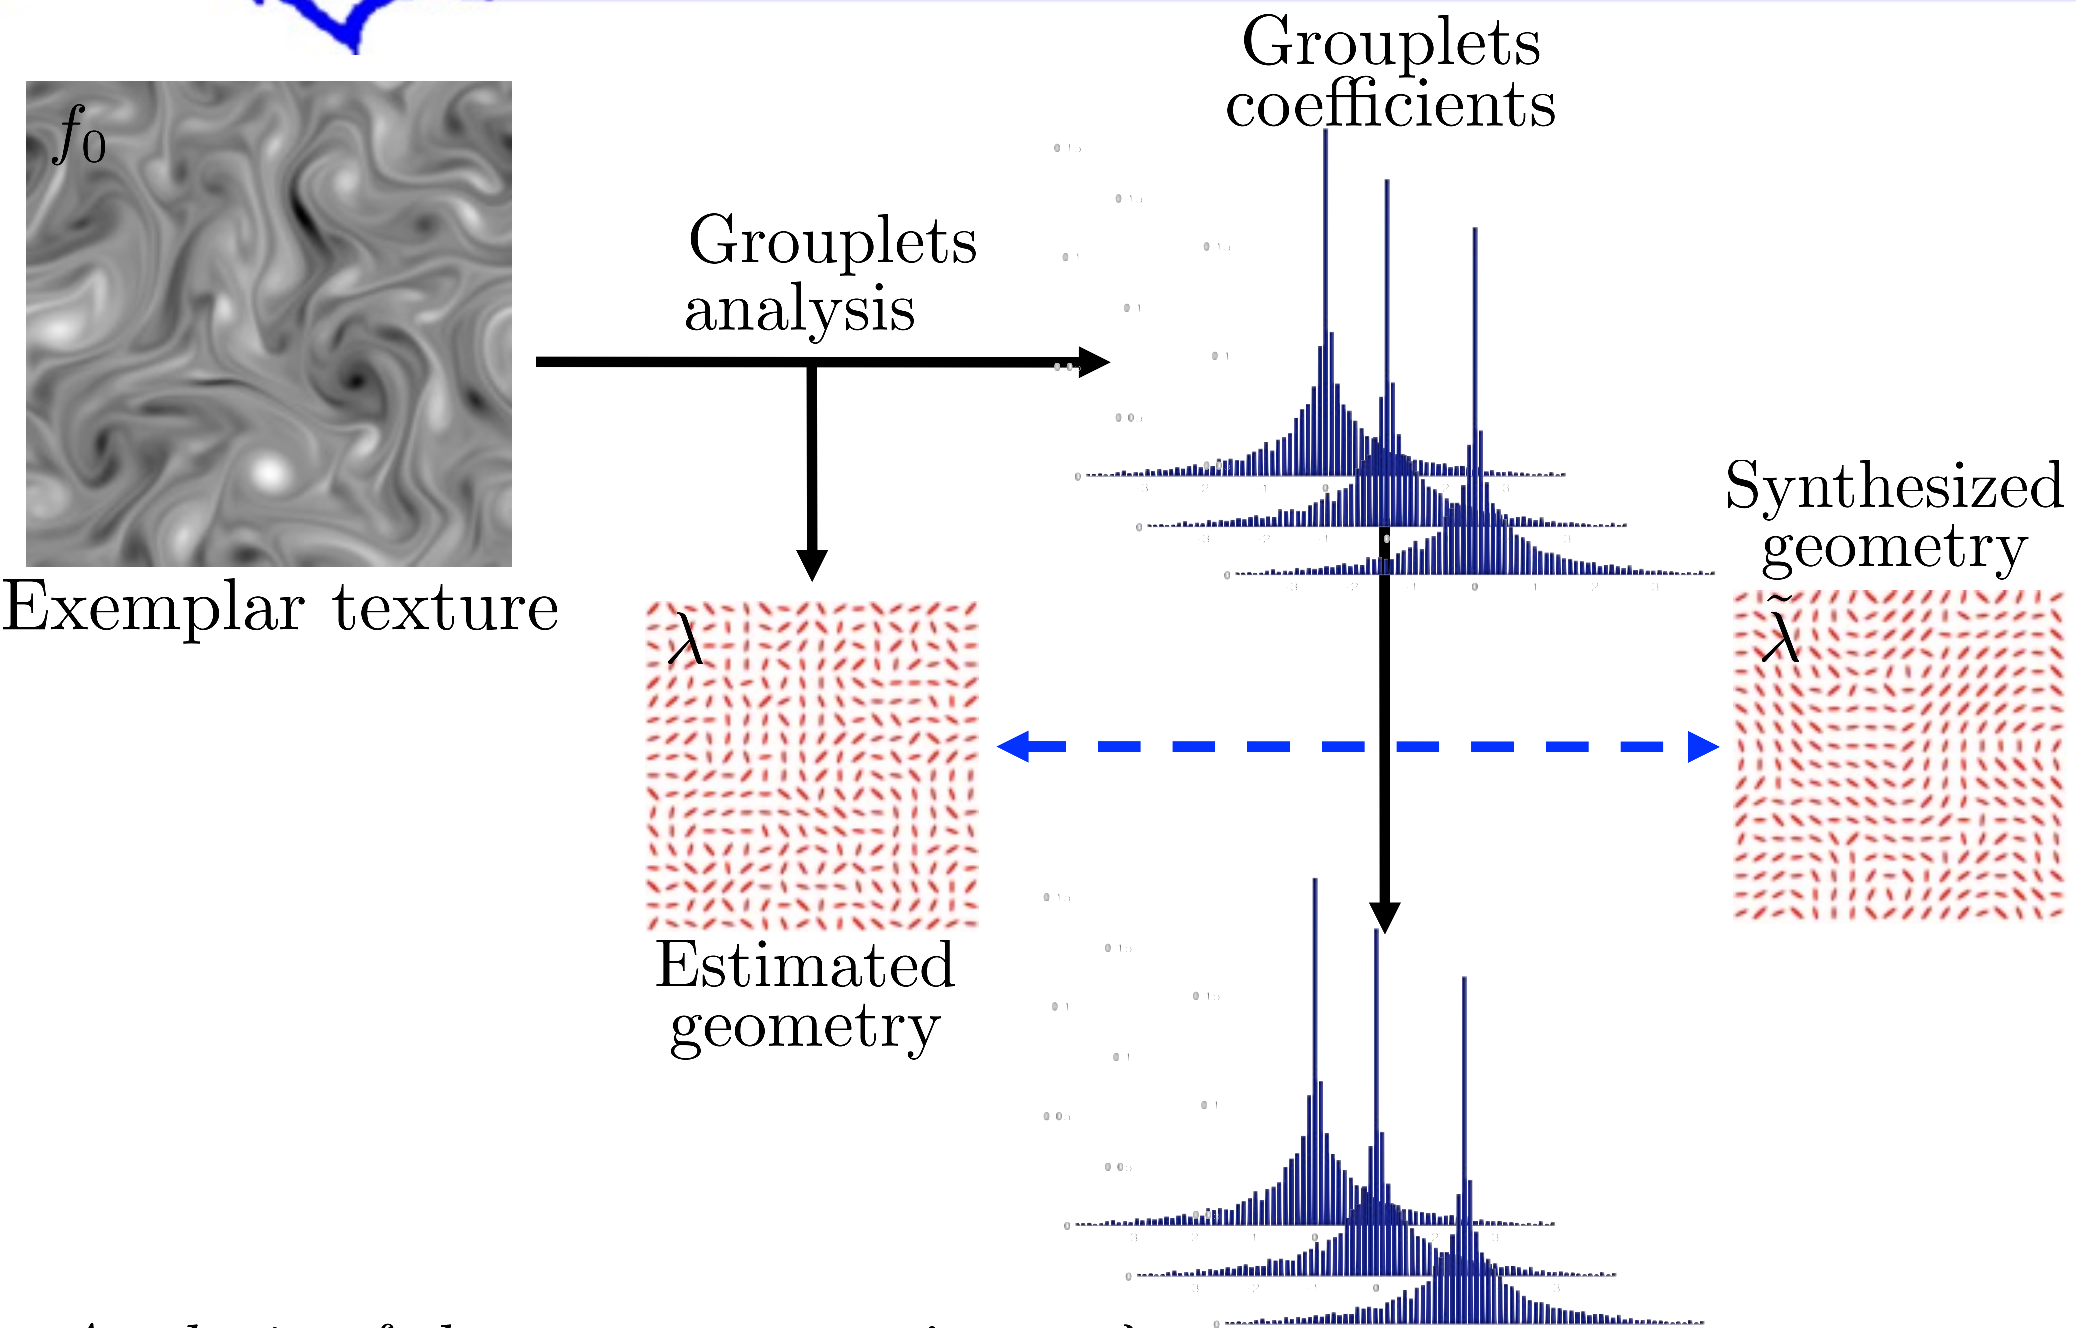

Analysis of the geometry: estimate λ
[Peyré, 2010]

Synthesis of the geometry: $\lambda \longleftrightarrow \tilde{\lambda}$

– Given by the user.

– Statistical model of (multiscale) association fields.

