Manifold Models
for Natural Images and Textures

Gabriel Peyré

CEREMADE
Université Paris Dauphine

Fixed and Adaptive Image Priors

Variational image prior: $J(f)$ depends on ∇f.

Sparsity in basis $\mathcal{B} = \{\psi_m\}_m$: $J(f) = \sum_m |\langle f, \psi_m \rangle|$.

Fixed and Adaptive Image Priors

Variational image prior: $J(f)$ depends on ∇f.

Sparsity in basis $\mathcal{B} = \{\psi_m\}_m$: $J(f) = \sum_m |\langle f, \psi_m \rangle|$.

Uniformly smooth C^α image.

$$J(f) = \int |\nabla f|^2, \quad \mathcal{B} = \text{Fourier}.$$

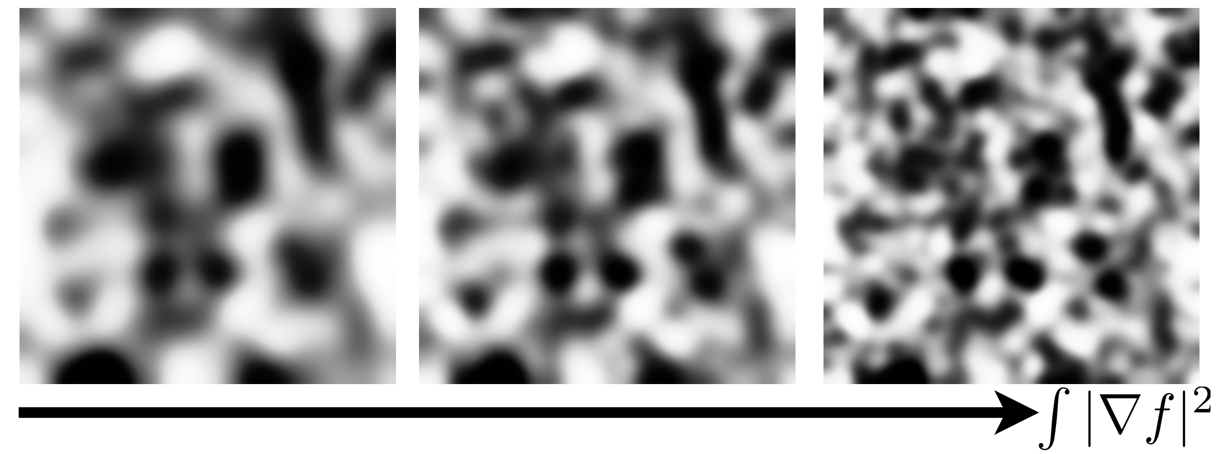

Fixed and Adaptive Image Priors

Variational image prior: $J(f)$ depends on ∇f.

Sparsity in basis $\mathcal{B} = \{\psi_m\}_m$: $J(f) = \sum_m |\langle f,\, \psi_m \rangle|$.

Uniformly smooth C^α image.

$J(f) = \int |\nabla f|^2, \quad \mathcal{B} =$Fourier.

Discontinuous image with bounded variation.

$J(f) = \int |\nabla f|, \quad \mathcal{B} =$wavelets.

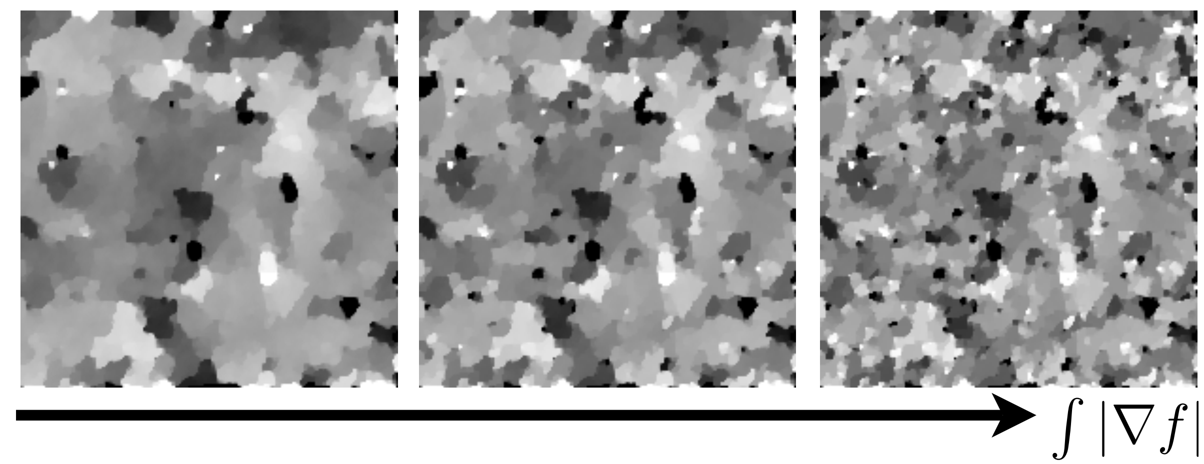

Fixed and Adaptive Image Priors

Variational image prior: $J(f)$ depends on ∇f.

Sparsity in basis $\mathcal{B} = \{\psi_m\}_m$: $J(f) = \sum_m |\langle f,\, \psi_m \rangle|$.

Uniformly smooth C^α image.

$J(f) = \int |\nabla f|^2$, $\quad \mathcal{B} =$Fourier.

Discontinuous image with bounded variation.

$J(f) = \int |\nabla f|$, $\quad \mathcal{B} =$wavelets.

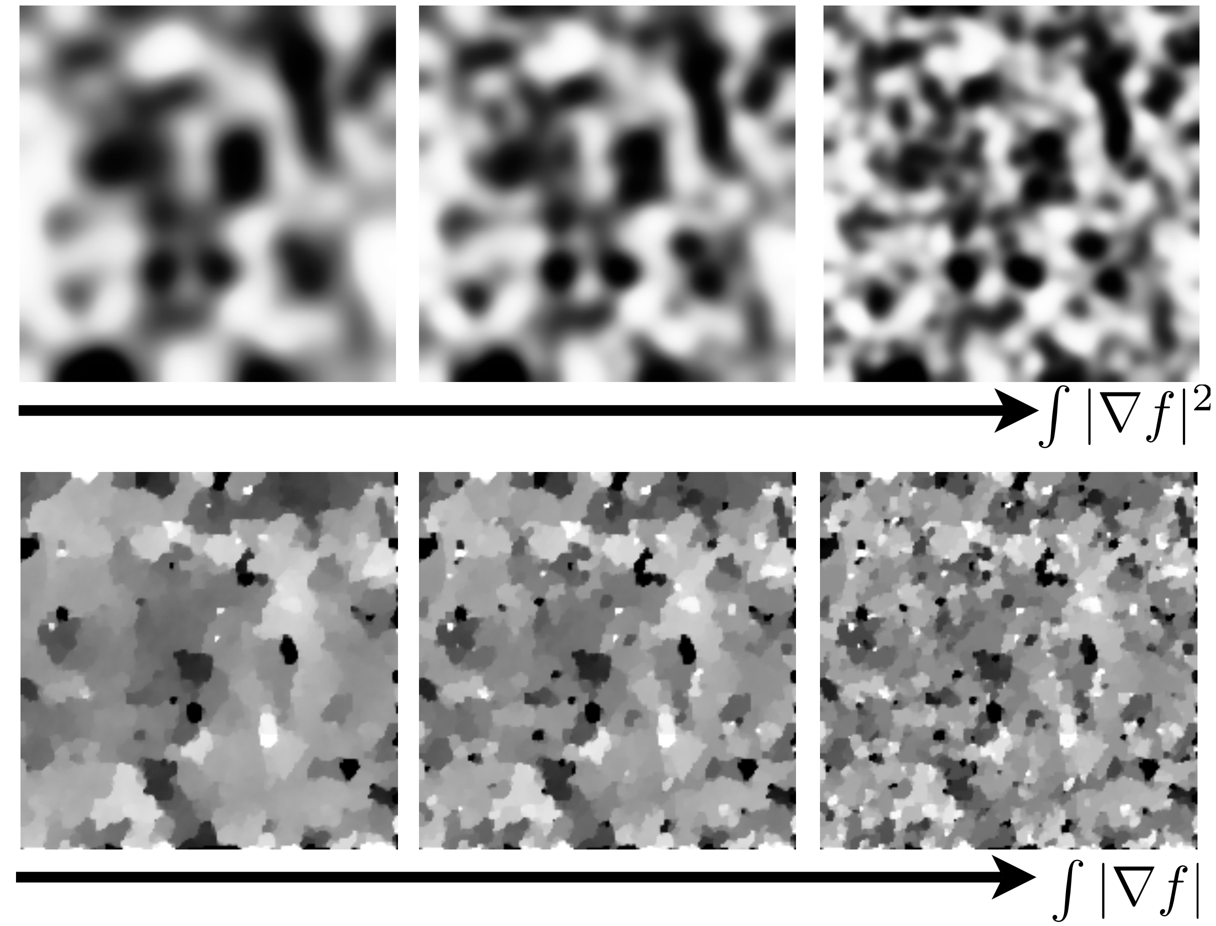

Adaptive prior: $J(f) = J_w(f)$, $w =$geometry.

– denoising: w estimated from noisy image.

– inverse pbm: adapt w to the image to recover

Fixed and Adaptive Image Priors

Variational image prior: $J(f)$ depends on ∇f.

Sparsity in basis $\mathcal{B} = \{\psi_m\}_m$: $J(f) = \sum_m |\langle f,\, \psi_m \rangle|$.

Uniformly smooth C^α image.

$$J(f) = \int |\nabla f|^2, \quad \mathcal{B} = \text{Fourier}.$$

Discontinuous image with bounded variation.

$$J(f) = \int |\nabla f|, \quad \mathcal{B} = \text{wavelets}.$$

Adaptive prior: $J(f) = J_w(f)$, w =geometry.

– denoising: w estimated from noisy image.

– inverse pbm: adapt w to the image to recover

Cartoon image with C^α edges.

$\longrightarrow$ Adaptivity to the edge **direction** w.

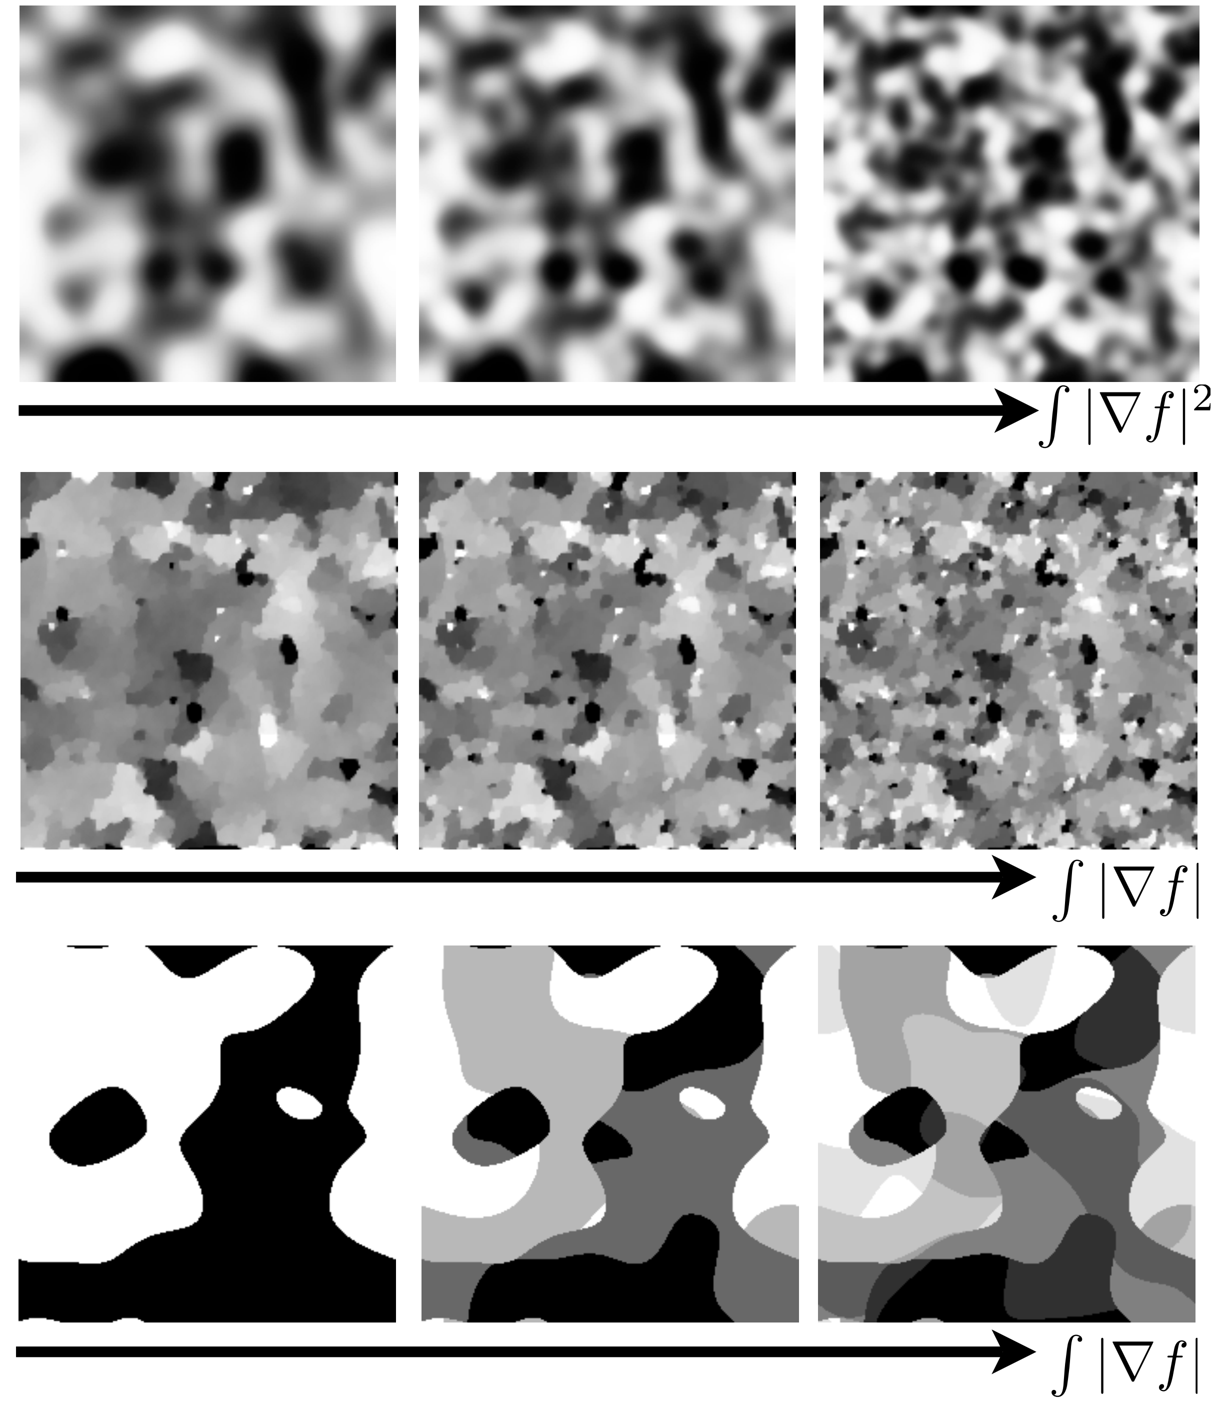

Fixed and Adaptive Image Priors

Variational image prior: $J(f)$ depends on ∇f.

Sparsity in basis $\mathcal{B} = \{\psi_m\}_m$: $J(f) = \sum_m |\langle f,\, \psi_m \rangle|$.

Uniformly smooth C^α image.

$J(f) = \int |\nabla f|^2, \quad \mathcal{B} =$Fourier.

Discontinuous image with bounded variation.

$J(f) = \int |\nabla f|, \quad \mathcal{B} =$wavelets.

Adaptive prior: $J(f) = J_w(f)$, $w =$geometry.

– denoising: w estimated from noisy image.

– inverse pbm: adapt w to the image to recover

Cartoon image with C^α edges.

$\longrightarrow$ Adaptivity to the edge **direction** w.

Locally parallel, turbulent textures.

$\longrightarrow$ Adaptivity to the texture **orientation** w.

$\int |\nabla f|^2$

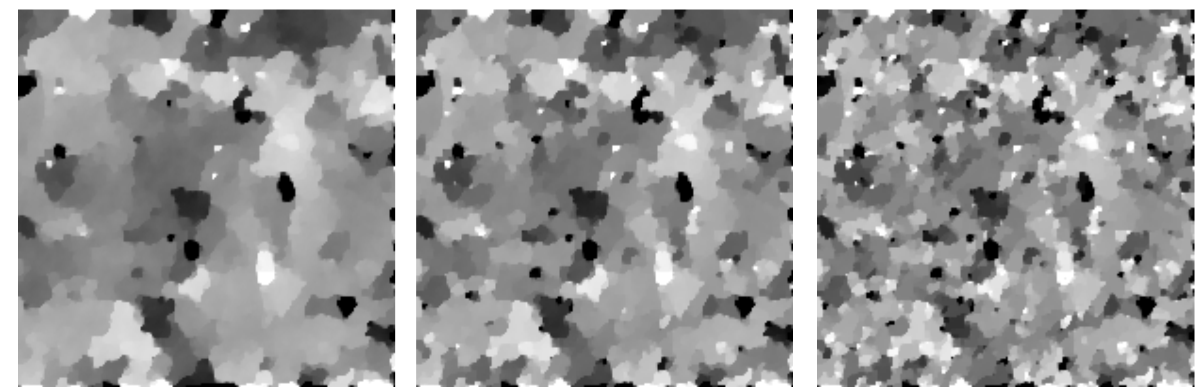

$\int |\nabla f|$

$\int |\nabla f|$

Fixed and Adaptive Image Priors

Variational image prior: $J(f)$ depends on ∇f.

Sparsity in basis $\mathcal{B} = \{\psi_m\}_m$: $J(f) = \sum_m |\langle f,\, \psi_m \rangle|$.

Uniformly smooth C^α image.

$$J(f) = \int |\nabla f|^2, \quad \mathcal{B} = \text{Fourier}.$$

Discontinuous image with bounded variation.

$$J(f) = \int |\nabla f|, \quad \mathcal{B} = \text{wavelets}.$$

Adaptive prior: $J(f) = J_w(f)$, w =geometry.

– denoising: w estimated from noisy image.

– inverse pbm: adapt w to the image to recover

Cartoon image with C^α edges.

$\longrightarrow$ Adaptivity to the edge **direction** w.

Locally parallel, turbulent textures.

$\longrightarrow$ Adaptivity to the texture **orientation** w.

Complex natural images: open question ...

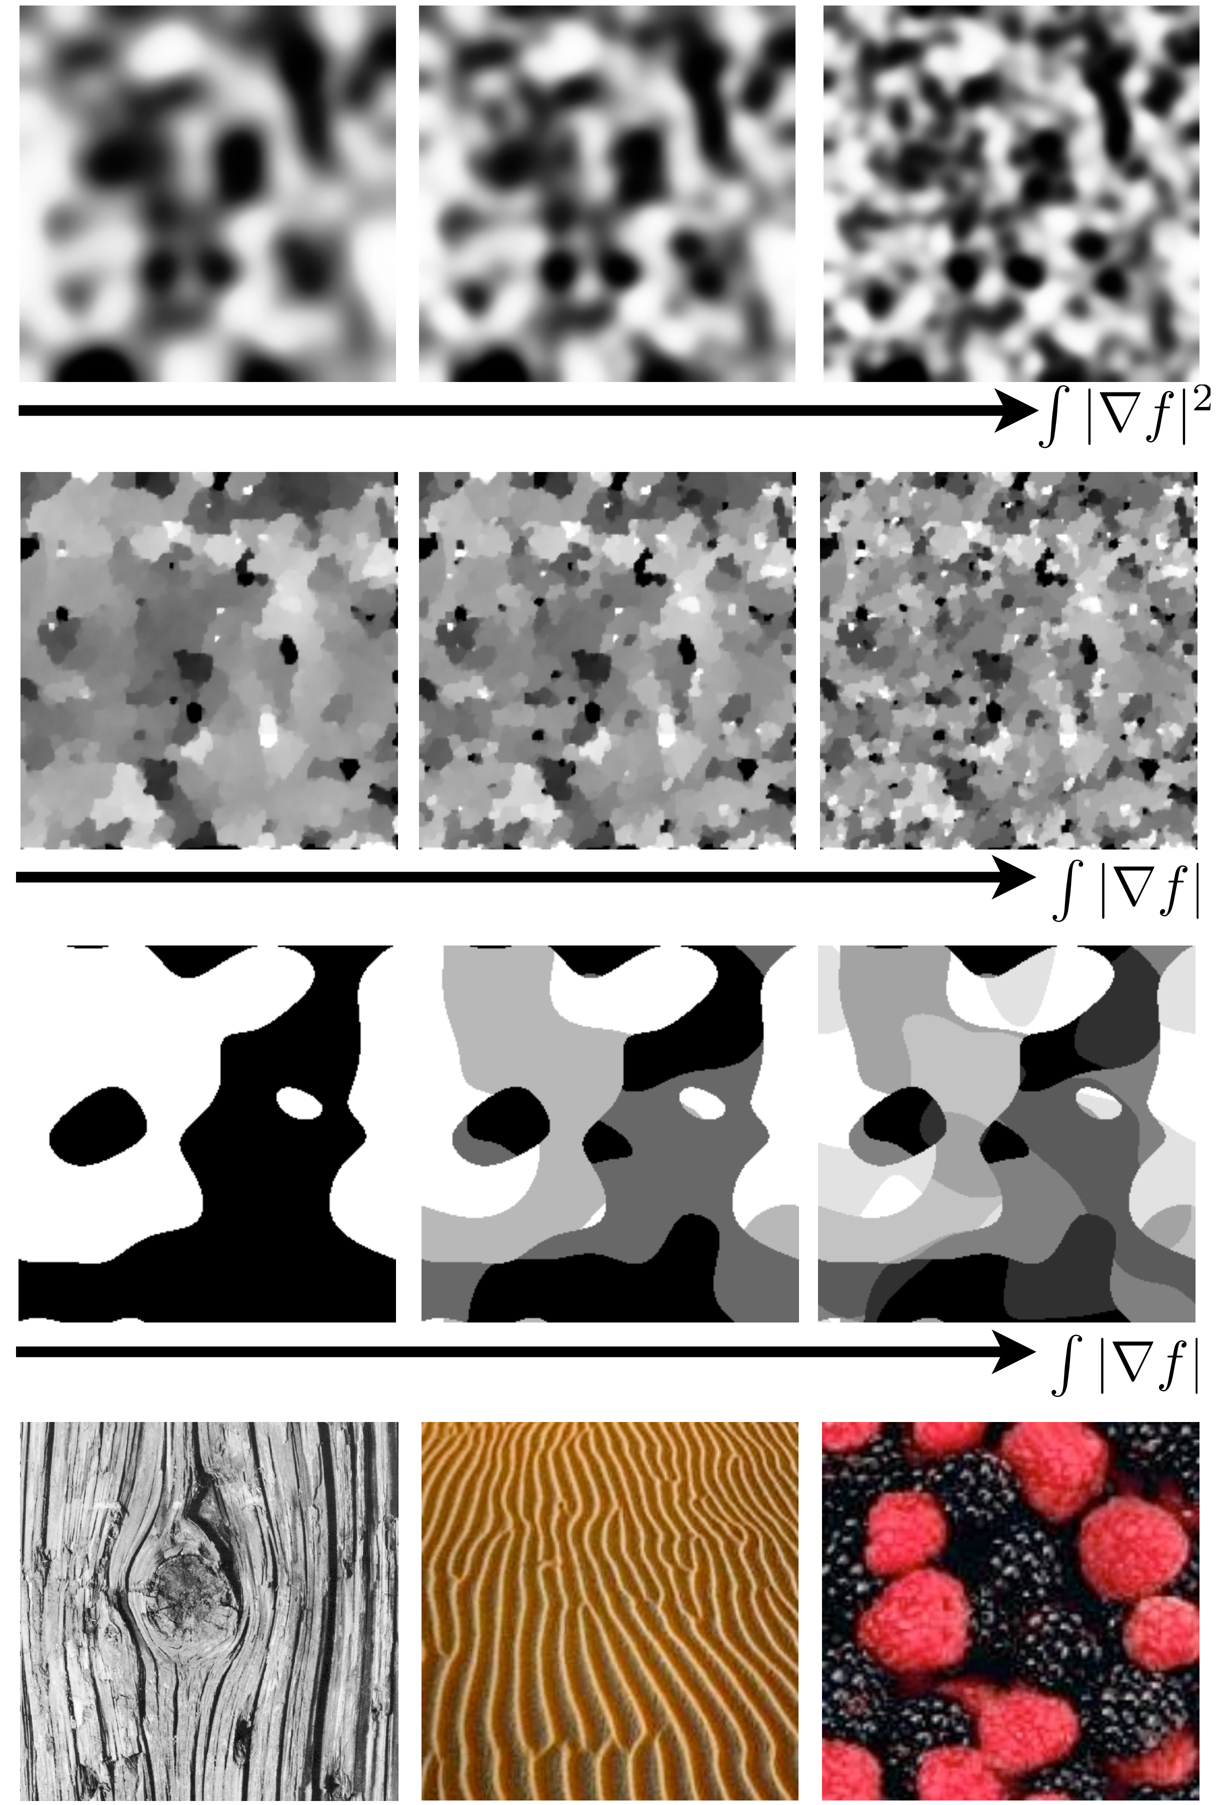

The Local Geometry of Images

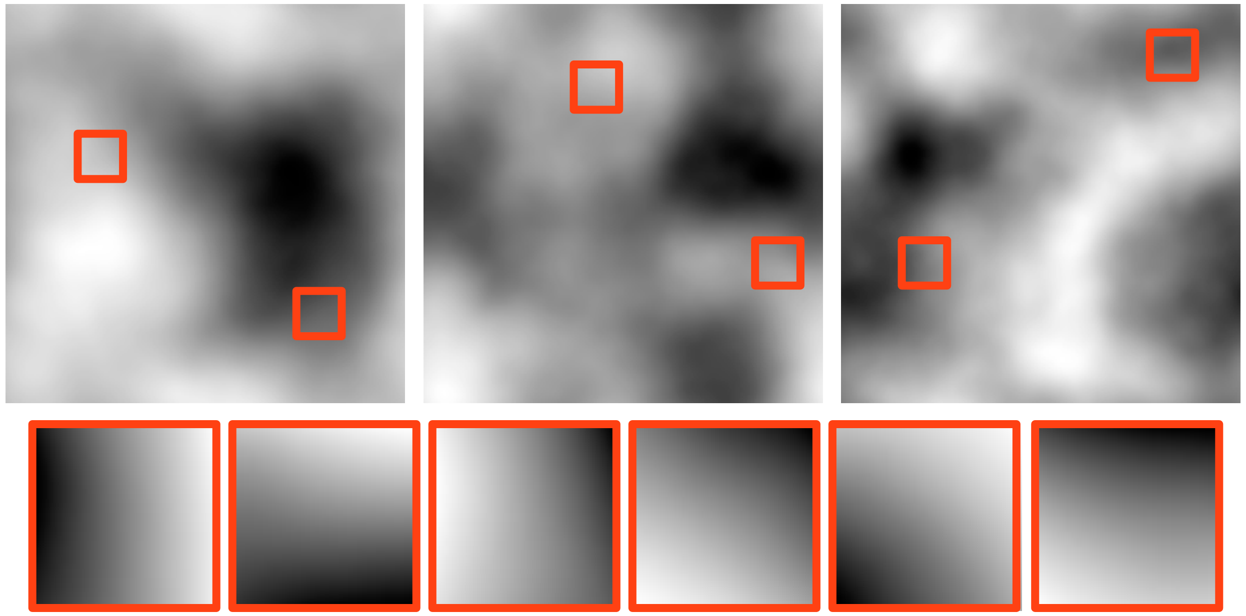

Model: C^2 uniformly regular image.
Patches: linear gradient of intensity.

The Local Geometry of Images

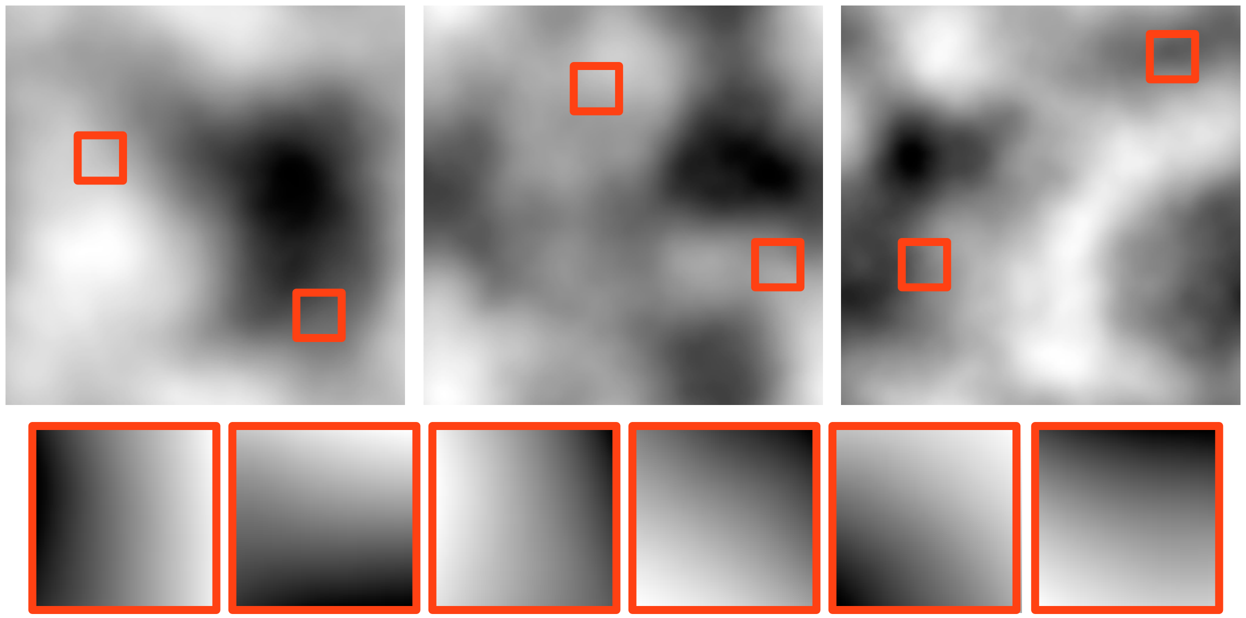

Model: C^2 uniformly regular image.
Patches: linear gradient of intensity.

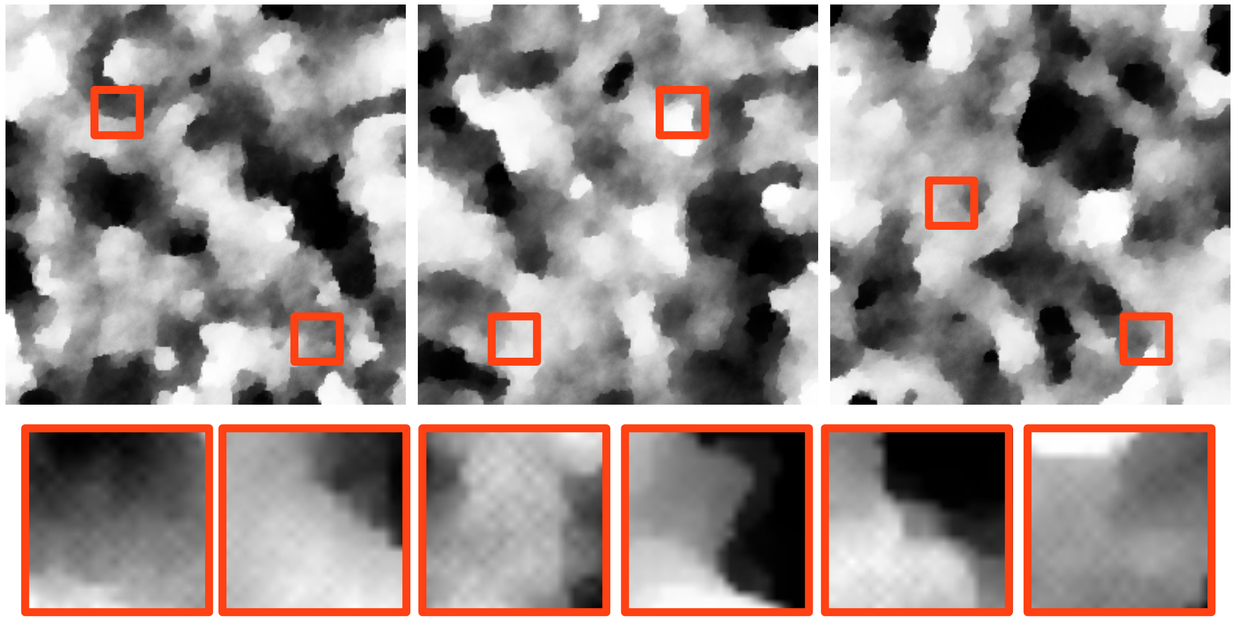

Model: bounded variation image.
Patches: finite length level sets.

The Local Geometry of Images

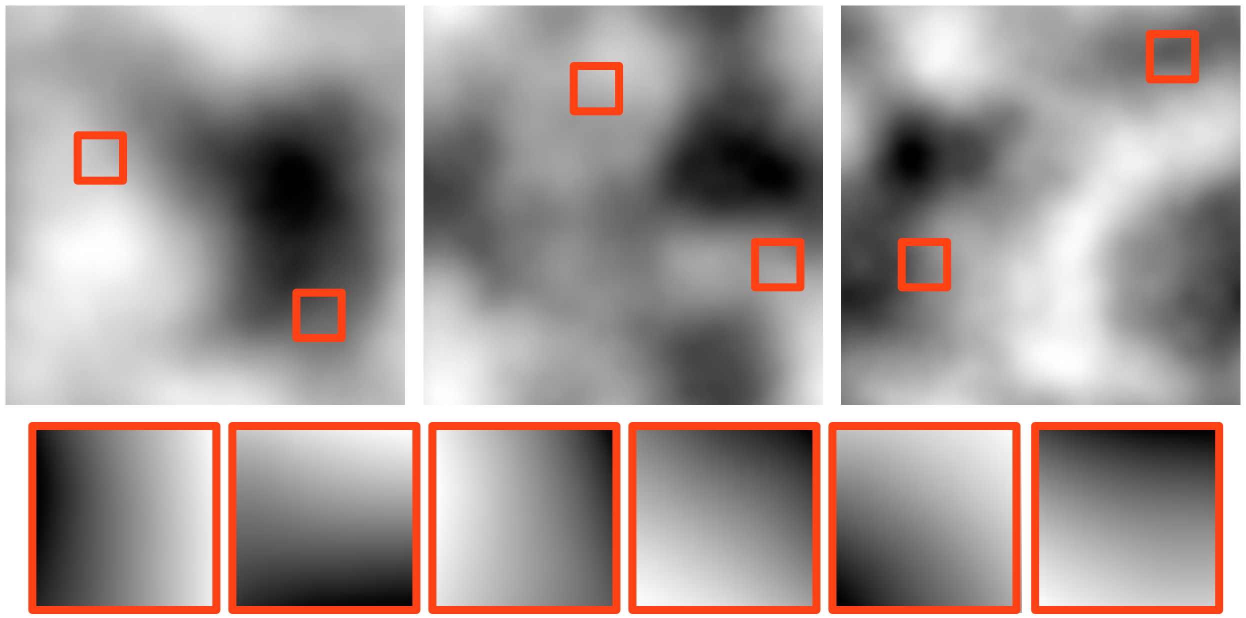

Model: C^2 uniformly regular image.
Patches: linear gradient of intensity.

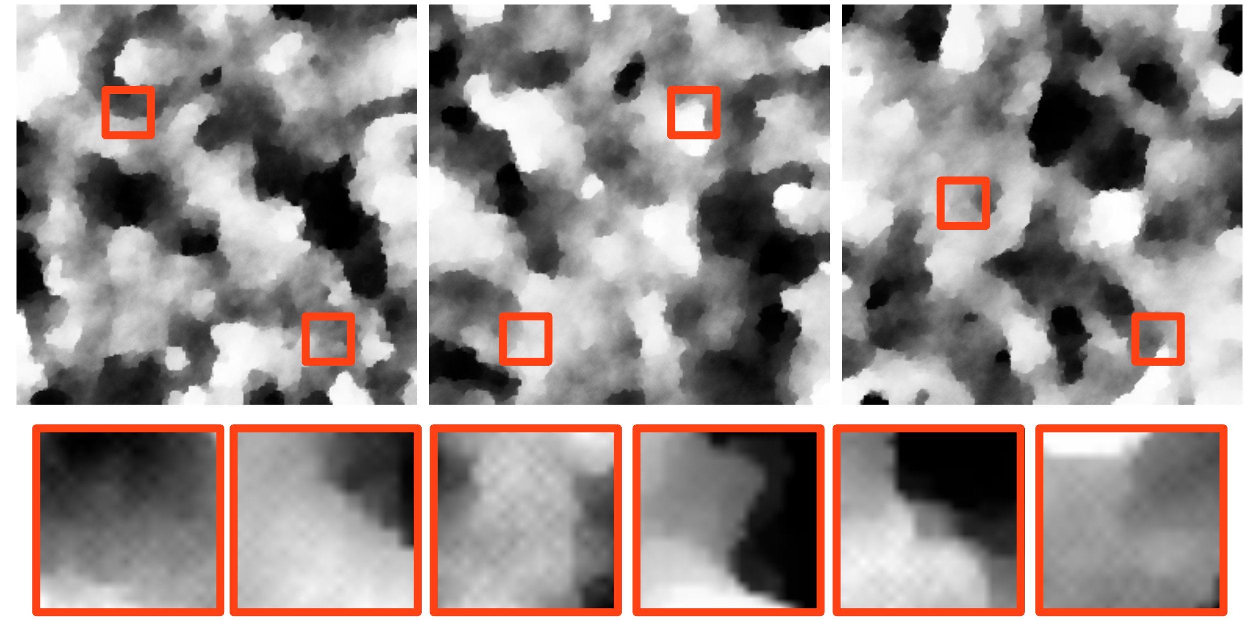

Model: bounded variation image.
Patches: finite length level sets.