Grouplets Texture Synthesis

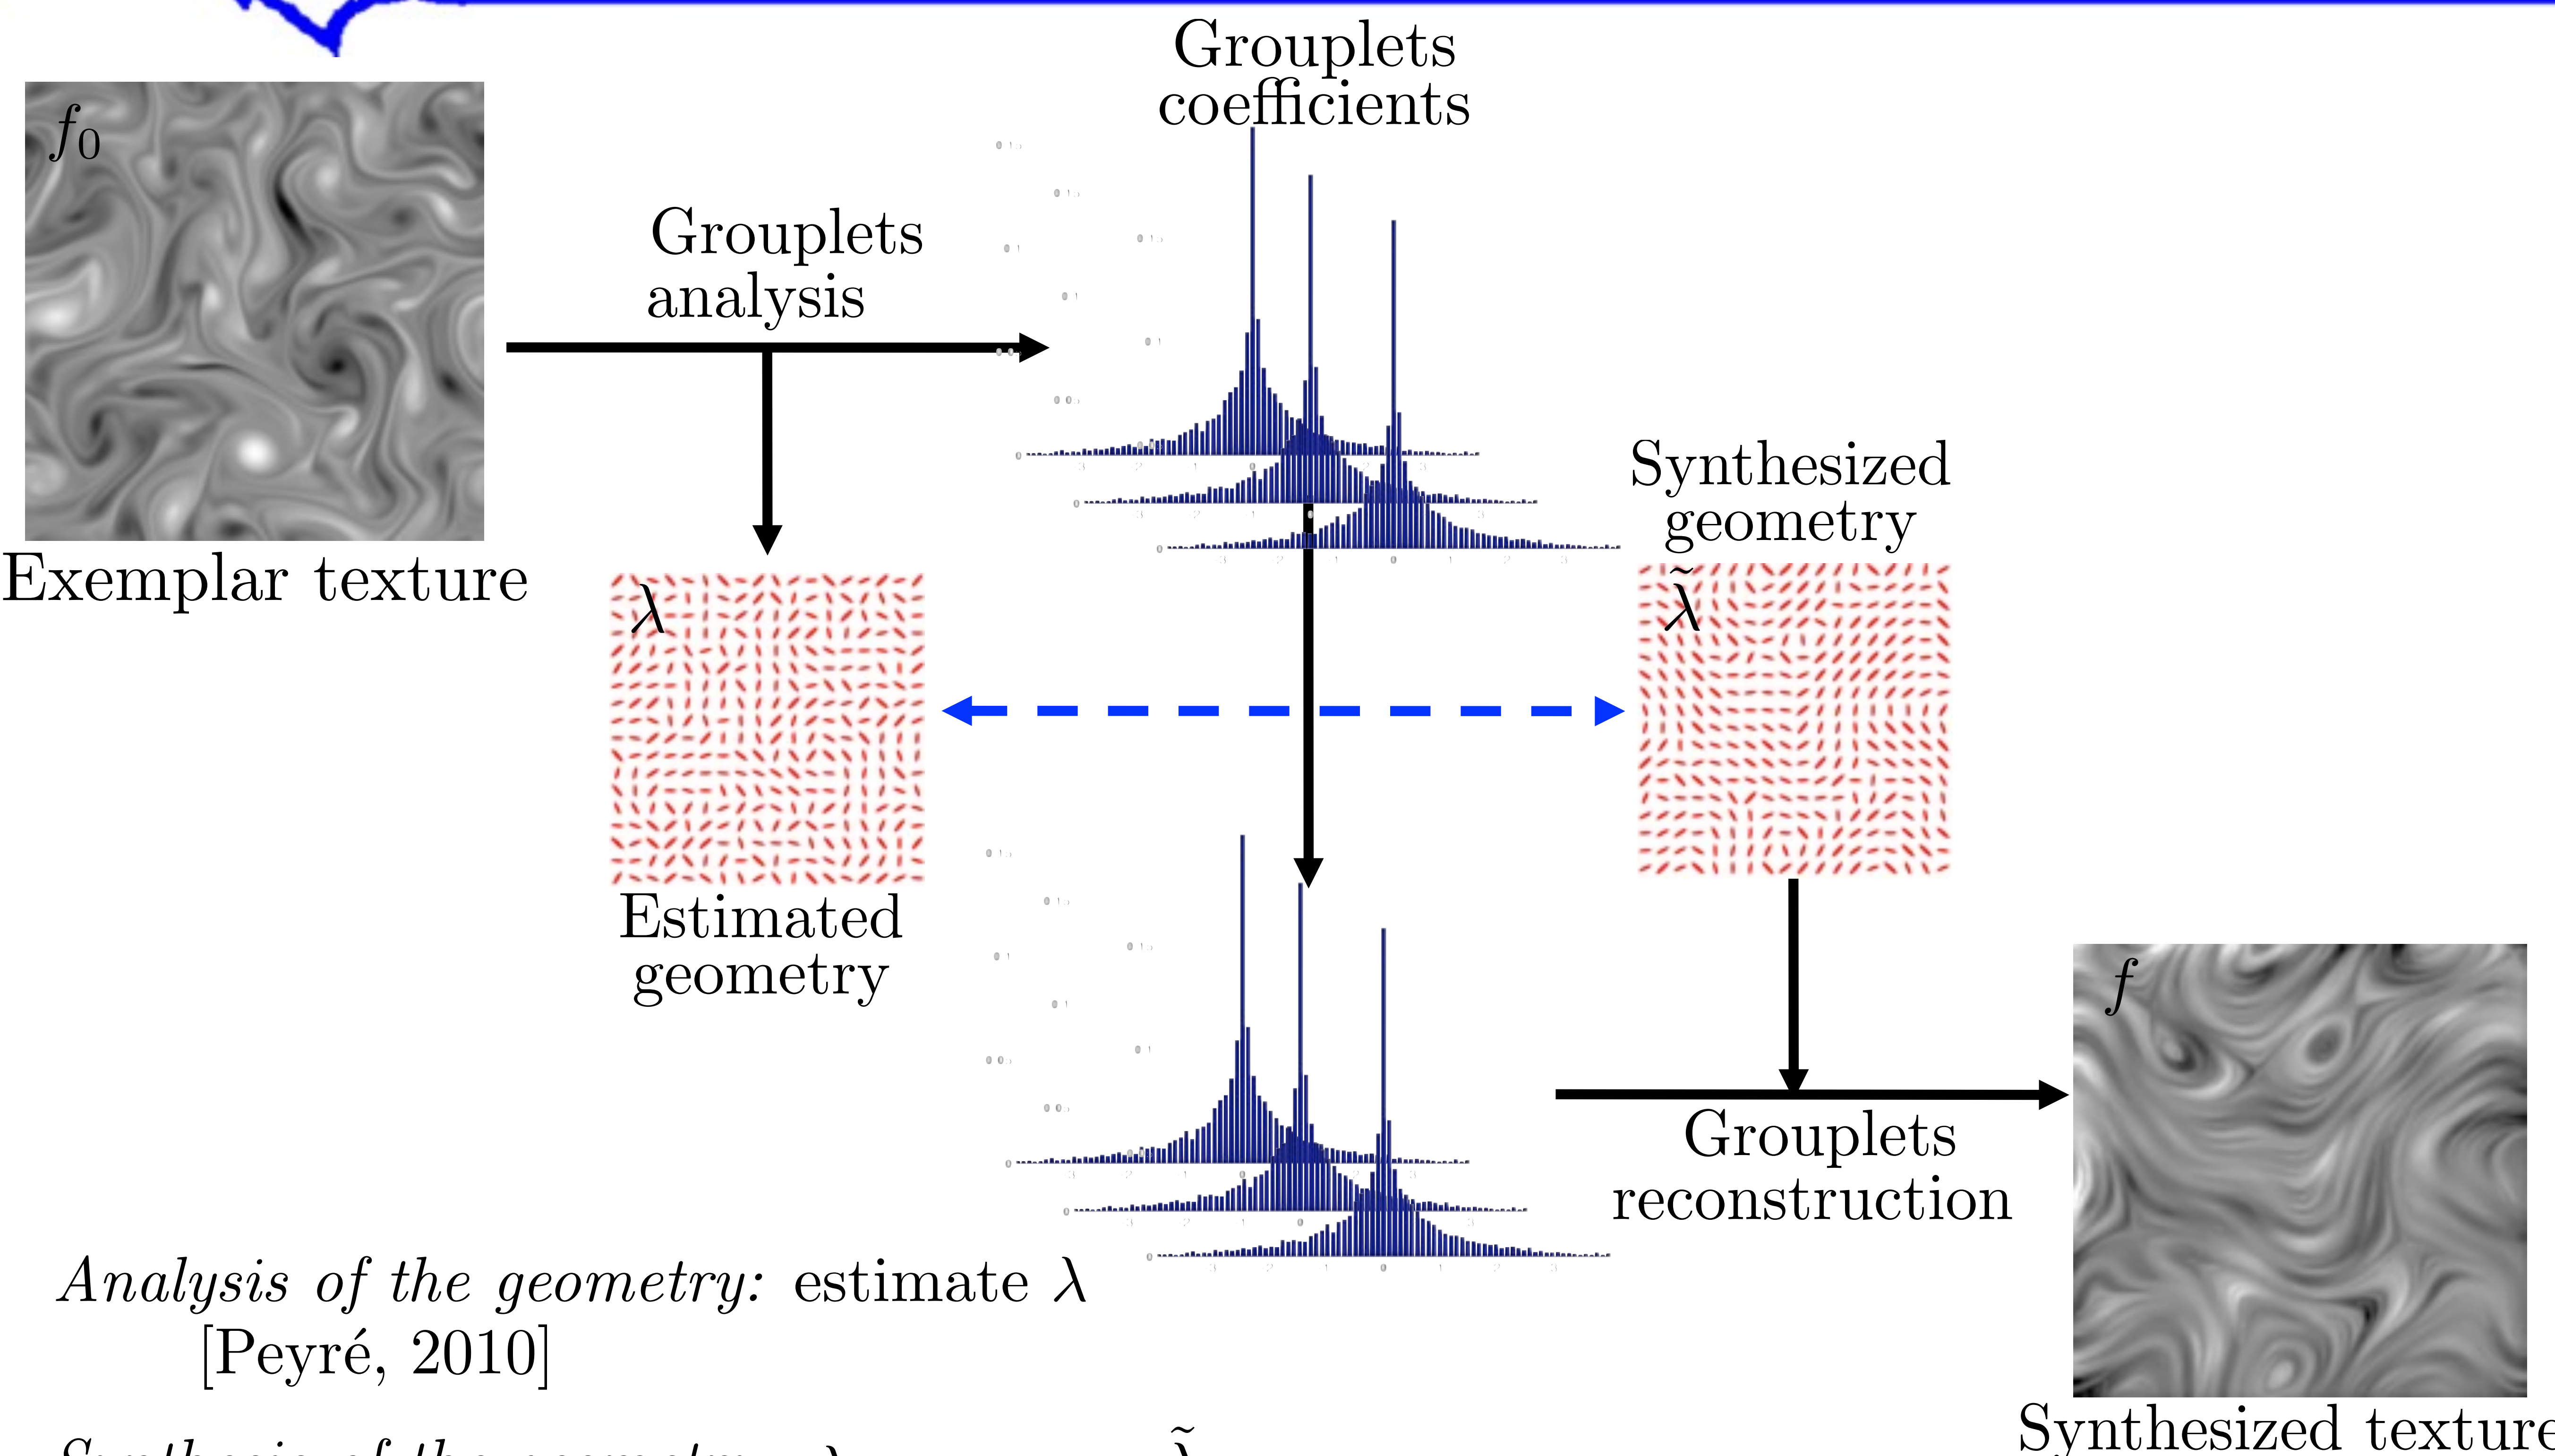

Analysis of the geometry: estimate λ
[Peyré, 2010]

Synthesis of the geometry: $\lambda \longleftrightarrow \tilde{\lambda}$

– Given by the user.

– Statistical model of (multiscale) association fields.

Grouplets Texture Synthesis

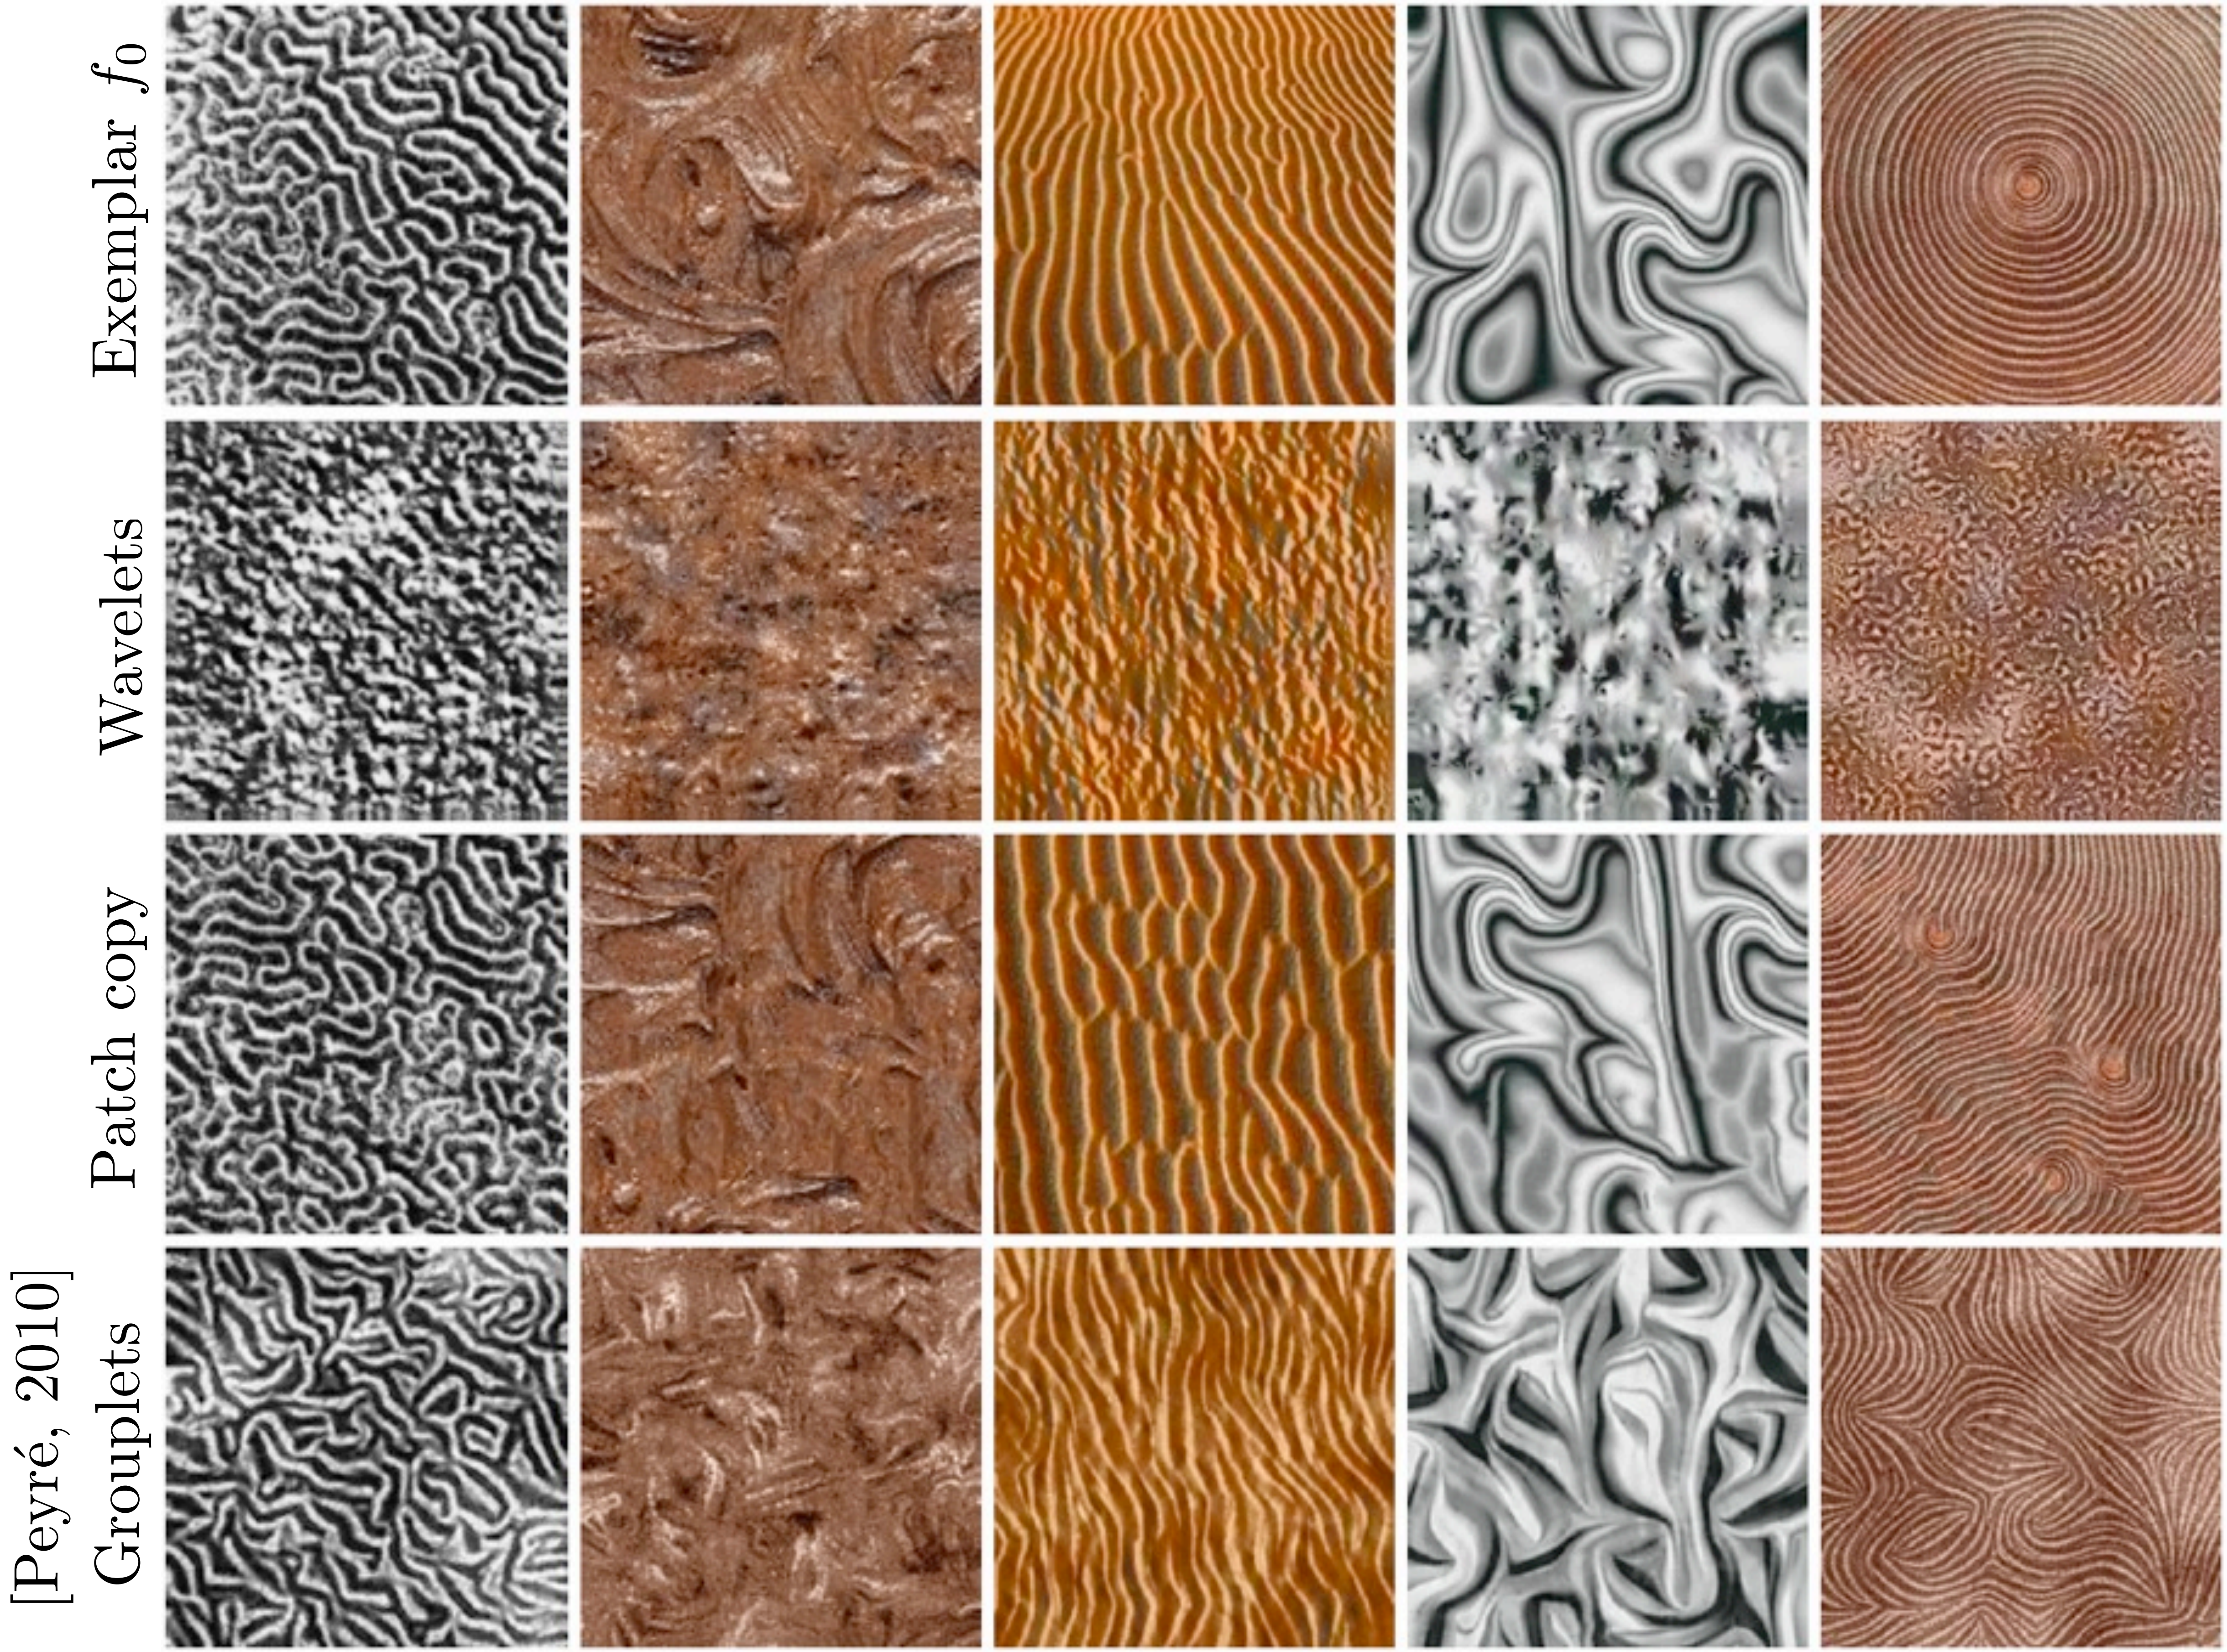