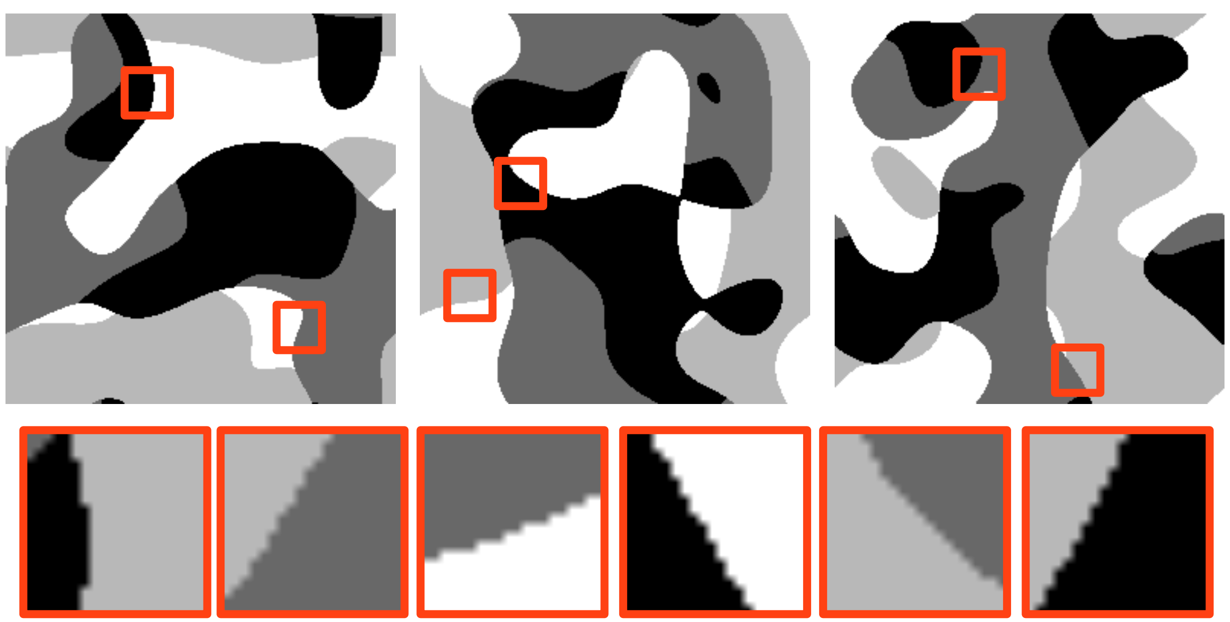

Model: cartoon image.
Patches: binary straight edge.

The Local Geometry of Images

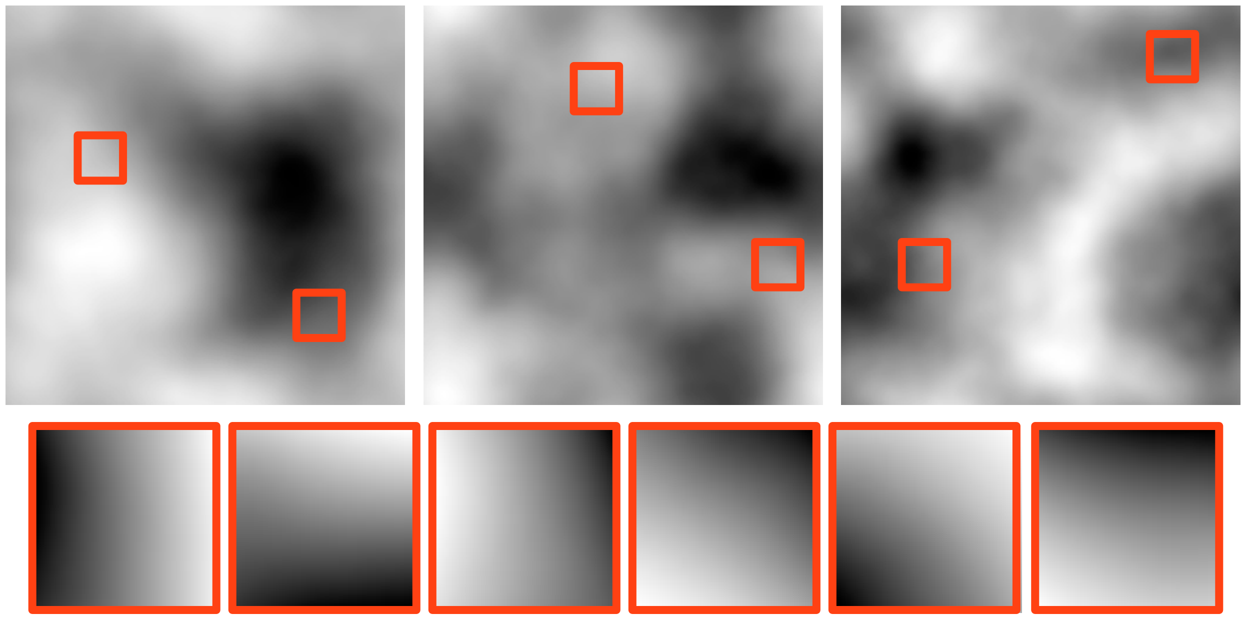

Model: C^2 uniformly regular image.
Patches: linear gradient of intensity.

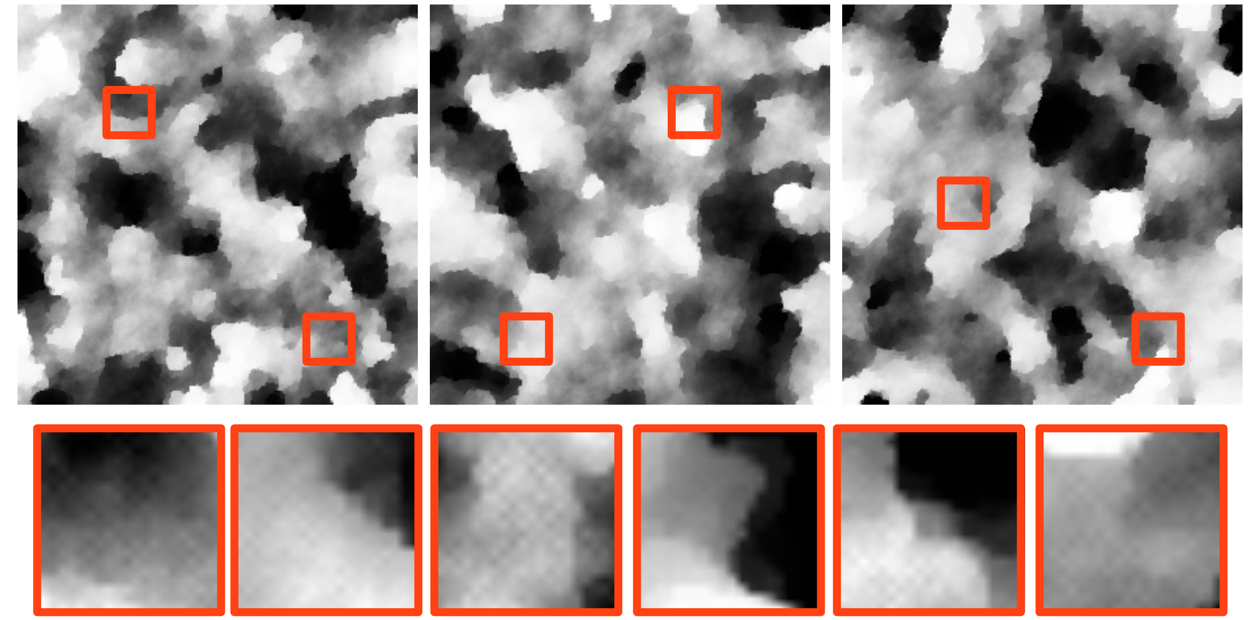

Model: bounded variation image.
Patches: finite length level sets.

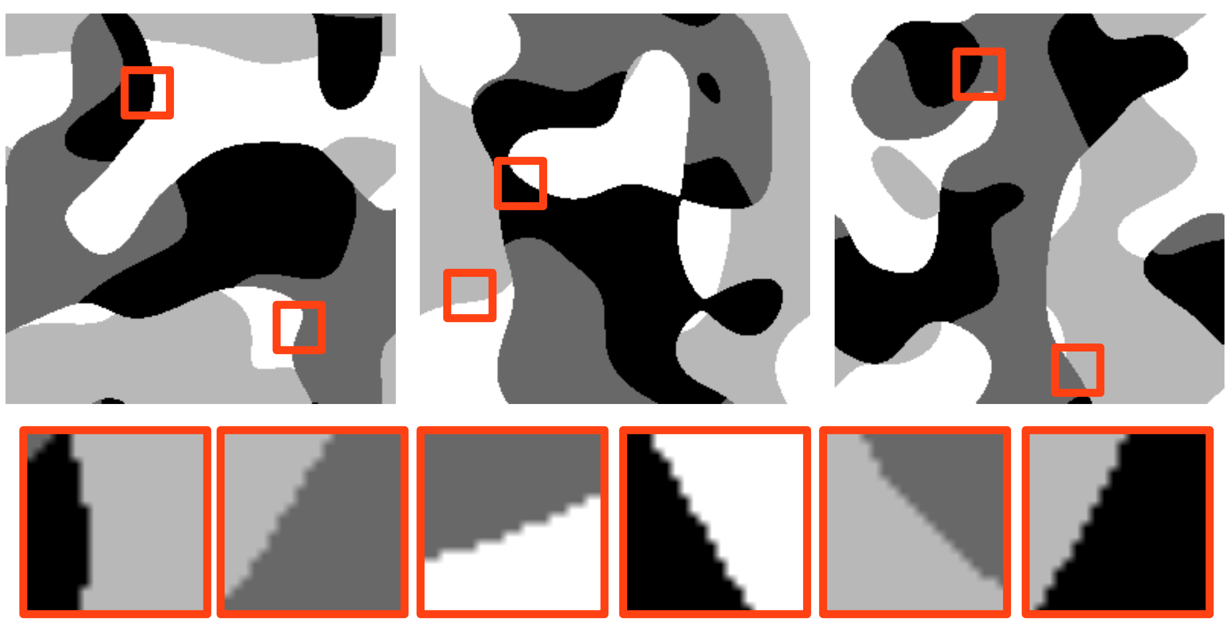

Model: cartoon image.
Patches: binary straight edge.

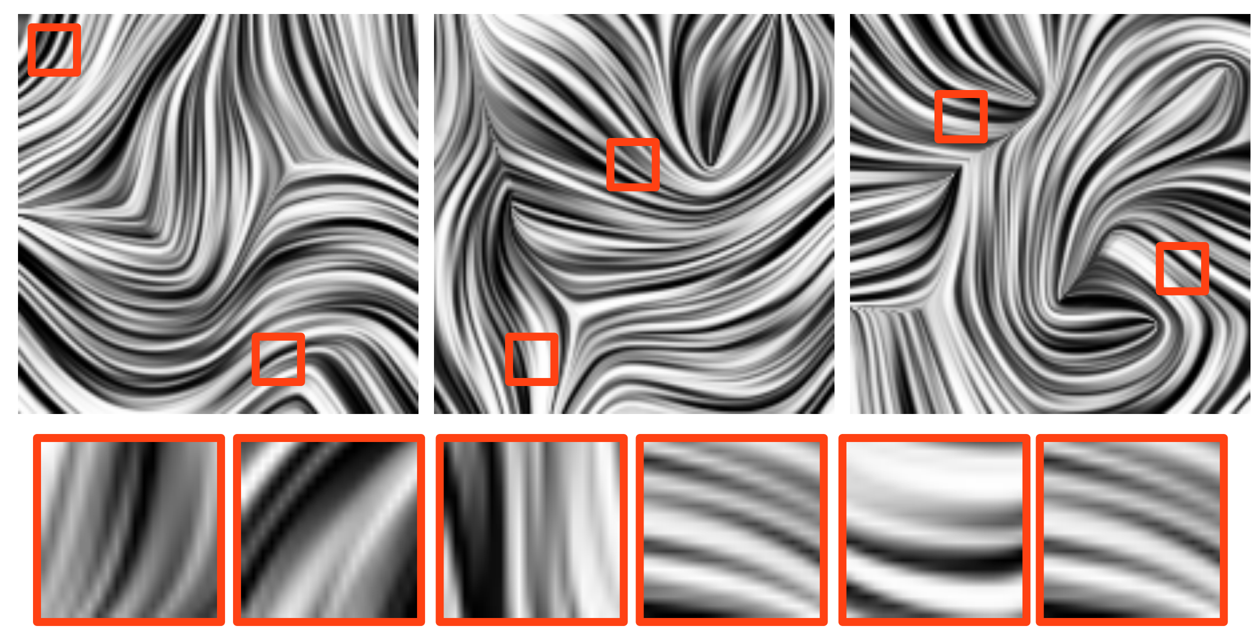

Model: locally parallel texture.
Patches: directional oscillations.

The Local Geometry of Images

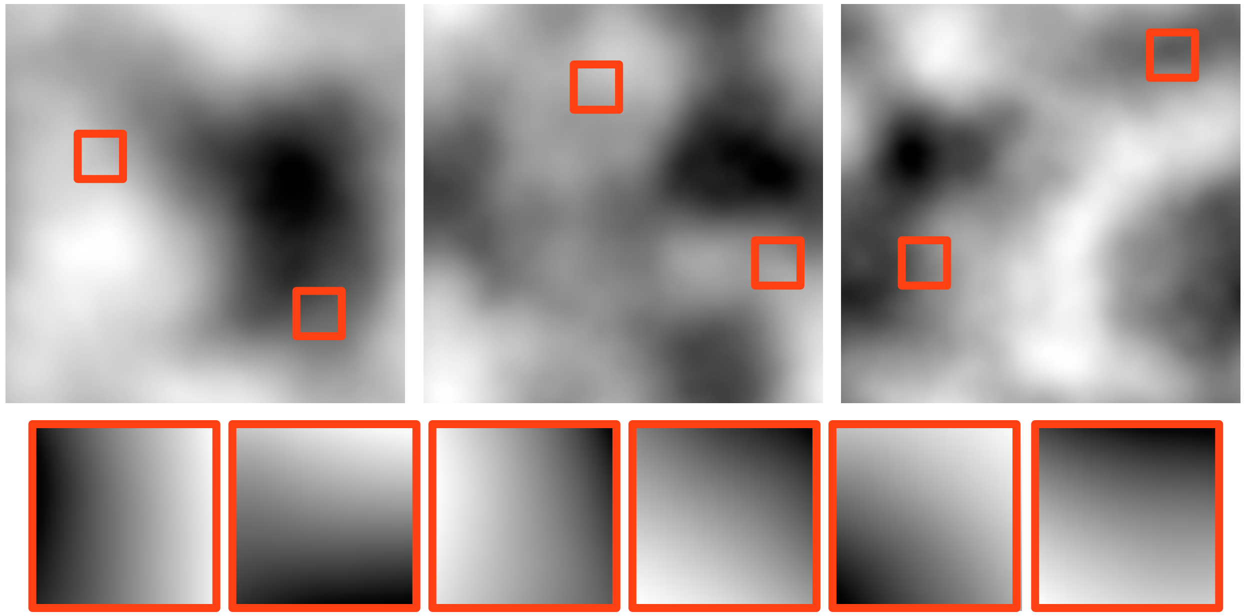

Model: C^2 uniformly regular image.
Patches: linear gradient of intensity.

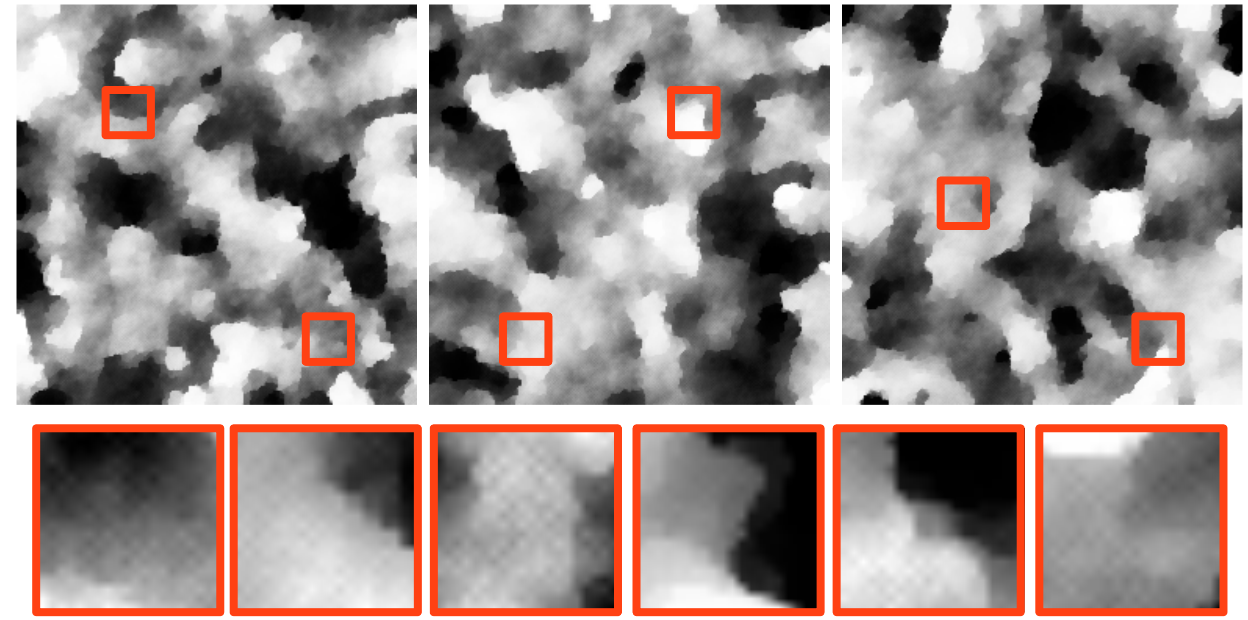

Model: bounded variation image.
Patches: finite length level sets.

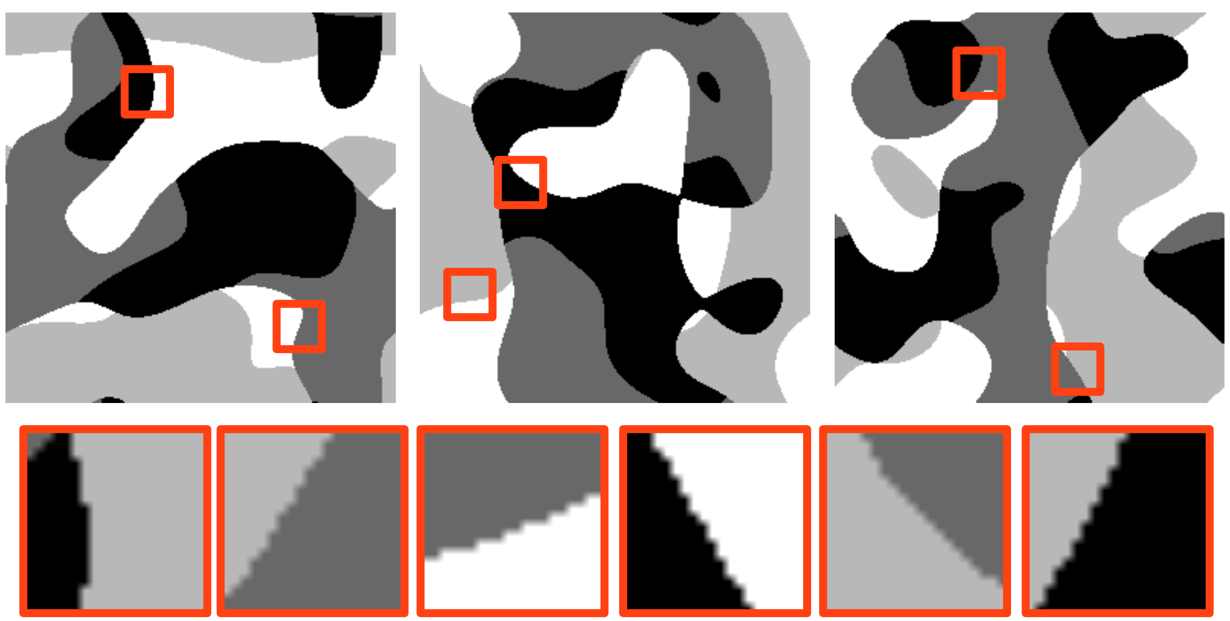

Model: cartoon image.
Patches: binary straight edge.