Texture Mixing

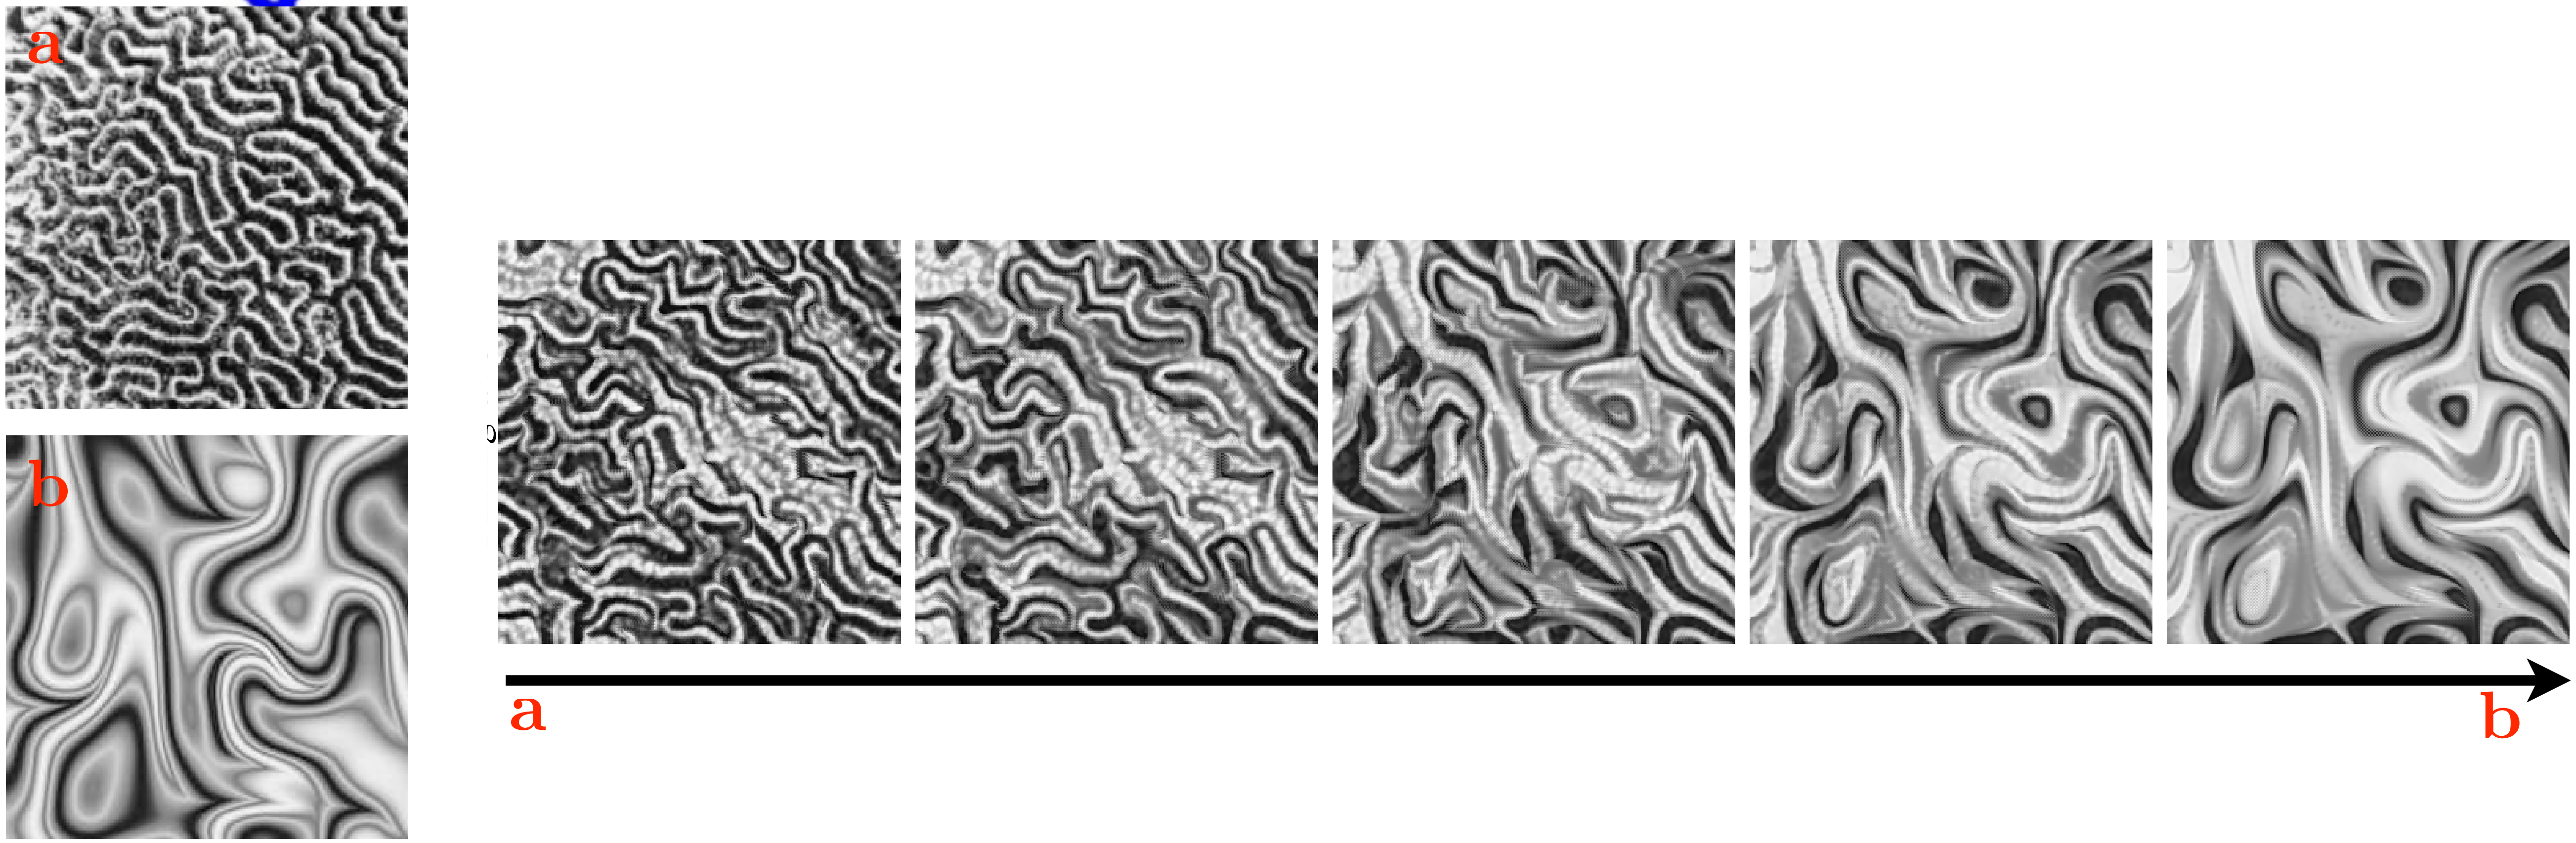

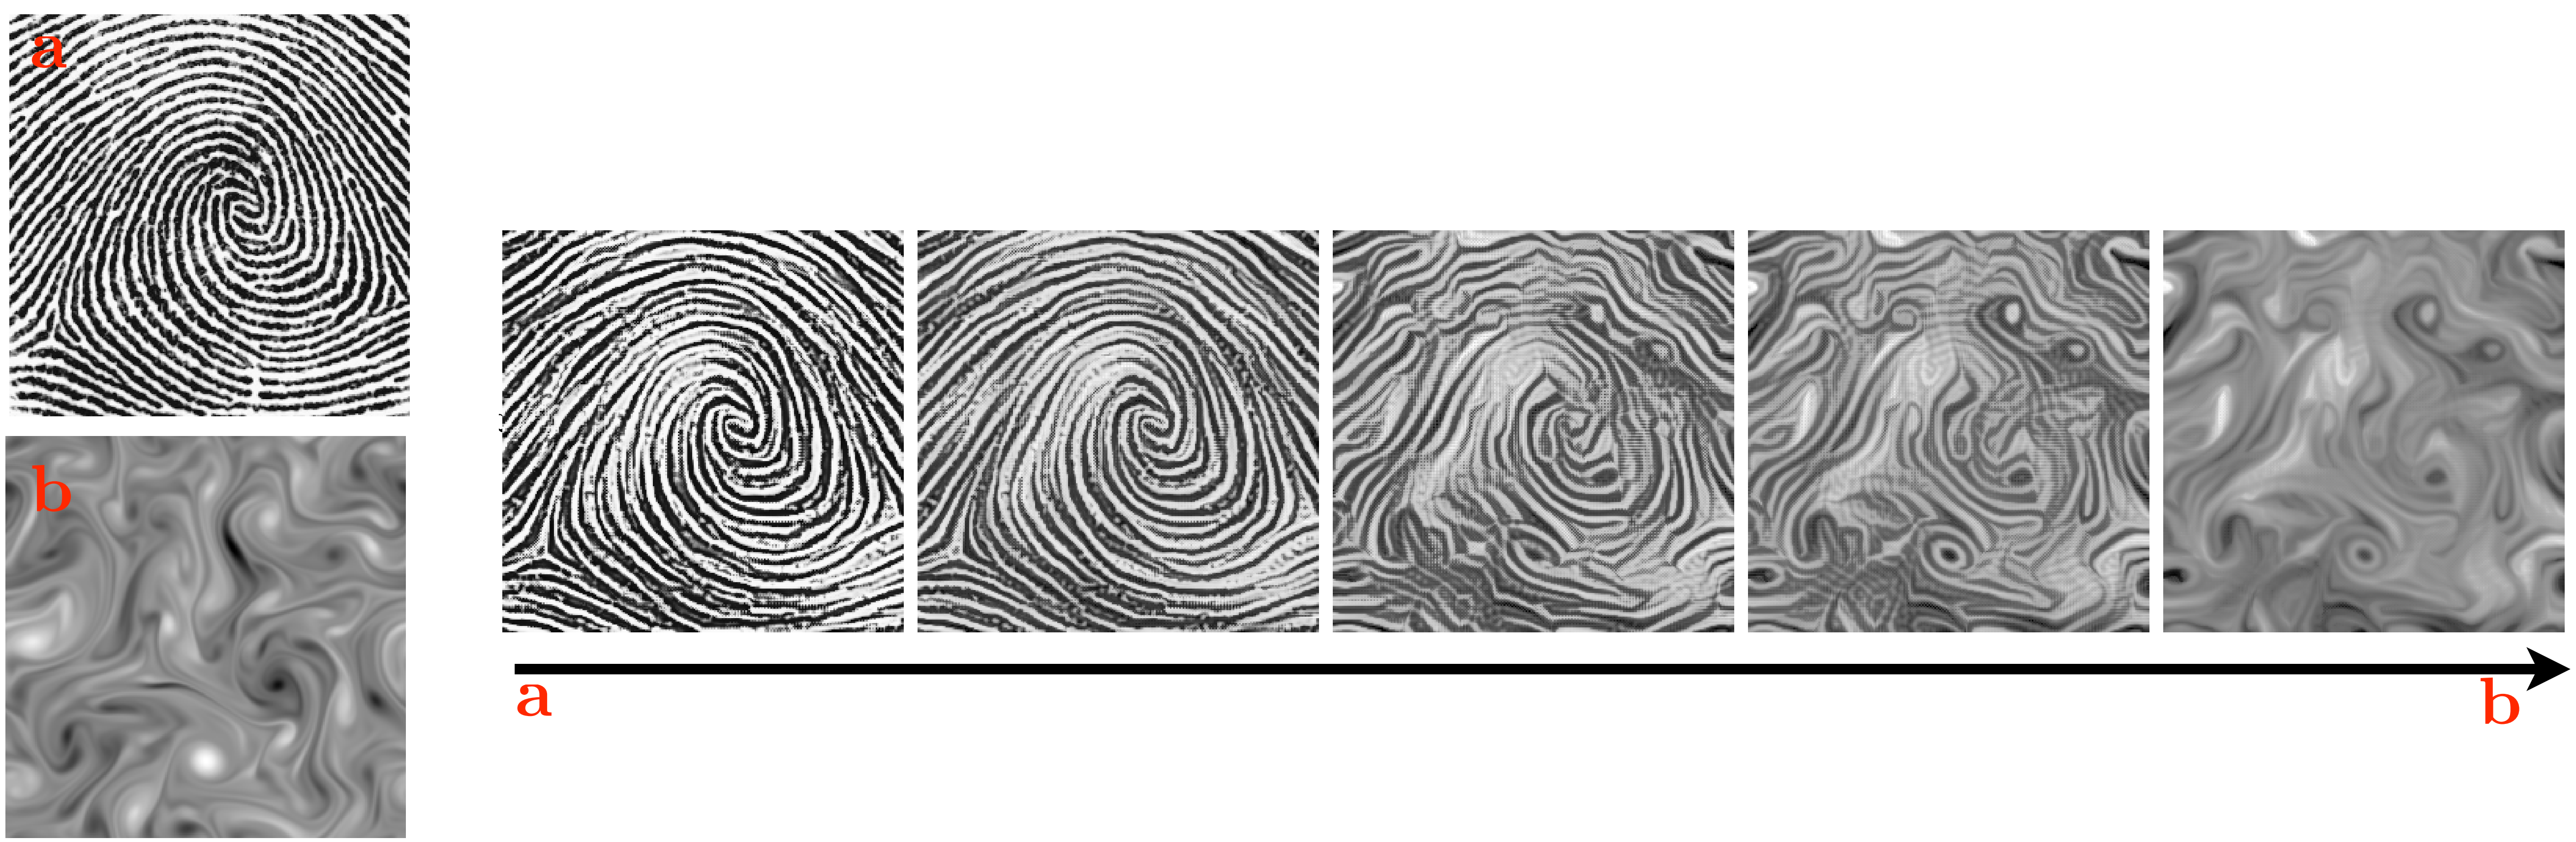

Conclusion

- *Sparsity:* approximate signals with few atoms.

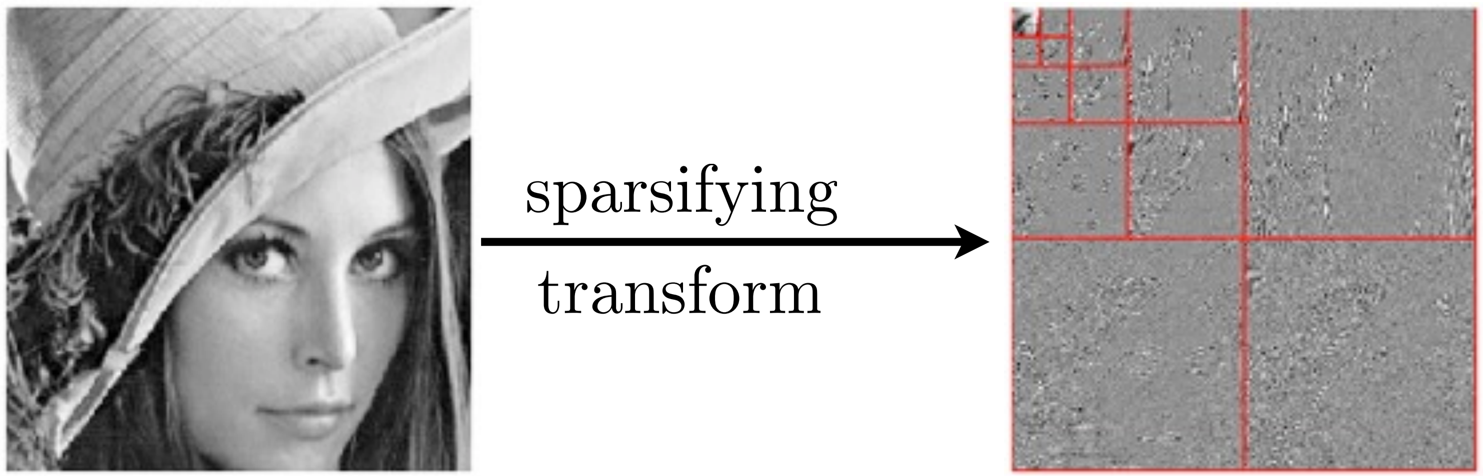

Conclusion

- *Sparsity:* approximate signals with few atoms.

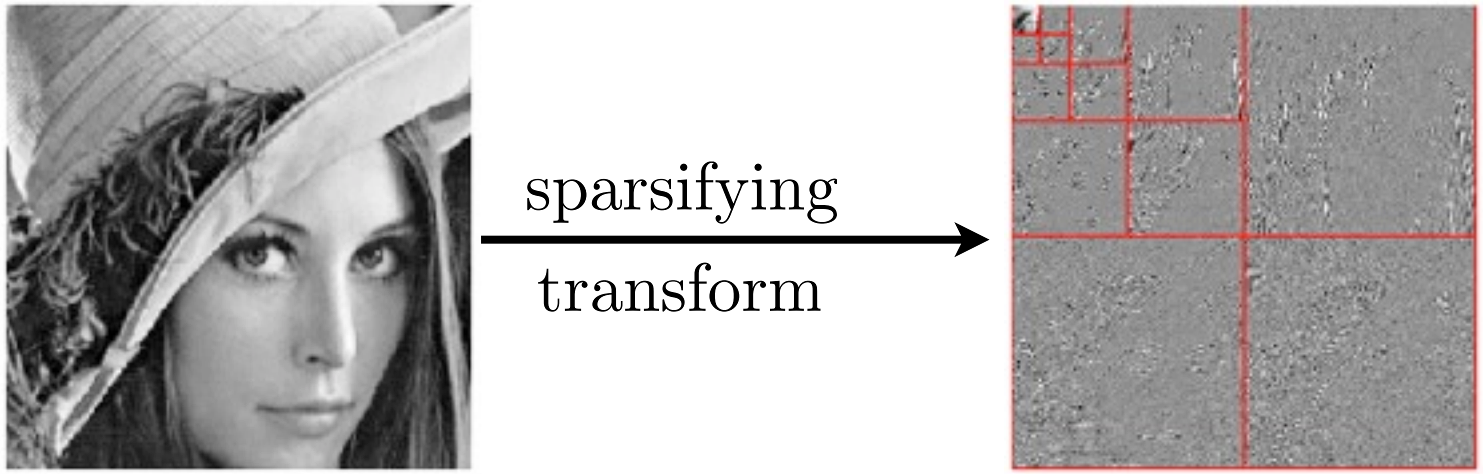

- *Sparse approximation:* compression / denoising.

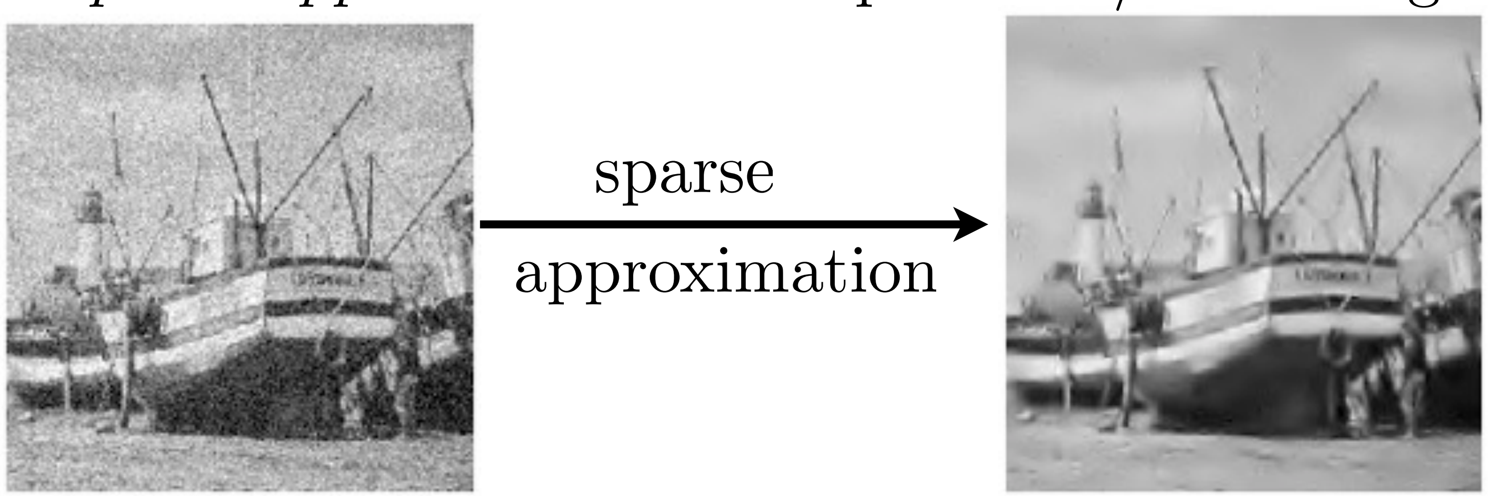

Conclusion

- *Sparsity:* approximate signals with few atoms.

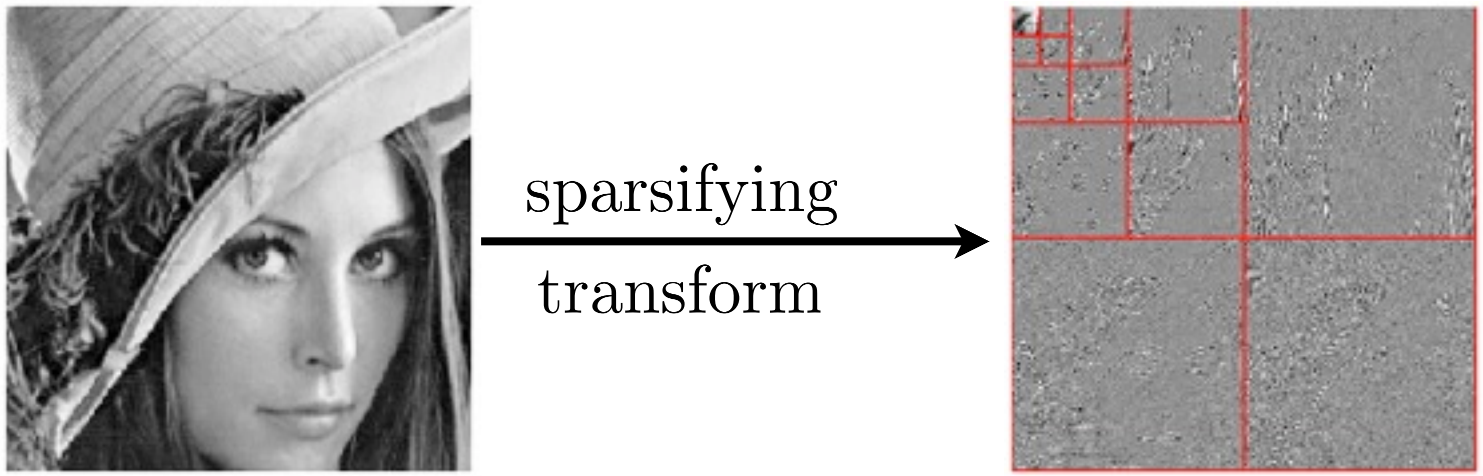

- *Sparse approximation:* compression / denoising.