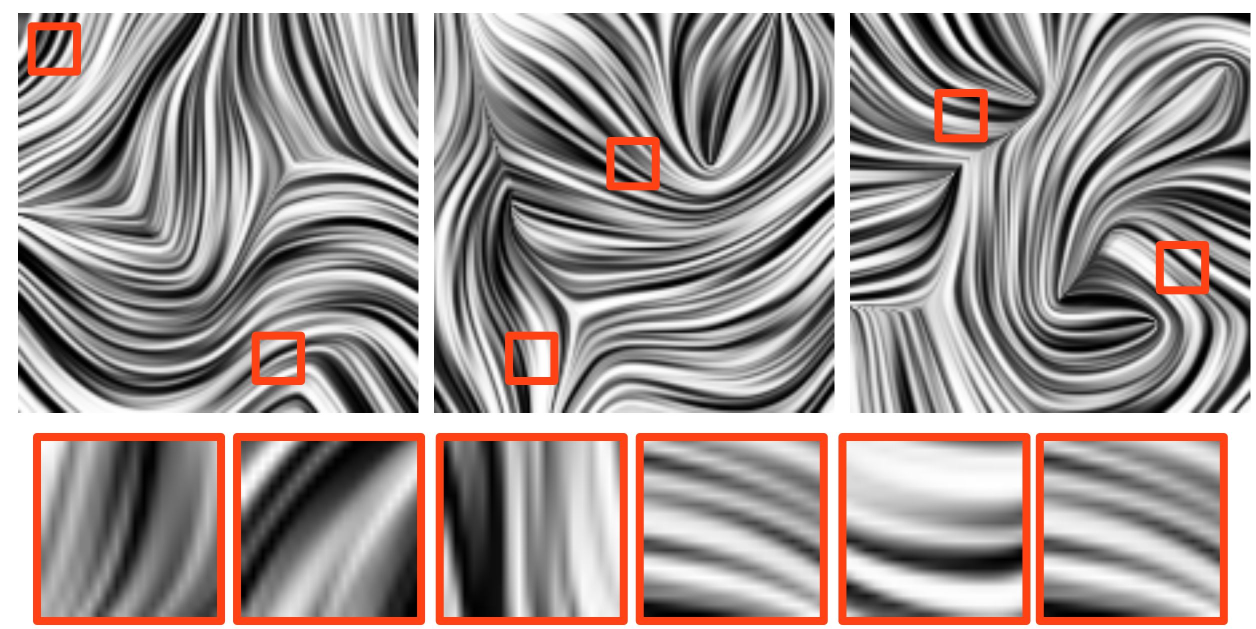

Model: locally parallel texture.
Patches: directional oscillations.

$\longrightarrow$ represent patches with a small number of parameters.

Overview

- Manifolds: Image Libraries vs. Patches
- Examples of Patch Manifolds
- Manifold Energies for Inverse Problems
- Non-adaptive Manifold Models
- Texture Synthesis with Manifold Models
- Adaptive Manifold Model

Manifold of Images Ensembles

Library of images of n pixels: $\{f_k\}_k \subset \mathbb{R}^n$.

Parameterized by a small number $m \ll n$ of parameters

Example: V/H rotation $\theta_v, \theta_h \quad \Longrightarrow \quad f_k(x) = f_0(R_{\theta_h,\theta_v} x)$.

Hypothesis: $\{f_k\} \subset \mathcal{M} \subset \mathbb{R}^n$ smooth manifold of dimension m.

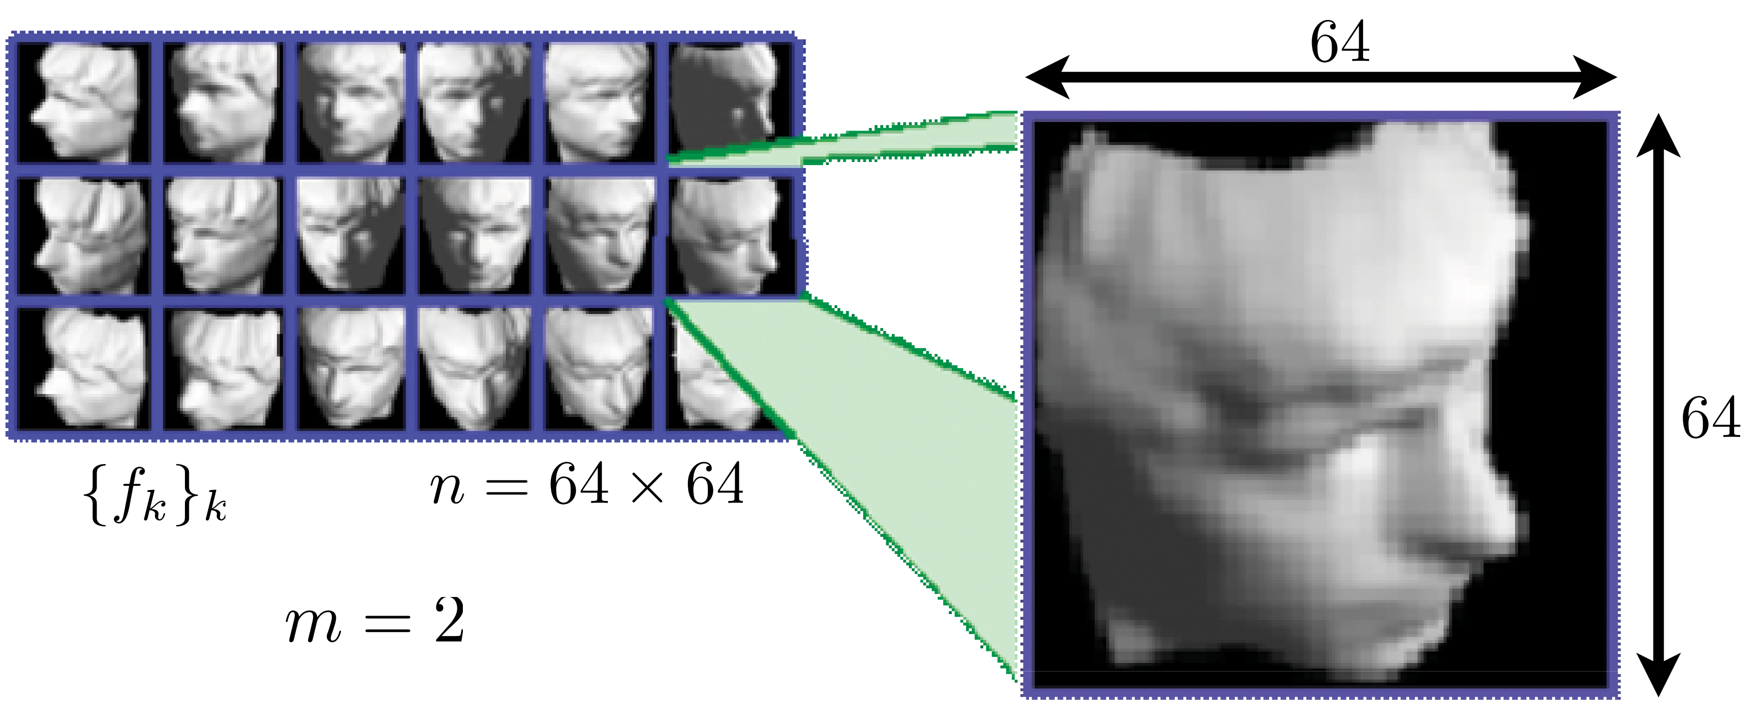

$$\varphi \left\{ \begin{array}{ccc} S^1 \times S^1 & \longrightarrow & \mathcal{M} \\ (\theta_h, \theta_v) & \longmapsto & f_0(R_{\theta_h,\theta_v} x) \end{array} \right.$$

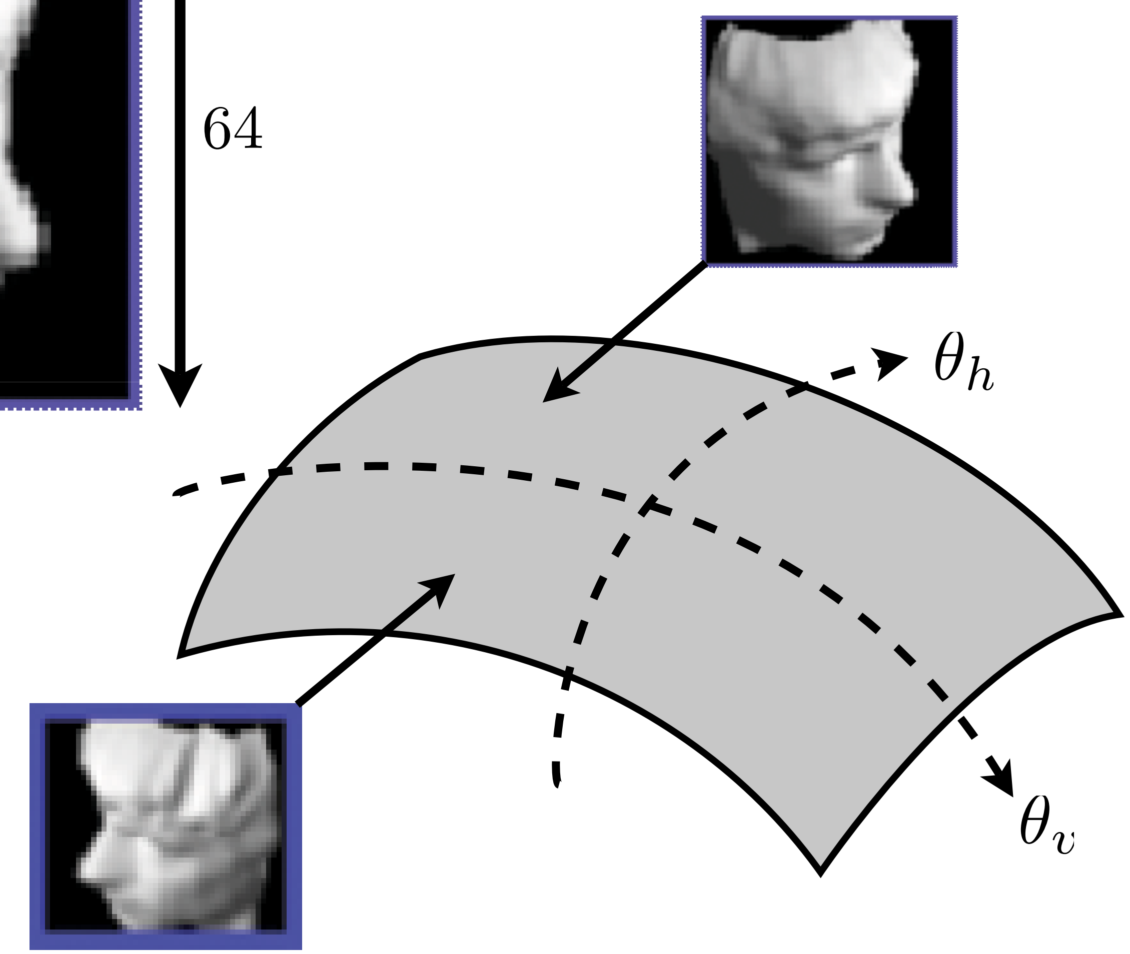

Image Models and Patch Manifolds

Patch extracted from f at location $x \in [0,1]^2$:

$$\forall\, |t| \leqslant \tau/2, \quad p_x(f)(t) = f(x+t)$$

f

$\{p_x(f)\}_x$

Image Models and Patch Manifolds

f $\{p_x(f)\}_x$

Patch extracted from f at location $x \in [0,1]^2$:

$$\forall\, |t| \leqslant \tau/2, \quad p_x(f)(t) = f(x+t)$$

Image model: exploit an image ensemble $\Theta \subset \mathrm{L}^2([0,1]^d)$,

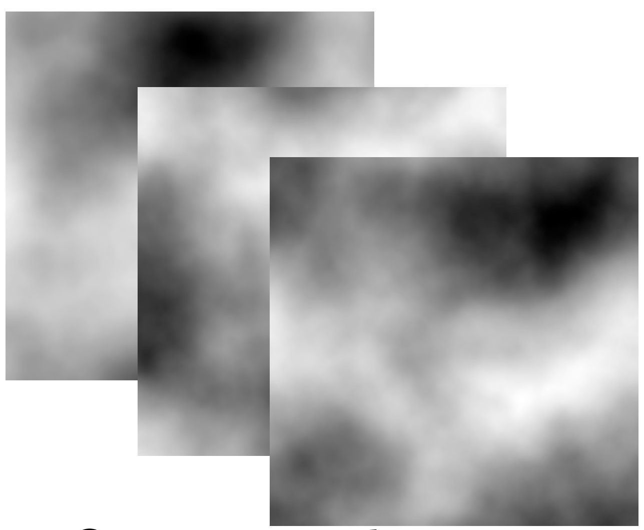

Θ =smooth images

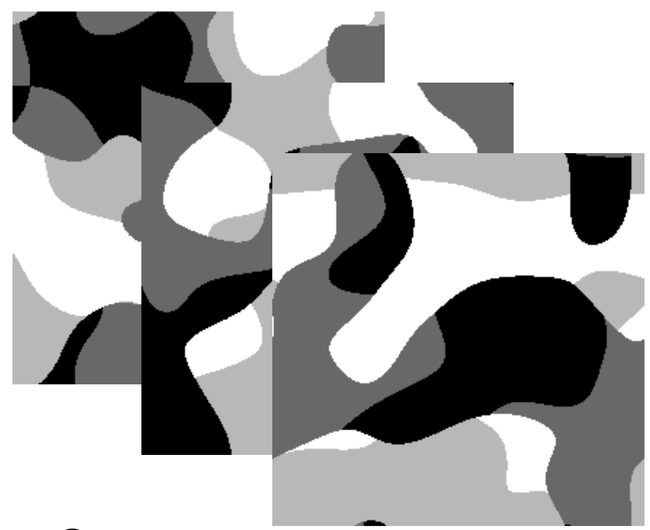

Θ =cartoon images

Θ =oscilating textures

Image Models and Patch Manifolds

Patch extracted from f at location $x \in [0,1]^2$:

$$\forall\, |t| \leqslant \tau/2, \quad p_x(f)(t) = f(x+t)$$

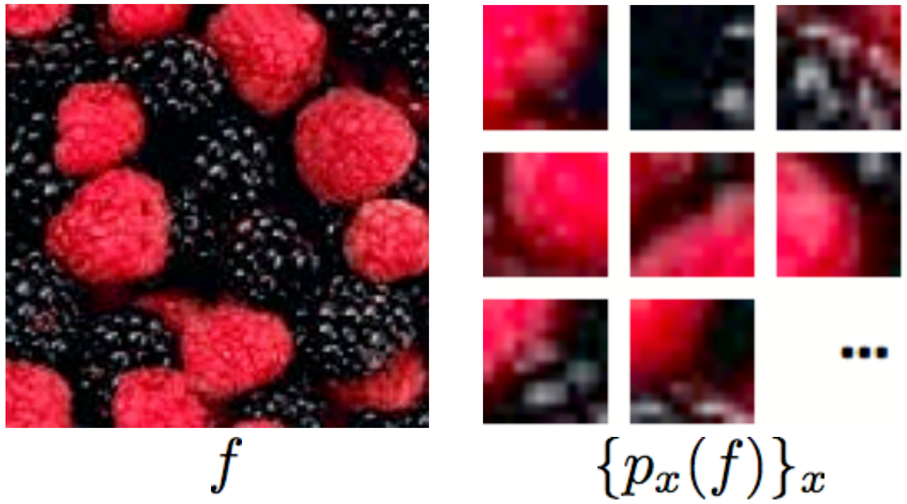

Image model: exploit an image ensemble $\Theta \subset \mathrm{L}^2([0,1]^d)$,

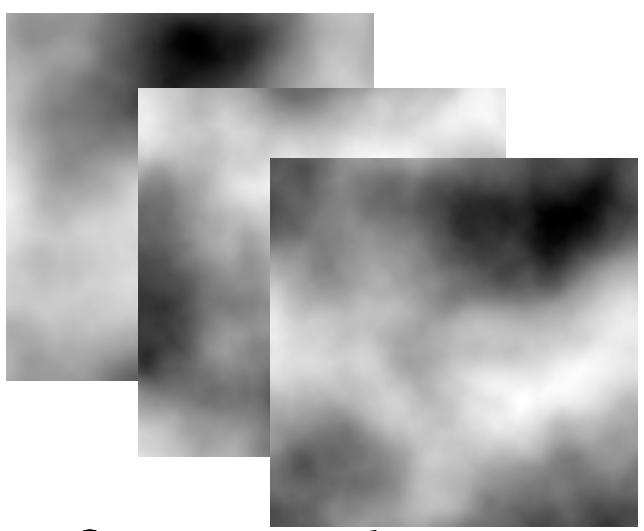

Θ =smooth images

Θ =cartoon images

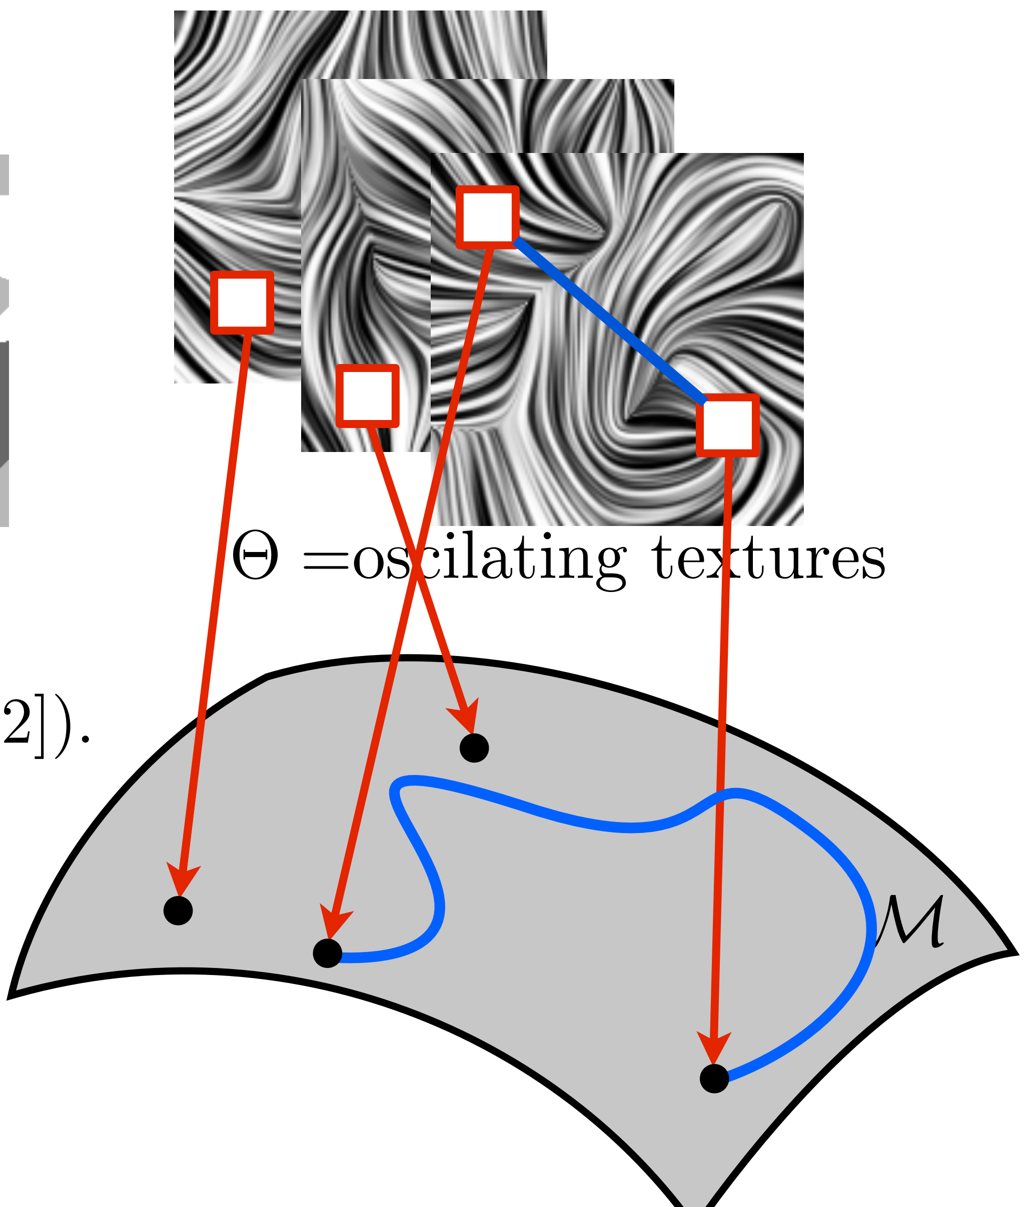

Θ =oscilating textures

$$\mathcal{M} = \{p_x(g) \setminus x \in [0,1]^d \quad \text{and} \quad g \in \Theta\} \subset \mathrm{L}^2([-\tau/2, \tau/2]).$$

What is the topology / geometry of $\mathcal{M}$?