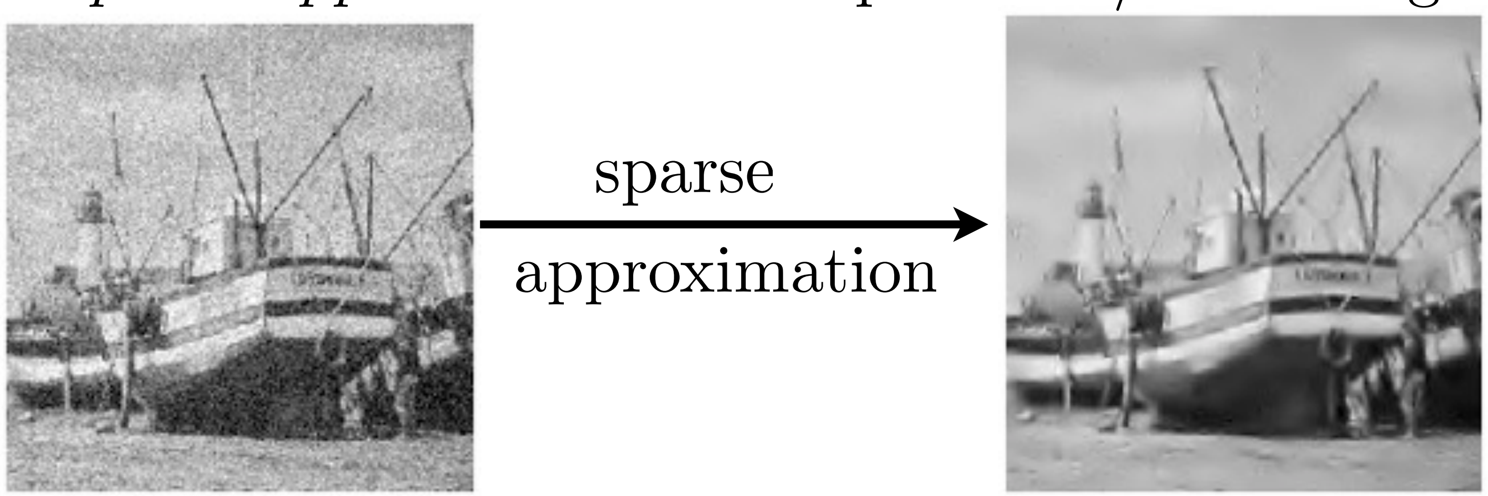

- *Quest for the best representation:*

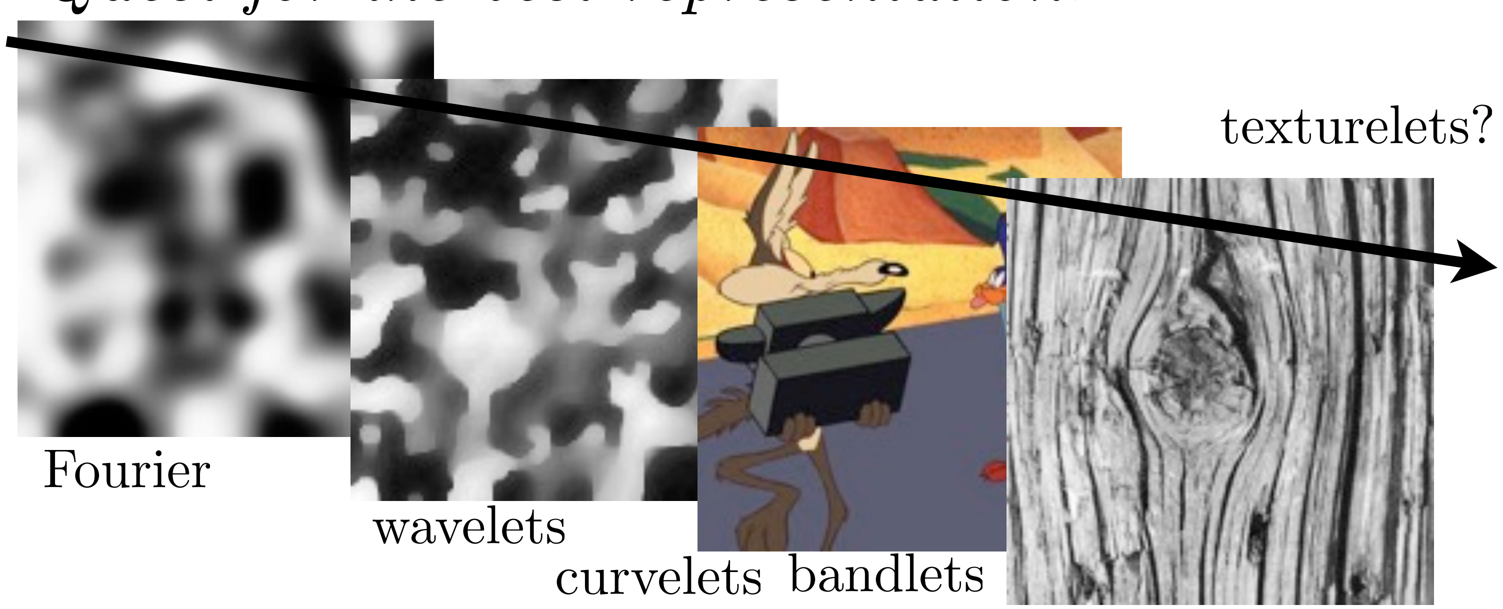

Conclusion

- *Sparsity:* approximate signals with few atoms.

sparsifying
transform

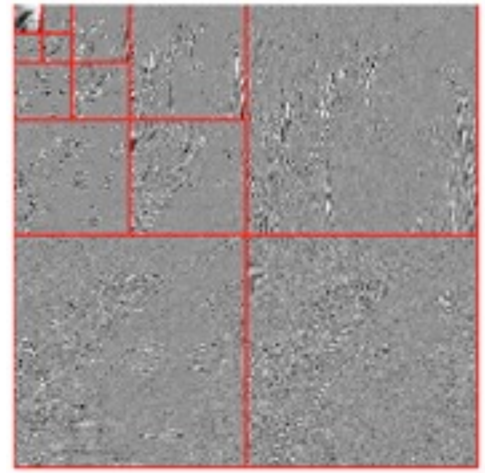

- *Sparse approximation:* compression / denoising.