Use it for synthesis of geometrical images.

Non-adaptive setting: $\mathcal{M}$ is fixed.

Image Models and Patch Manifolds

Non-adaptive processing: exploit a signal ensemble $\Theta \subset \mathrm{L}^2([0,1]^d)$,

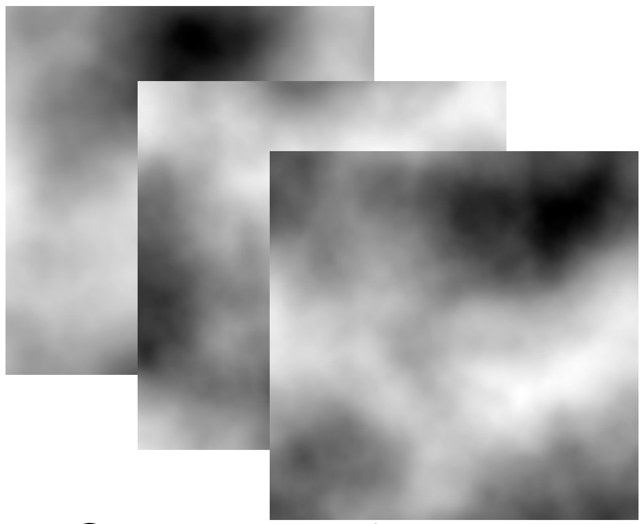

$\Theta =$smooth images

$\Theta =$cartoon images

$\Theta =$oscilating textures

$$\mathcal{M} = \left\{ p_x(g) \setminus x \in [0,1]^d \quad \text{and} \quad g \in \Theta \right\} \subset \mathrm{L}^2([-\tau/2, \tau/2]).$$

Image Models and Patch Manifolds

Non-adaptive processing: exploit a signal ensemble $\Theta \subset \mathrm{L}^2([0,1]^d)$,

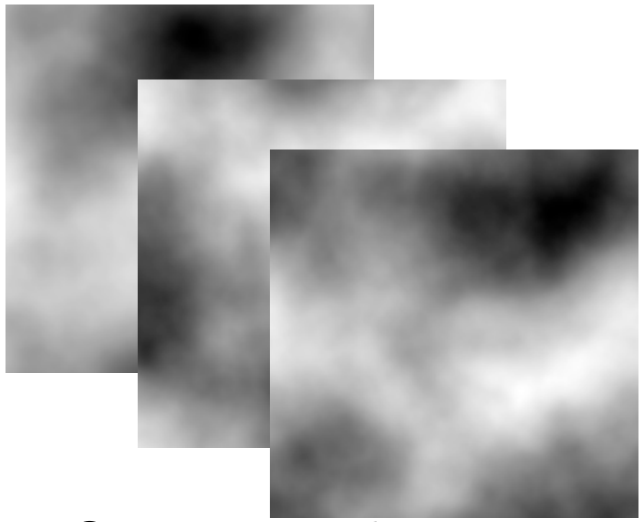

Θ =smooth images

Θ =cartoon images

Θ =oscilating textures

$$\mathcal{M} = \left\{ p_x(g) \setminus x \in [0,1]^d \quad \text{and} \quad g \in \Theta \right\} \subset \mathrm{L}^2([-\tau/2, \tau/2]).$$

Adaptive processing: $\mathcal{M} = \mathcal{M}_f$ is estimated from some $f \in \mathrm{L}^2([0,1]^d)$

Estimating $\mathcal{M}_f \iff$ estimating connexions between the points $\{p_x(f)\}_x$.

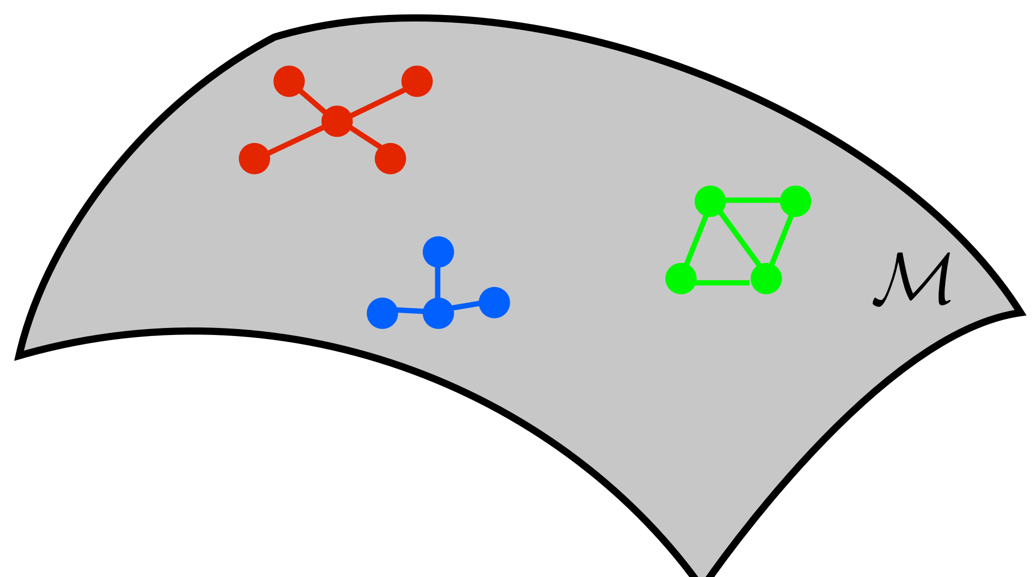

$\longrightarrow$ use $\mathcal{M}$ or $\mathcal{M}_f$ to regularize image processing problems.

Overview

- Manifolds: Image Libraries vs. Patches
- Examples of Patch Manifolds
- Manifold Energies for Inverse Problems
- Non-adaptive Manifold Models
- Texture Synthesis with Manifold Models
- Adaptive Manifold Model

Manifold of Smooth Images

$$\Theta = \{ f \in \mathrm{C}^2 \setminus \|f\|_\infty \leqslant C_1, \|\nabla f\|_\infty \leqslant C_2 \}$$

Patch $\approx$ linear gradient of intensity.

$$p_x(f)(t) \approx a(x) + \langle b(x),\, t \rangle$$

where $\quad a(x) = f(x) \quad$ and $\quad b(x) = \nabla_x f$

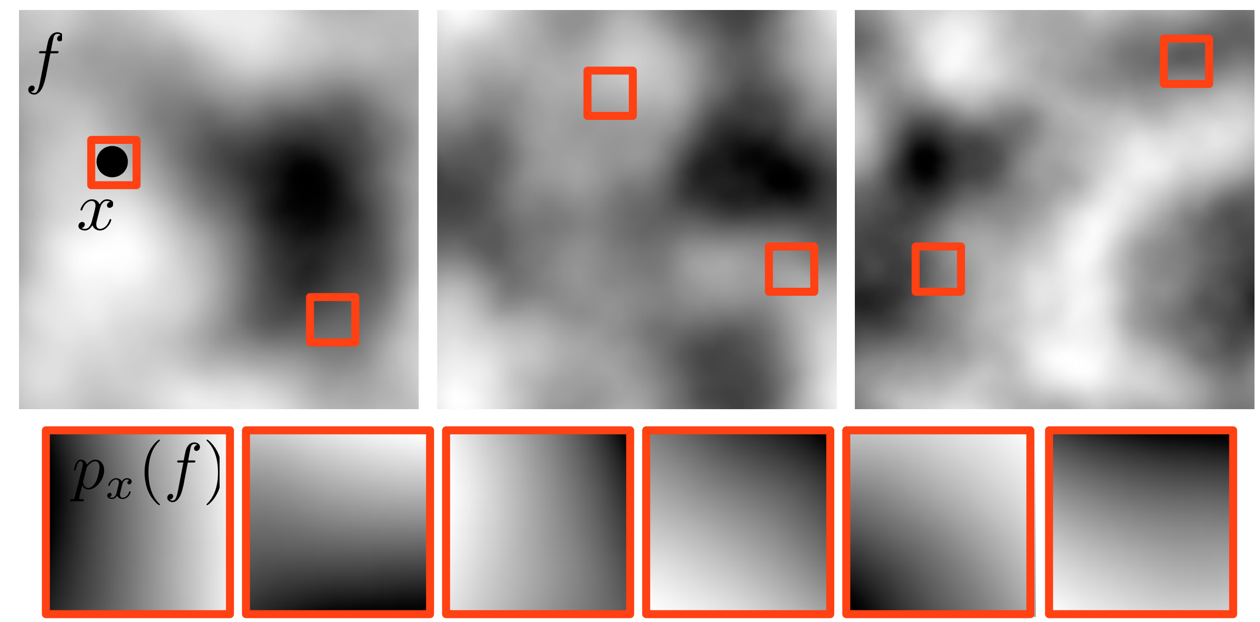

Manifold of Smooth Images

$$\Theta = \{ f \in \mathrm{C}^2 \setminus \|f\|_\infty \leqslant C_1, \|\nabla f\|_\infty \leqslant C_2 \}$$

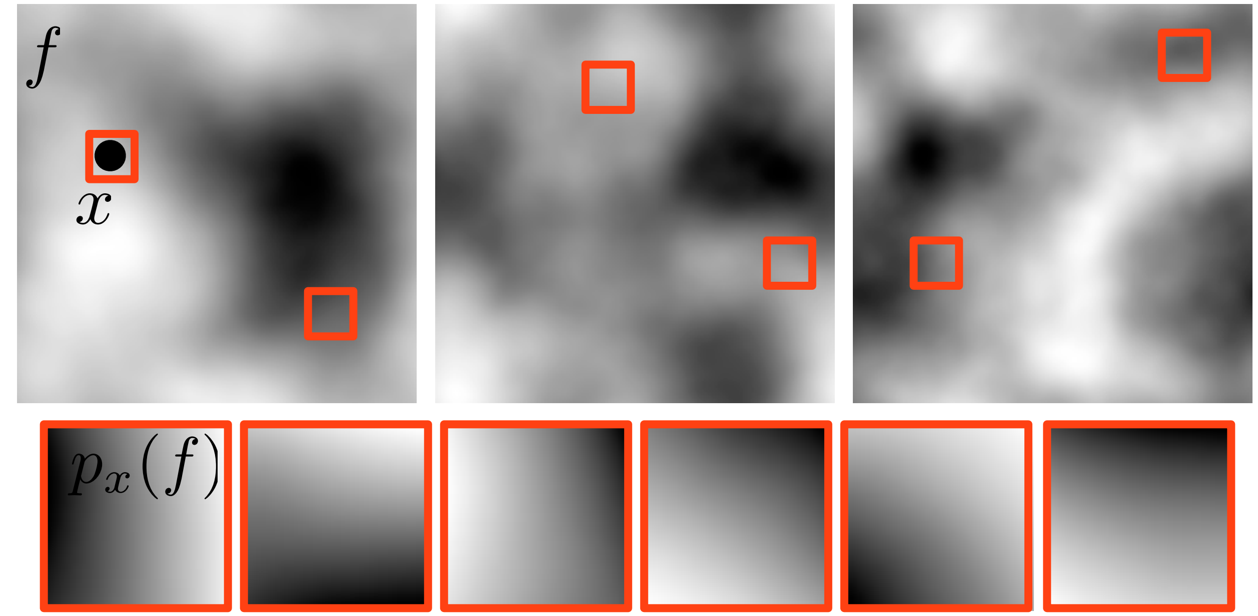

Patch $\approx$ linear gradient of intensity.

$$p_x(f)(t) \approx a(x) + \langle b(x),\, t \rangle$$

where $\quad a(x) = f(x) \quad$ and $\quad b(x) = \nabla_x f$

Manifold of affine patches: $\quad \mathcal{M} = \{ t \mapsto a + \langle b,\, t \rangle \setminus |a| \leqslant C_1, |b| \leqslant C_2 \}$

$\mathcal{M} \simeq [-C_1, C_1] \times [-C_2, C_2] \times [-C_2, C_2]$ "3D cube"

$\mathcal{M}$ is a flat (Euclidean) manifold.

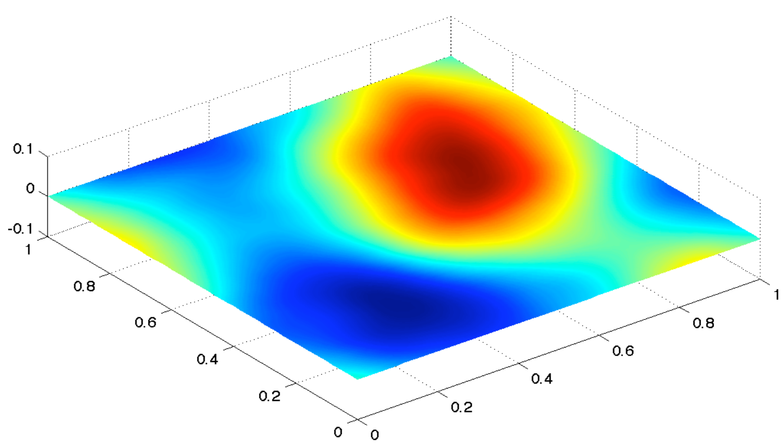

$x \mapsto f(x)$

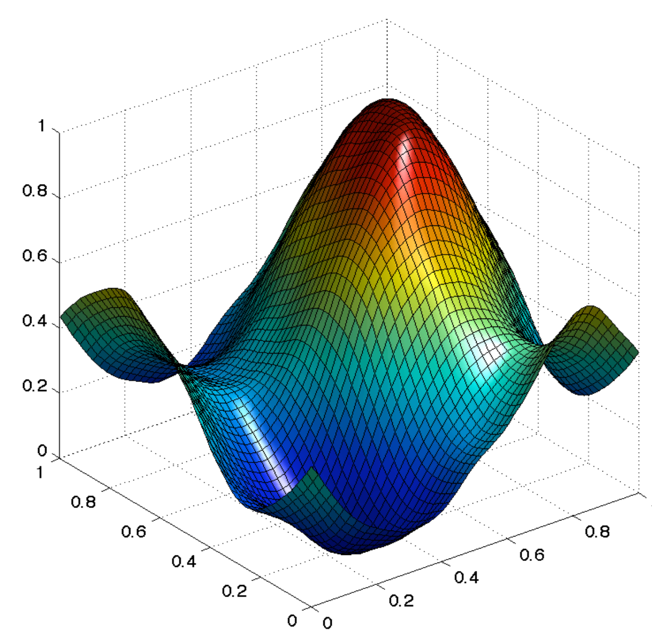

$x \mapsto (x, f(x))$

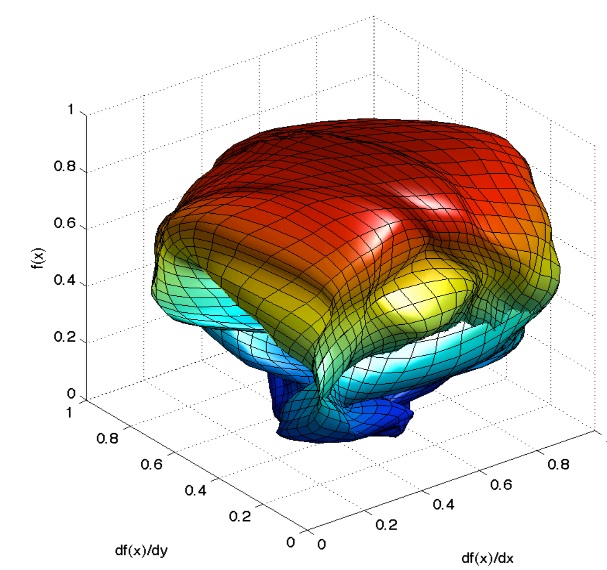

$x \mapsto p_x(f)$

Manifold of Cartoon Images

$$\Theta_{\text{cartoon}} = \{f \setminus f \text{ is } \mathrm{C}^\alpha \text{ outside } \mathrm{C}^\alpha \text{ curves}\}\,.$$

$$\Theta = \{f = 1_\Omega \setminus \partial\Omega \text{ a } \mathrm{C}^\alpha \text{ curve }\}\,.$$

$$p_x(f)(t) = P_{\theta(x),\delta(x)}(t)$$

$$\text{where} \quad \begin{cases} P_{\theta,\delta}(t) = P_{0,0}(R_\theta(t-\delta)) \\ P_{0,0}(x) = 1_{x_1 \geqslant 0}(x) \end{cases}$$

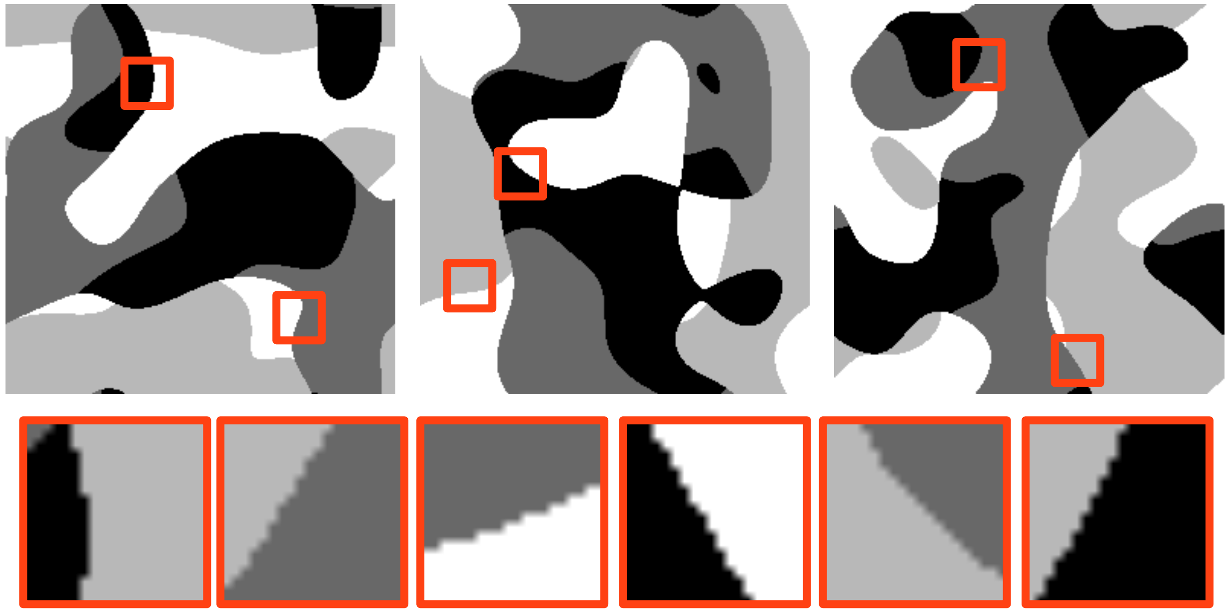

Manifold of Cartoon Images

$$\Theta_{\text{cartoon}} = \{ f \setminus f \text{ is } \mathrm{C}^\alpha \text{ outside } \mathrm{C}^\alpha \text{ curves} \}.$$

$$\Theta = \{ f = 1_\Omega \setminus \partial\Omega \text{ a } \mathrm{C}^\alpha \text{ curve } \}.$$

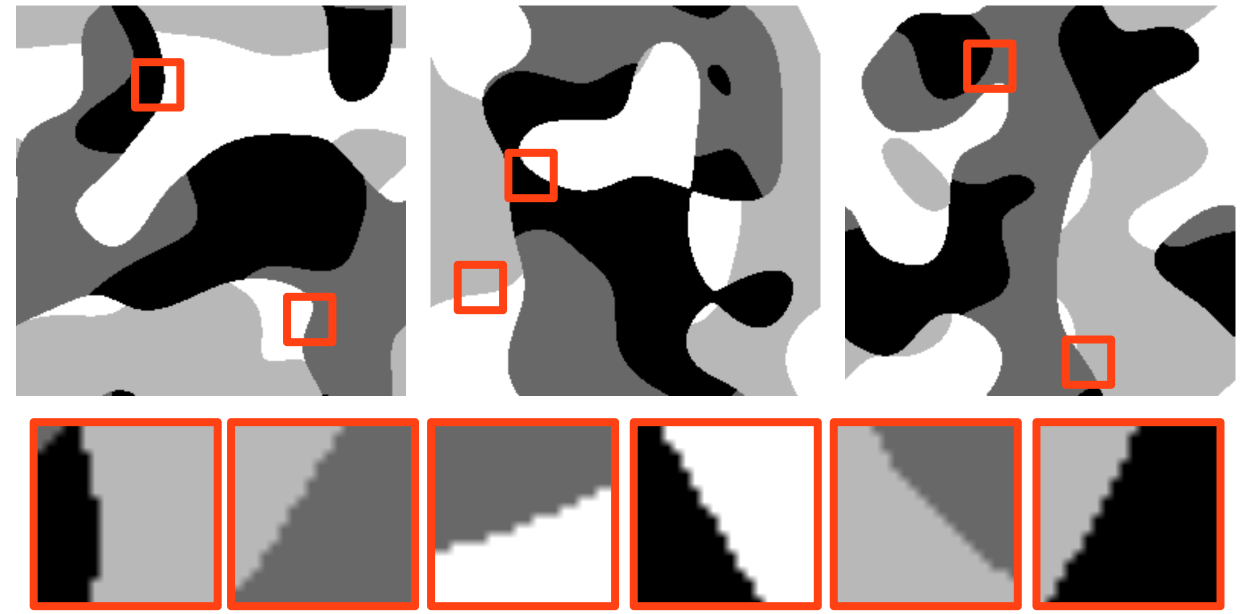

$$p_x(f)(t) = P_{\theta(x),\delta(x)}(t)$$

$$\text{where} \quad \begin{cases} P_{\theta,\delta}(t) = P_{0,0}(R_\theta(t-\delta)) \\ P_{0,0}(x) = 1_{x_1 \geqslant 0}(x) \end{cases}$$

Manifold of binary edges:

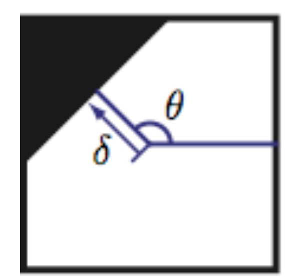

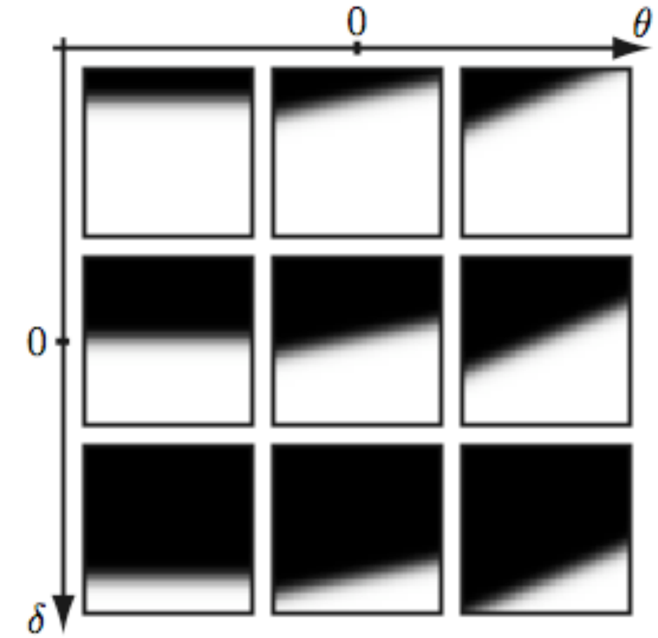

$$\mathcal{M} = \{ P_{\theta,\delta} \setminus \theta \in [0, 2\pi), \delta \in \mathbb{R} \}$$

$$\mathcal{M} \simeq S^1 \times \mathbb{R} \qquad \text{(cylinder)}$$

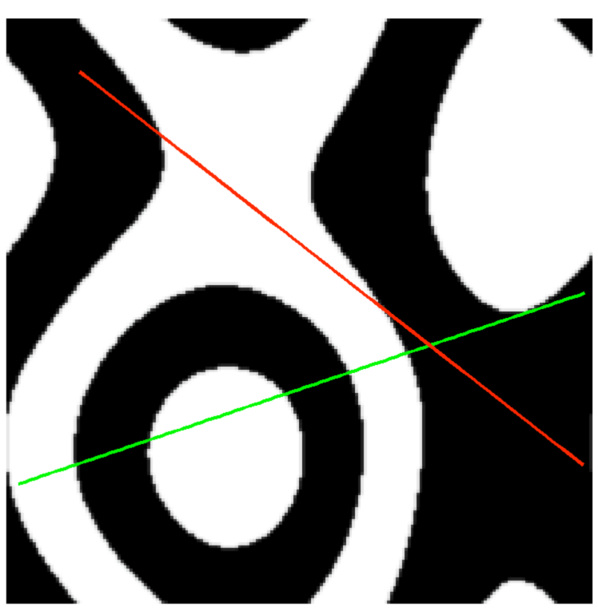

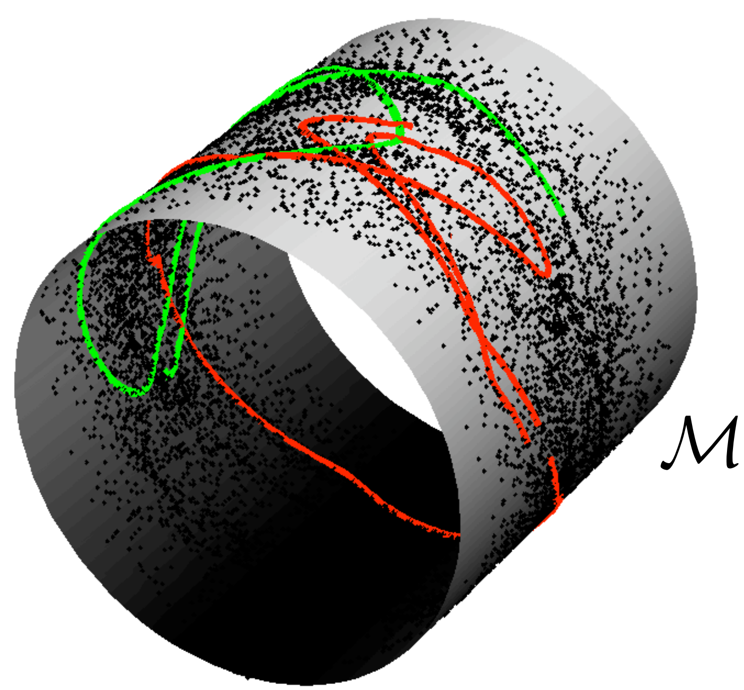

Manifold of Locally Stationary Sounds

$$\Theta \overset{\text{def.}}{=} \{x \mapsto f(x) = A(x)\cos(\Psi(x)) \setminus \|A'\|_\infty \leqslant A_{\max} \quad \text{and} \quad \|\Psi''\|_\infty \leqslant \Psi_{\max}.\}$$

$$\mathcal{M} = \left\{ P_{(A,\rho,\delta)} \setminus A \geqslant 0 \quad \text{and} \quad \rho \geqslant 0 \quad \text{and} \quad \delta \in \mathrm{S}^1 \right\}$$

where $\quad P_{(A,\rho,\delta)}(x) \overset{\text{def.}}{=} A\cos(\rho x + \delta).$

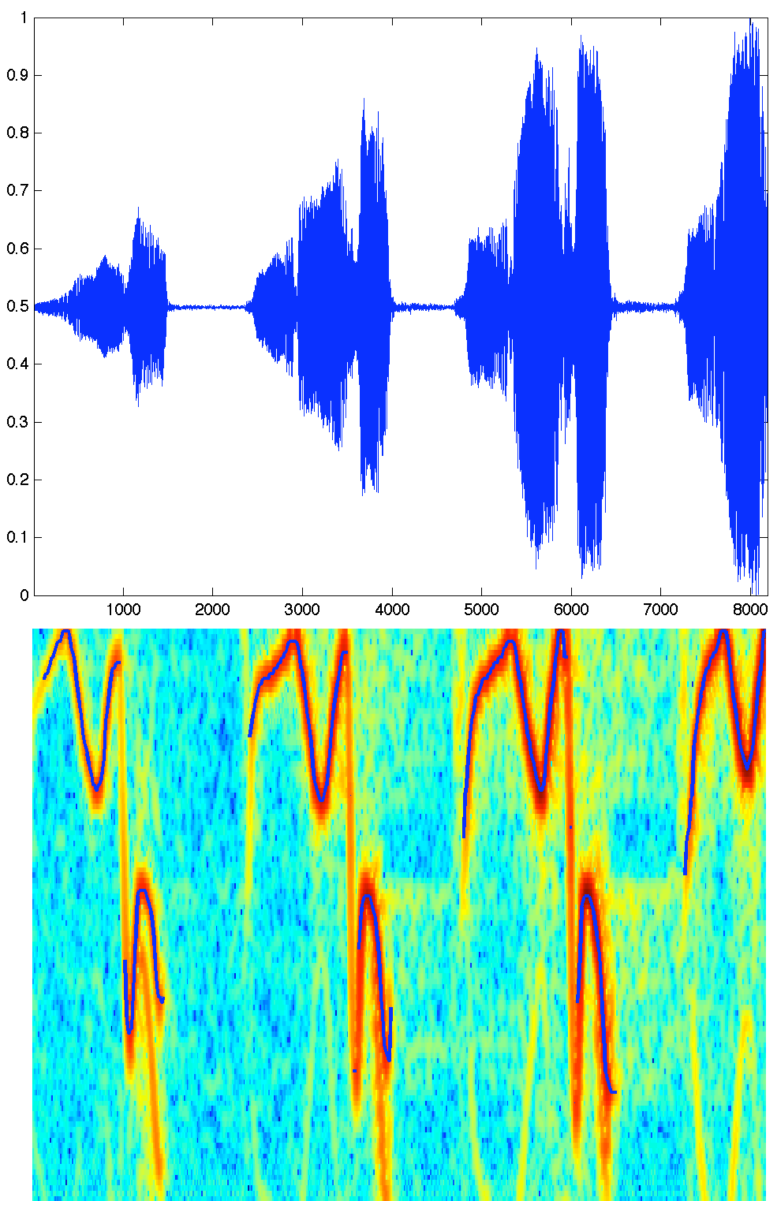

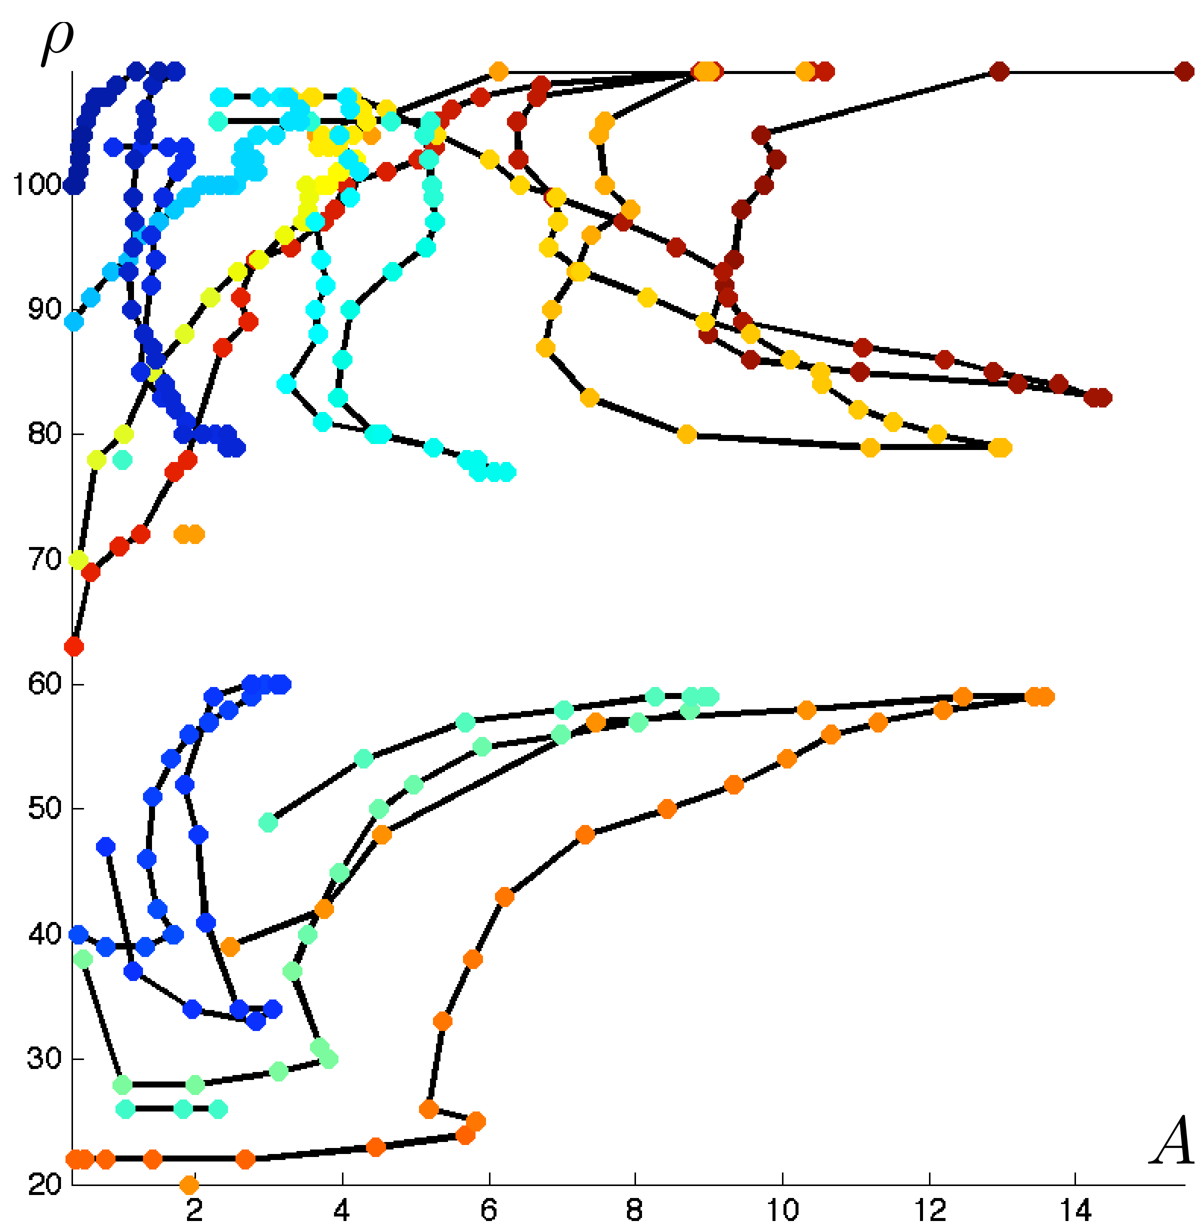

Manifold of Locally Parallel Textures

$$f(x) = A(x)\cos(\Phi(x))$$

Phase Φ slowly varying.

Orientation: $\nabla_x \Phi$

Amplitude: $A(x)$

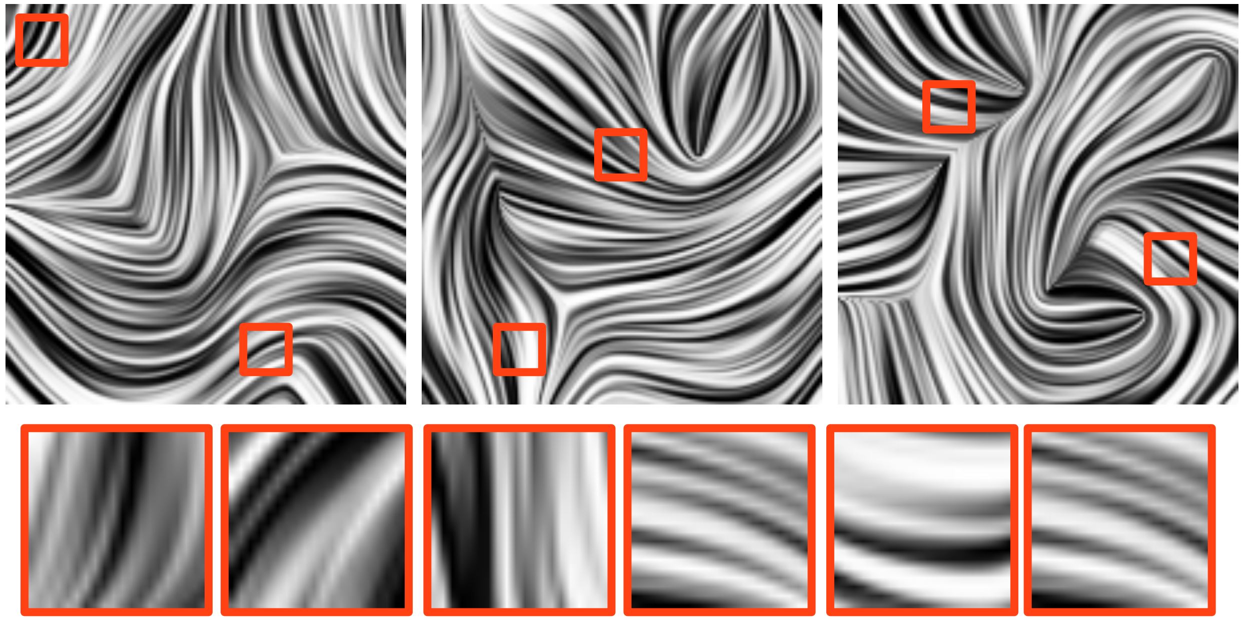

$$\mathcal{M} = \left\{ A P_{\rho,\theta,\delta} \setminus A \leqslant C_1,\ \rho \leqslant C_2,\ \theta \in \tilde{S}^1,\ \delta \in S^1 \right\}$$

Manifold of Locally Parallel Textures

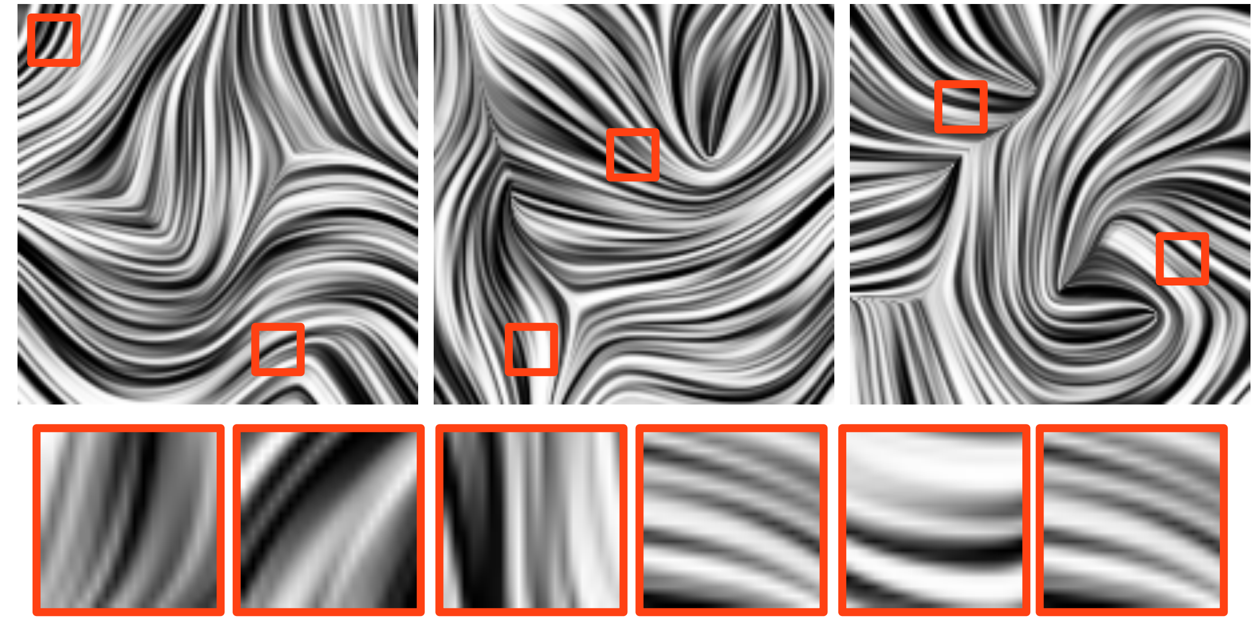

$$f(x) = A(x)\cos(\Phi(x))$$

Phase Φ slowly varying.