sparse
approximation

- *Quest for the best representation:*

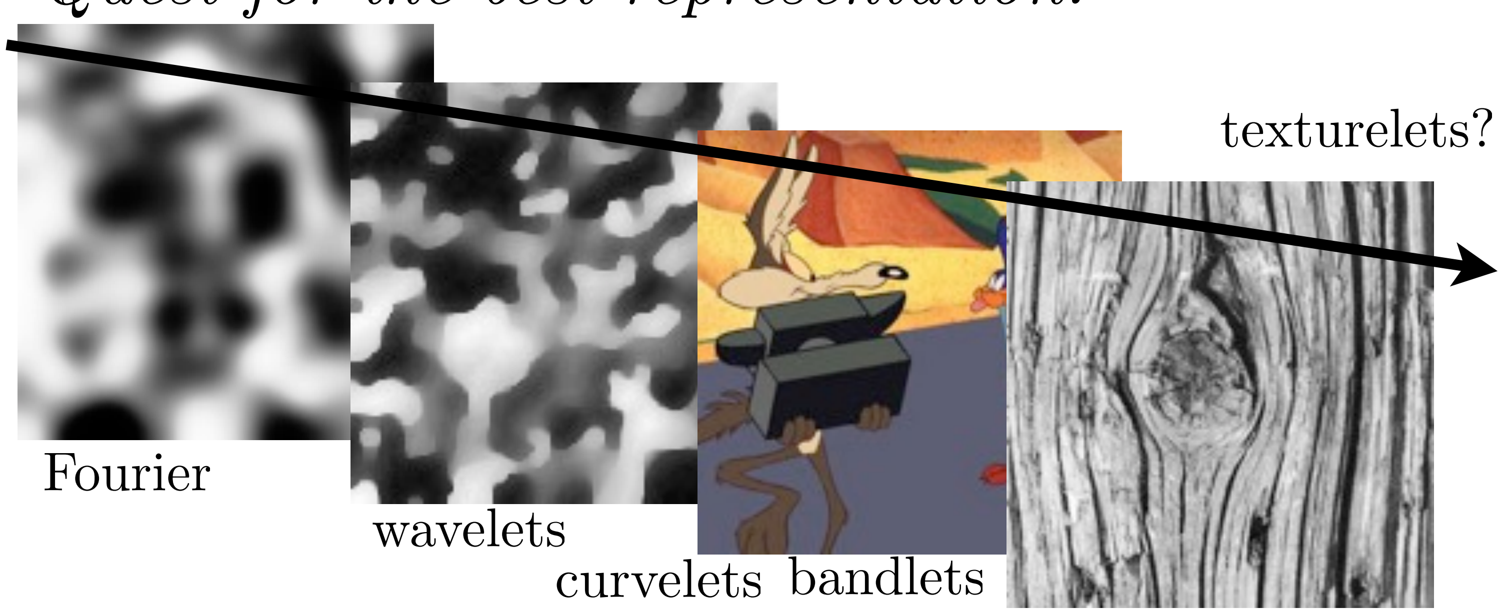

- *Inverse problems regularization:*

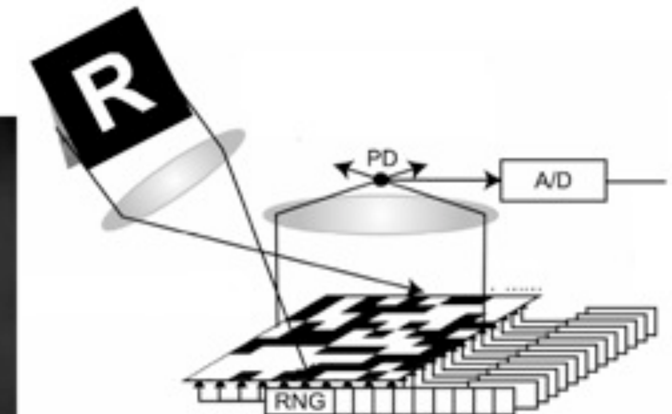

Compressed sensing

Convex sparsity prior: ℓ^1

$$\sum_m |\langle f,\, \psi_m \rangle|$$

More sparsity $\Rightarrow$ better prior

$\Rightarrow$ better recovery.

Conclusion

- *Sparsity:* approximate signals with few atoms.

sparsifying $\longrightarrow$ transform

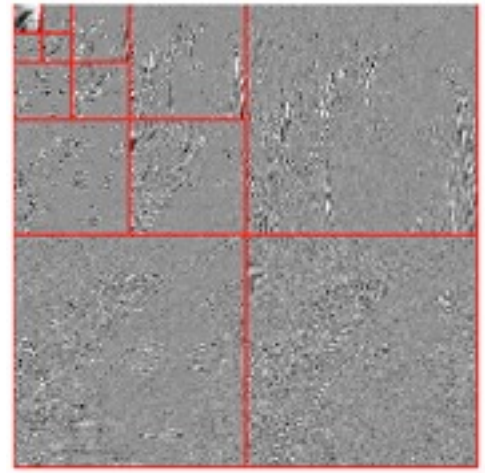

- *Sparse approximation:* compression / denoising.

sparse $\longrightarrow$ approximation

- *Quest for the best representation:*

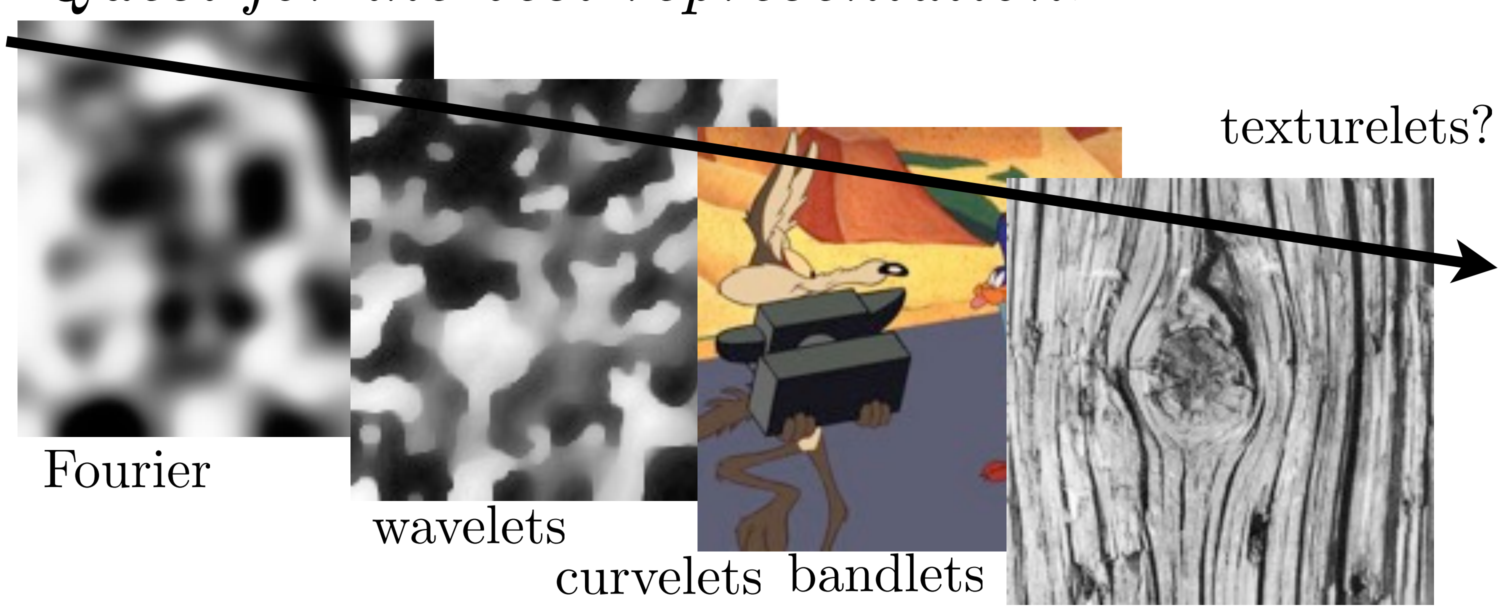

- *Inverse problems regularization:*

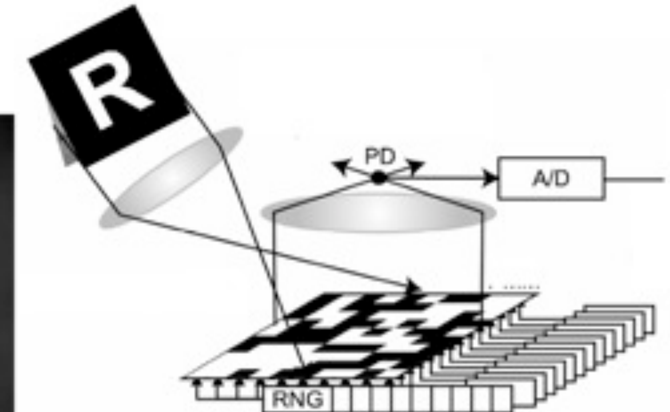

Compressed sensing

Convex sparsity prior: ℓ^1

$$\sum_m |\langle f,\, \psi_m \rangle|$$

More sparsity $\Rightarrow$ better prior

$\Rightarrow$ better recovery.

- *Texture synthesis:*

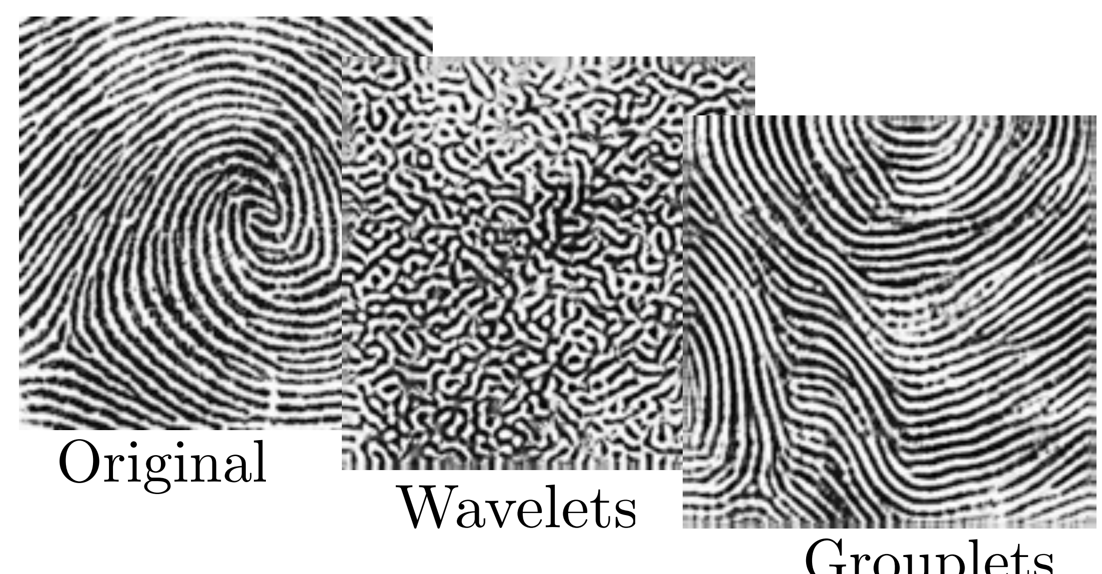

More sparsity $\Rightarrow$ better synthesis