Orientation: $\nabla_x \Phi$

Amplitude: $A(x)$

$$p_x(f) \approx A(x) P_{\rho(x),\theta(x),\delta(x)} \quad \text{where} \quad P_{\rho,\theta,\delta}(t) = \cos(\rho\langle t,\, \theta\rangle + \delta)$$

$$\mathcal{M} = \left\{ A P_{\rho,\theta,\delta} \setminus A \leqslant C_1,\ \rho \leqslant C_2,\ \theta \in \tilde{S}^1,\ \delta \in S^1 \right\}$$

$$\mathcal{M} \simeq [0, C_1] \times [0, C_2] \times \tilde{S}^1 \times S^1 \qquad \theta \in \tilde{S}^1 \quad \text{(orientation but no direction)}$$

$(A(x), \rho(x), \theta(x), \delta(x))$ can be estimated with a local Fourier transform.

$f(x)$

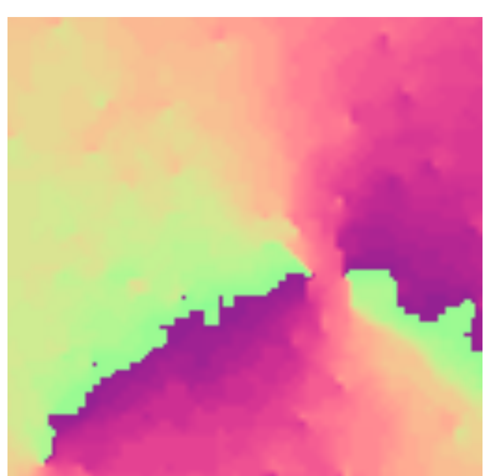

$\theta(x)$

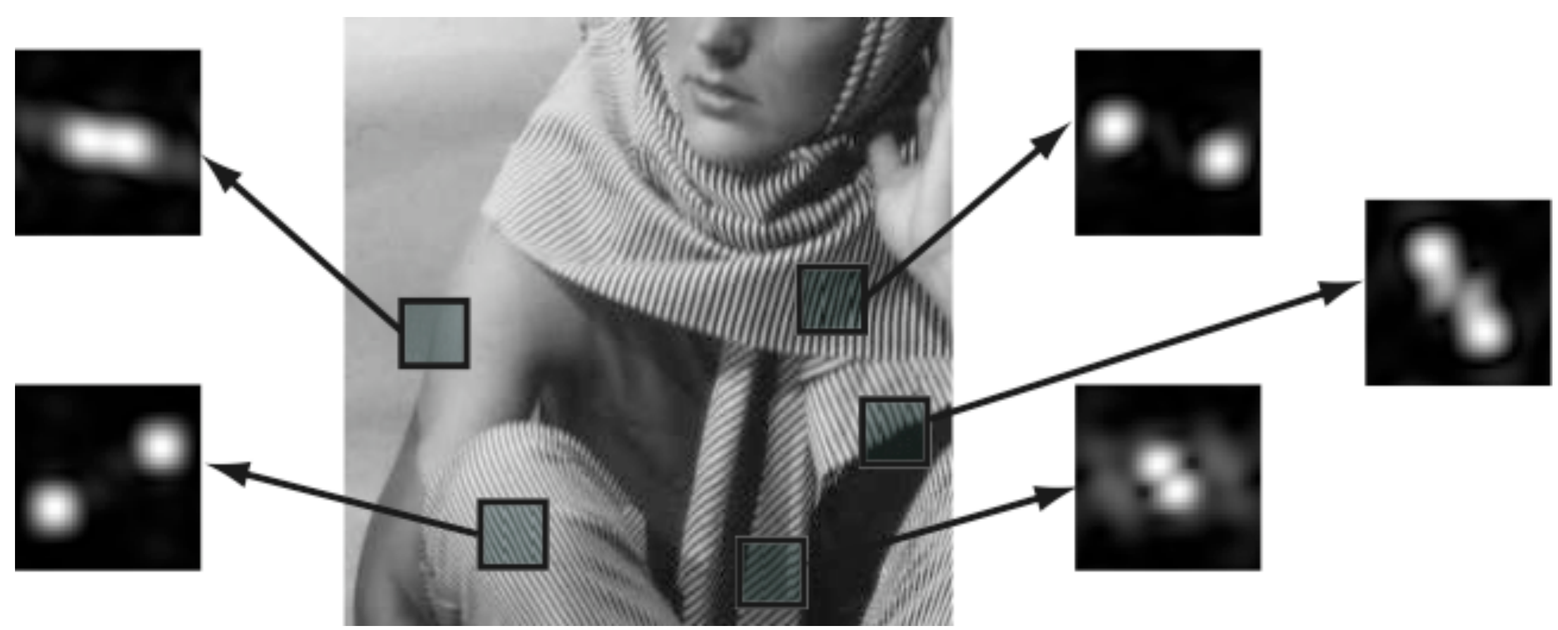

Sparse Texture Ensemble

Dictionary: collection of atoms: $\Psi = (\psi_i)_{i=0}^{q-1}$.

Redundancy: $\psi_i \in \mathbb{R}^n$ and $q \gg n$.

Examples: translation invariant wavelets, Gabor frame, etc.

Linear expansion of a patch: $p_x(f) = \sum_i s_x(i)\psi_i = \Psi s_x$.

Sparsity: only a few $s_x(i)$ are non-zero.

$$\|s_x\|_{\ell^0} = \#\{i \setminus s_x(i) \neq 0\} \leqslant k$$

$$\Theta = \{f \setminus \forall\, x,\ p_x(f) = \Psi s_x \quad \text{with} \quad \|s_x\|_{\ell^0} \leqslant k\}$$

$$\mathcal{M} = \{\Psi s \setminus \|s\|_{\ell^0} \leqslant k\}$$

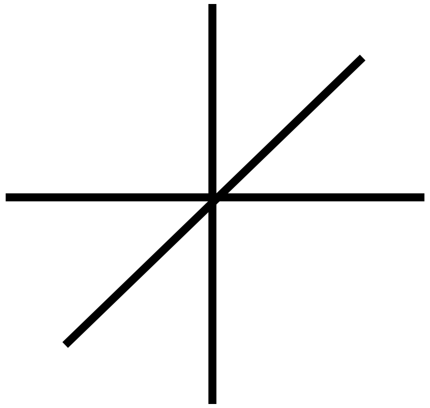

$k = 1$
(coordinate axes)

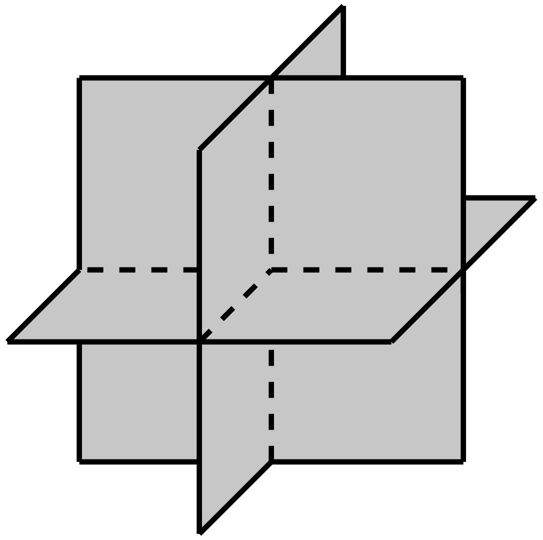

$k = 2$

$\mathcal{M}$ is not a smooth manifold, union of k-dimensional spaces.

Dictionary Learning

Input: set of patch exemplar $P = (p_i)_i$.

Learning Ψ*:*

$$\min_{\Psi=(\psi_j)_j, S=(s_i)_i} \|P - \Psi S\|^2 = \sum_i \|p_i - \Psi s_i\|^2 \quad \text{subject to} \quad \begin{cases} \|s_i\|_{\ell^0} \leqslant k, \\ \|\psi_j\|_{\ell^2} = 1. \end{cases}$$

Step #1: Ψ fixed,

$$s_i \leftarrow \mathrm{Proj}_{\mathcal{M}}(p_i) = \operatorname*{argmin}_s \|p_i - \Psi s\| \quad \text{subject to} \quad \|s\|_{\ell^0} \leqslant k.$$

$\longrightarrow$ sparse coding, approximation with pursuits (MP, OMP, BP, etc).

Step #2: $(p_i)_i$ fixed, Ψ computed by linear best fit

$$\Psi \leftarrow P S^+ \text{ where } S^+ = (S^* S)^{-1} S^*.$$

Input patches $(p_i)_i$

Learning Ψ $\longrightarrow$

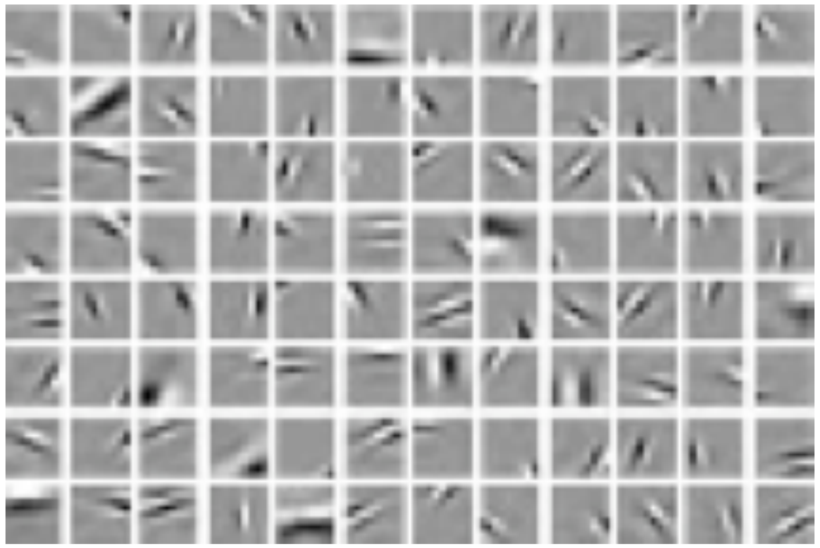

Atoms $(\psi_j)_j$

[Olshausen, Field, 1996] natural images leads to oriented wavelets Ψ.

Other algorithms: MOD [Engan et al., 1999], K-SVD [Aharon, Elad, 2006], etc.

Dictionary Learned from a Texture

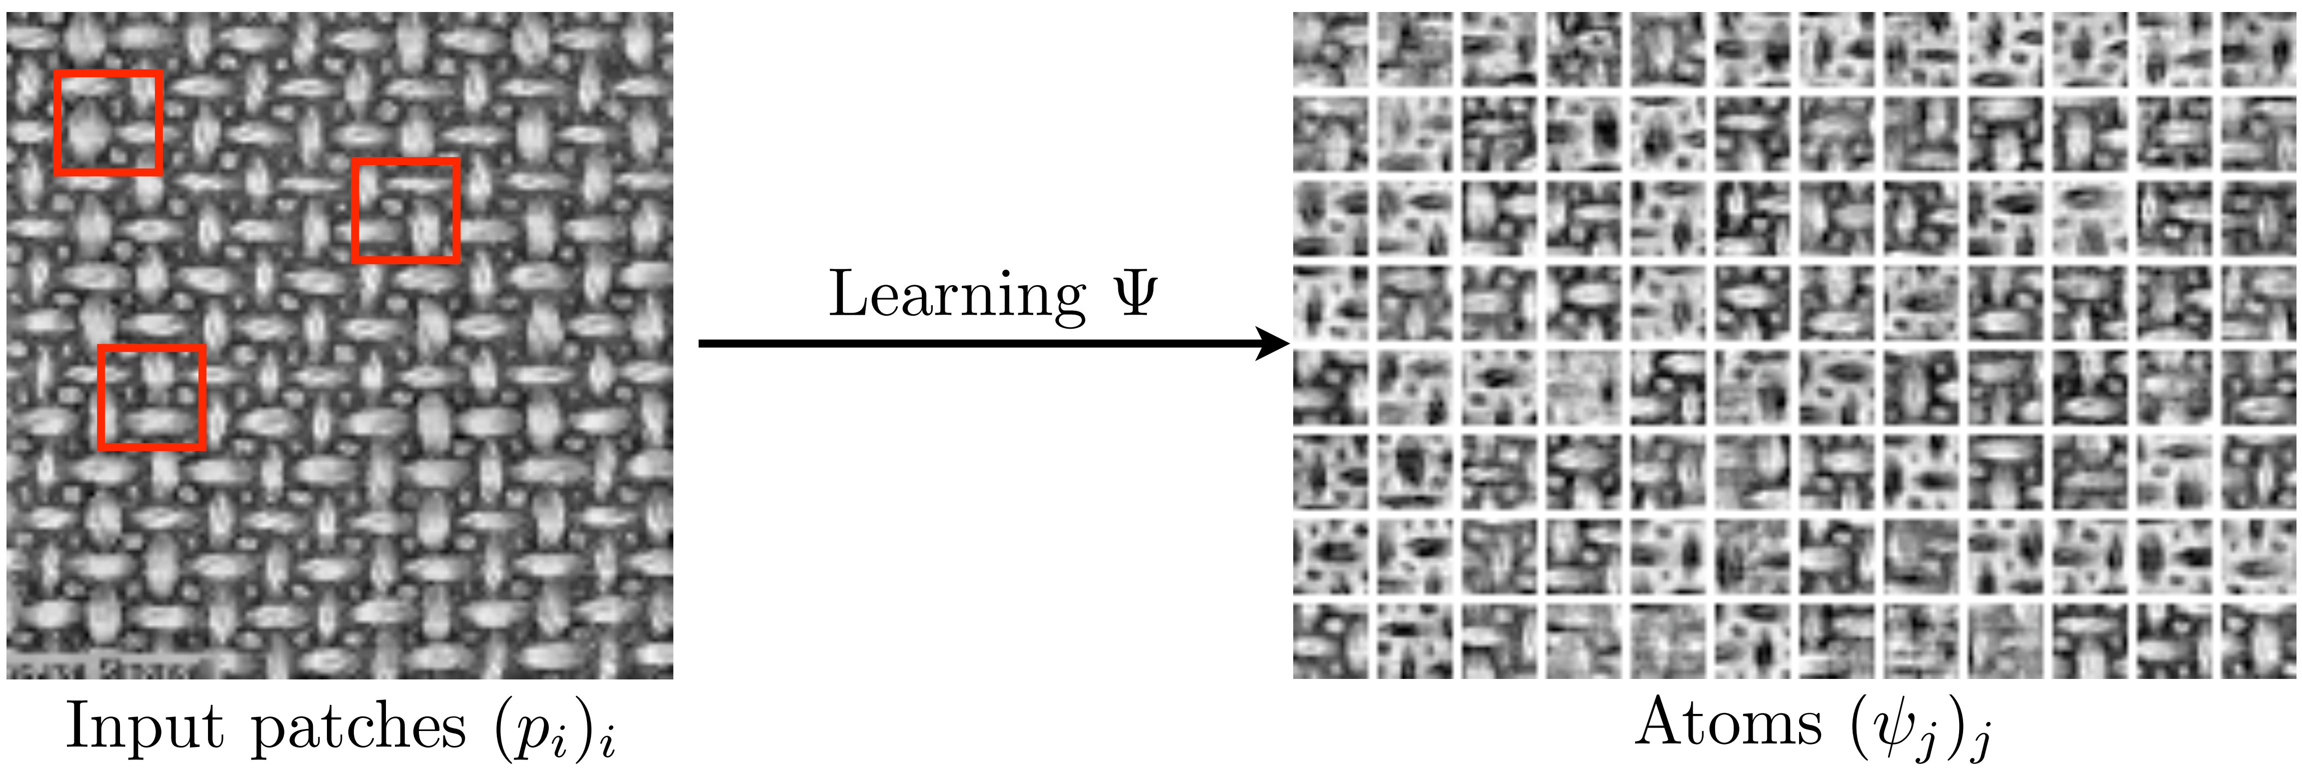

$\longrightarrow$ texturelets-like atoms representing texture patterns.

$\longrightarrow$ works for homogeneous texture (otherwise Gabor-like atoms).

Overview

- Manifolds: Image Libraries vs. Patches
- Examples of Patch Manifolds
- Manifold Energies for Inverse Problems
- Non-adaptive Manifold Models
- Texture Synthesis with Manifold Models
- Adaptive Manifold Models

Inverse Problems

Recovering f from q noisy measurements $y = \Phi f$+noise.

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

Inverse Problems

Recovering f from q noisy measurements $y = \Phi f$+noise.

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

Denoising: $\Phi =$identity, $q = N$.

Inverse Problems

Recovering f from q noisy measurements $y = \Phi f$+noise.

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

Denoising: Φ =identity, $q = N$.

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

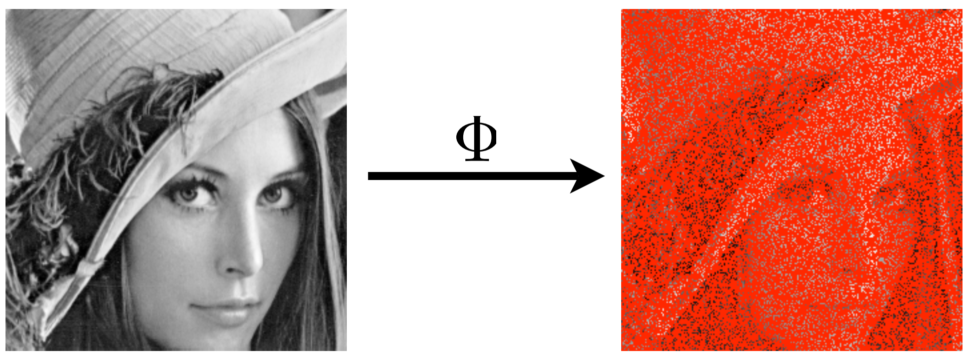

$$(\Phi f)(x) = \begin{cases} 0 & \text{if} \quad x \in \Omega, \\ f(x) & \text{if} \quad x \notin \Omega. \end{cases}$$

Inverse Problems

Recovering f from q noisy measurements $y = \Phi f$+noise.

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

Denoising: Φ =identity, $q = N$.

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

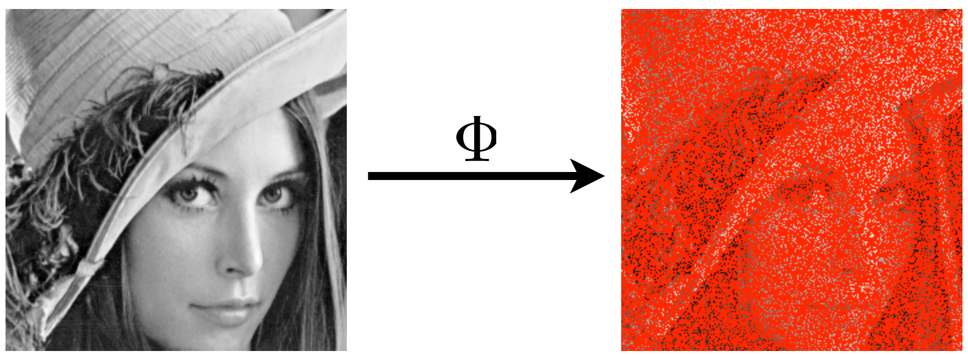

$$(\Phi f)(x) = \begin{cases} 0 & \text{if } \quad x \in \Omega, \\ f(x) & \text{if } \quad x \notin \Omega. \end{cases}$$

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

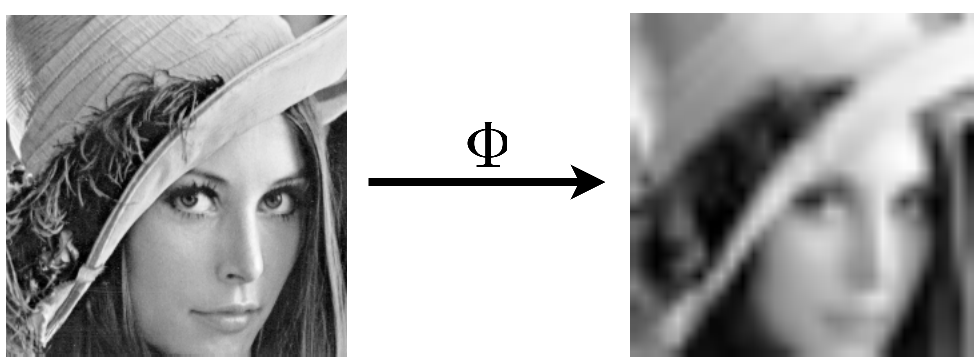

Inverse Problems

Recovering f from q noisy measurements $y = \Phi f$+noise.

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

Denoising: Φ =identity, $q = N$.

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

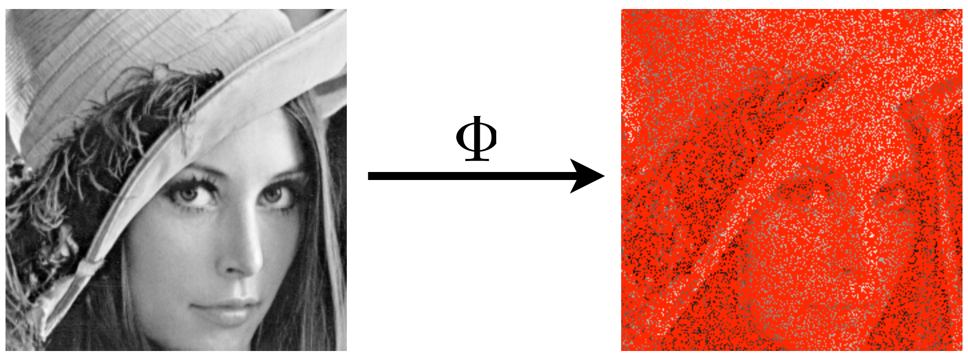

$$(\Phi f)(x) = \begin{cases} 0 & \text{if} \quad x \in \Omega, \\ f(x) & \text{if} \quad x \notin \Omega. \end{cases}$$

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

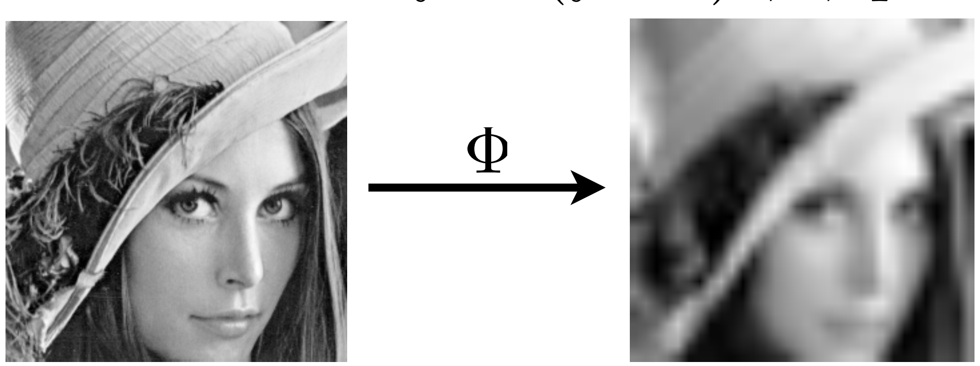

Compressed sampling: $(\Phi f)_i = \langle f, \varphi_i \rangle$, φ_i random vector.

$\Phi f \in \mathbb{R}^q$ is a "compressed" version of f.

CS theory [Candès, Tao, Donoho, 2004]:

f can be well recovered if f is sparse in an ortho-basis.

Inverse Problems Regularization

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

Penalized inversion: $$f^\star = \underset{g}{\operatorname{argmin}} \frac{1}{2}\|\Phi g - y\|^2 + \lambda J(g)$$

λ should be adapted to the measurement noise $\|\Phi f - y\|$ and the prior $J(f)$

$\Longrightarrow$ difficult in practice ...

Inverse Problems Regularization

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

Penalized inversion: $$f^\star = \underset{g}{\operatorname{argmin}} \frac{1}{2}\|\Phi g - y\|^2 + \lambda J(g)$$

λ should be adapted to the measurement noise $\|\Phi f - y\|$ and the prior $J(f)$

$\Longrightarrow$ difficult in practice ...

Sobolev regularization: $$J(f) = \int \|\nabla_x f\|^2 \mathrm{d}x$$

Total variation regularization: $$J(f) = \int \|\nabla_x f\| \mathrm{d}x$$

Sparse wavelets regularization: $J(f) = \sum_i |\langle f,\, \psi_i \rangle|$ where $\{\psi_i\}_i$ wavelet basis.

Inverse Problems Regularization

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

Penalized inversion: $$f^\star = \underset{g}{\operatorname{argmin}} \frac{1}{2}\|\Phi g - y\|^2 + \lambda J(g)$$

λ should be adapted to the measurement noise $\|\Phi f - y\|$ and the prior $J(f)$

$\Longrightarrow$ difficult in practice ...

Sobolev regularization: $$J(f) = \int \|\nabla_x f\|^2 \mathrm{d}x$$

Total variation regularization: $$J(f) = \int \|\nabla_x f\| \mathrm{d}x$$

Sparse wavelets regularization: $J(f) = \sum_i |\langle f, \psi_i \rangle|$ where $\{\psi_i\}_i$ wavelet basis.

Manifold regularization:

Non-adaptive regularization: $\mathcal{M}$ fixed from a image model $f \in \Theta$.

$J_{\mathcal{M}}(g)$ measures how much patches $\mathcal{C}_f = (p_x(f))_x$ are close to $\mathcal{M}$.

Adaptive regularization: $\mathcal{M} = \mathcal{M}_f = (p_x(f))_x$ estimated from some f.

$J_w(g)$ measures the smoothness of g with respect to the geometry of $\mathcal{M}_f$.

w is a graph that represent the geometry of $\mathcal{M}_f$.

Overview

- Manifolds: Image Libraries vs. Patches
- Examples of Patch Manifolds
- Manifold Energies for Inverse Problems
- Non-adaptive Manifold Models
- Texture Synthesis with Manifold Models
- Adaptive Manifold Model

Non-adaptive Manifold Energies

Setting #1: manifold $\mathcal{M}$ defined by an a priori model $f \in \Theta$.

$$\mathcal{M} = \left\{ p_x(g) \setminus x \in [0,1]^d \quad \text{and} \quad g \in \Theta \right\} \subset \mathrm{L}^2([-\tau/2, \tau/2]).$$

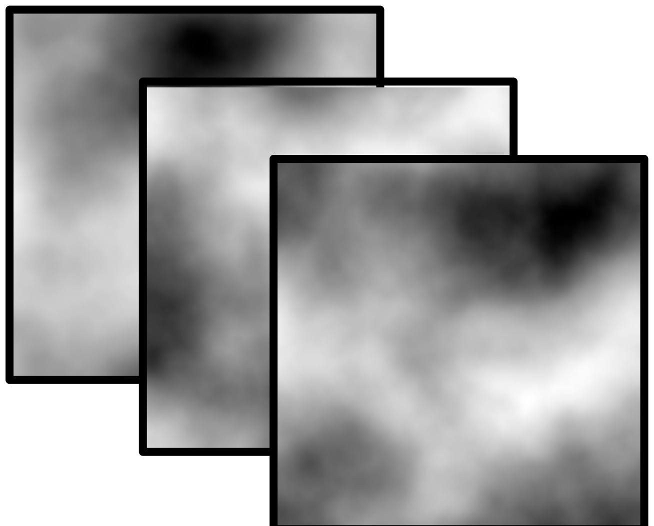

Θ =smooth images

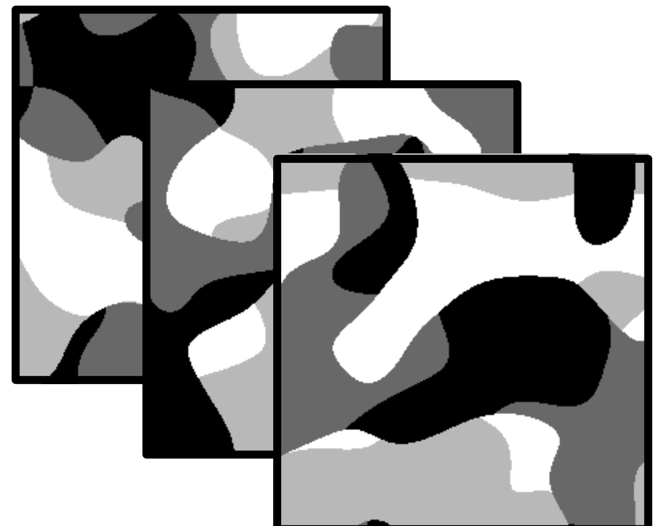

Θ =cartoon images

Θ =oscilating textures

Non-adaptive Manifold Energies

Setting #1: manifold $\mathcal{M}$ defined by an a priori model $f \in \Theta$.

$$\mathcal{M} = \{p_x(g) \setminus x \in [0,1]^d \quad \text{and} \quad g \in \Theta\} \subset \mathrm{L}^2([-\tau/2, \tau/2]).$$

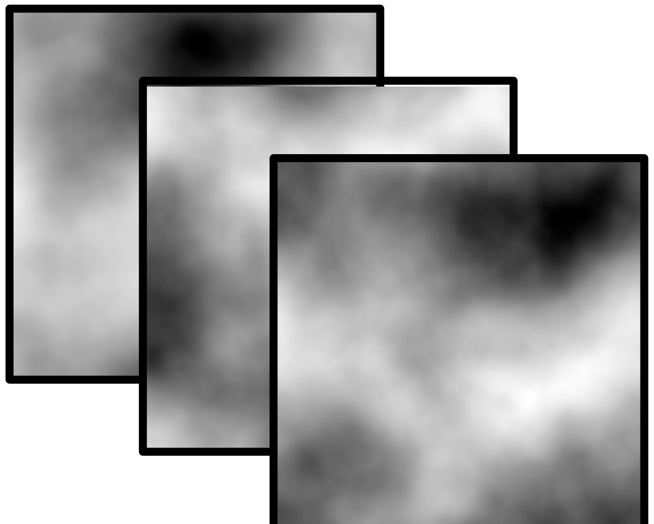

Θ =smooth images

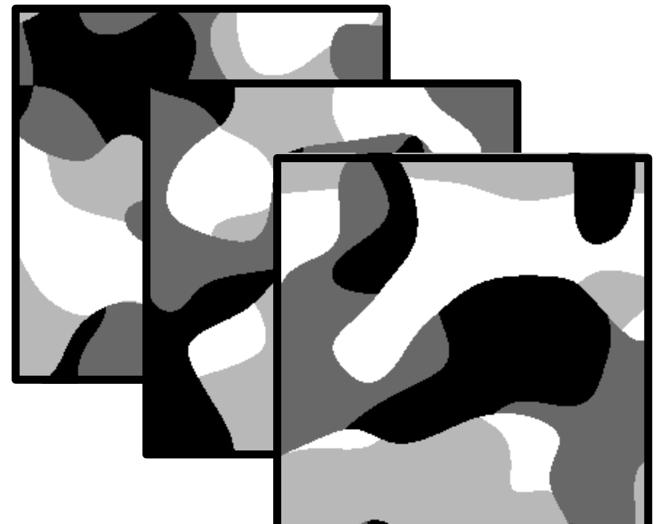

Θ =cartoon images

Θ =oscilating textures

Non-adaptive manifold energy: $J_{\mathcal{M}}(g) = \int \mathrm{dist}_{\mathcal{M}}(p_x(g))\mathrm{d}x$

where $\mathrm{dist}_{\mathcal{M}}(p) = \min_{q \in \mathcal{M}} \|p - q\| = \|p - \mathrm{Proj}_{\mathcal{M}}(p)\|$

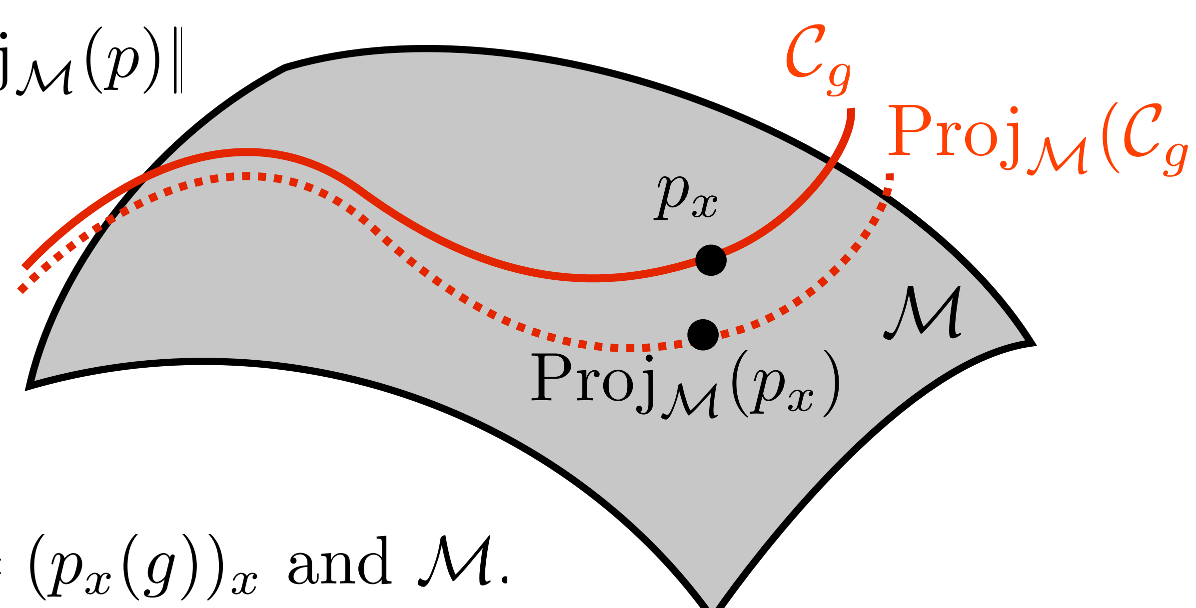

$\longrightarrow$ $J_{\mathcal{M}}(g)$ is small if $\forall\, x,\ p_x(g)$ is close to $\mathcal{M}$.

$\longrightarrow$ Average distance between the "surface" $\mathcal{C}_g = (p_x(g))_x$ and $\mathcal{M}$.

Non-adaptive Manifold Energy Minimization

Manifold energy: $J_{\mathcal{M}}(g) = \sum_x \mathrm{dist}_{\mathcal{M}}(p_x(g))$

Regularized inversion:
$$f^\star = \underset{g}{\mathrm{argmin}}\ \|y - \Phi g\|^2 + \lambda J_{\mathcal{M}}(g)$$
$$\{f^\star, (p_x^\star)\} = \underset{g,(p_x)_x}{\mathrm{argmin}}\ \|y - \Phi g\|^2 + \lambda \sum_x \|p_x(g) - p_x\|^2$$

Include patches $(p_x)_x$

Non-adaptive Manifold Energy Minimization

Manifold energy: $J_{\mathcal{M}}(g) = \sum_x \text{dist}_{\mathcal{M}}(p_x(g))$

Regularized inversion:
$$f^\star = \underset{g}{\text{argmin}}\; \|y - \Phi g\|^2 + \lambda J_{\mathcal{M}}(g)$$

$$\{f^\star, (p^\star_x)\} = \underset{g,(p_x)_x}{\text{argmin}}\; \|y - \Phi g\|^2 + \lambda \sum_x \|p_x(g) - p_x\|^2$$

Include patches $(p_x)_x$

Step #1: the image $f^\star$ is fixed, $\quad p^\star_x \leftarrow \text{Proj}_{\mathcal{M}}(p_x(f^\star))$.

Step #2: $(p^\star_x)_x$ fixed, $f^\star$ computed by linear best fit

$$(\Phi^*\Phi + \lambda \text{Id})\, f^\star = \Phi^* y + \lambda \bar{p}^\star$$

where $\quad \bar{p}^\star(x) = \frac{1}{\tau^2} \sum_{|x-y| \leqslant \tau/2} p^\star_y(x - y)$

Non-adaptive Manifold Energy Minimization

Manifold energy: $J_{\mathcal{M}}(g) = \sum_x \mathrm{dist}_{\mathcal{M}}(p_x(g))$

Regularized inversion:

$$f^\star = \underset{g}{\mathrm{argmin}}\ \|y - \Phi g\|^2 + \lambda J_{\mathcal{M}}(g)$$

$$\{f^\star, (p_x^\star)\} = \underset{g,(p_x)_x}{\mathrm{argmin}}\ \|y - \Phi g\|^2 + \lambda \sum_x \|p_x(g) - p_x\|^2$$

Include patches $(p_x)_x$

Step #1: the image $f^\star$ is fixed, $\quad p_x^\star \leftarrow \mathrm{Proj}_{\mathcal{M}}(p_x(f^\star))$.

Step #2: $(p_x^\star)_x$ fixed, $f^\star$ computed by linear best fit

$$(\Phi^*\Phi + \lambda \mathrm{Id})\, f^\star = \Phi^* y + \lambda \bar{p}^\star$$

where $\quad \bar{p}^\star(x) = \dfrac{1}{\tau^2} \displaystyle\sum_{|x-y|\leqslant \tau/2} p_y^\star(x-y)$

Manifold $\mathcal{M}$ of smooth patches.

$\mathcal{M} =$ 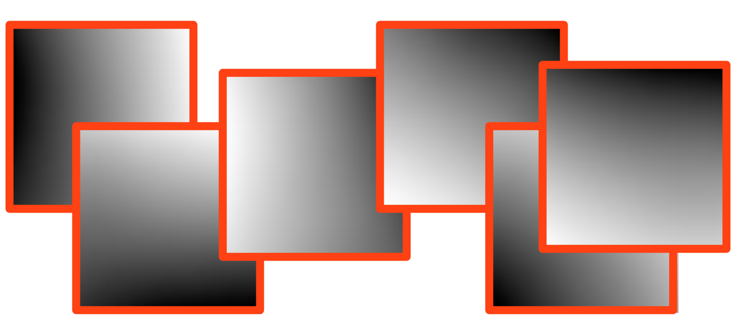

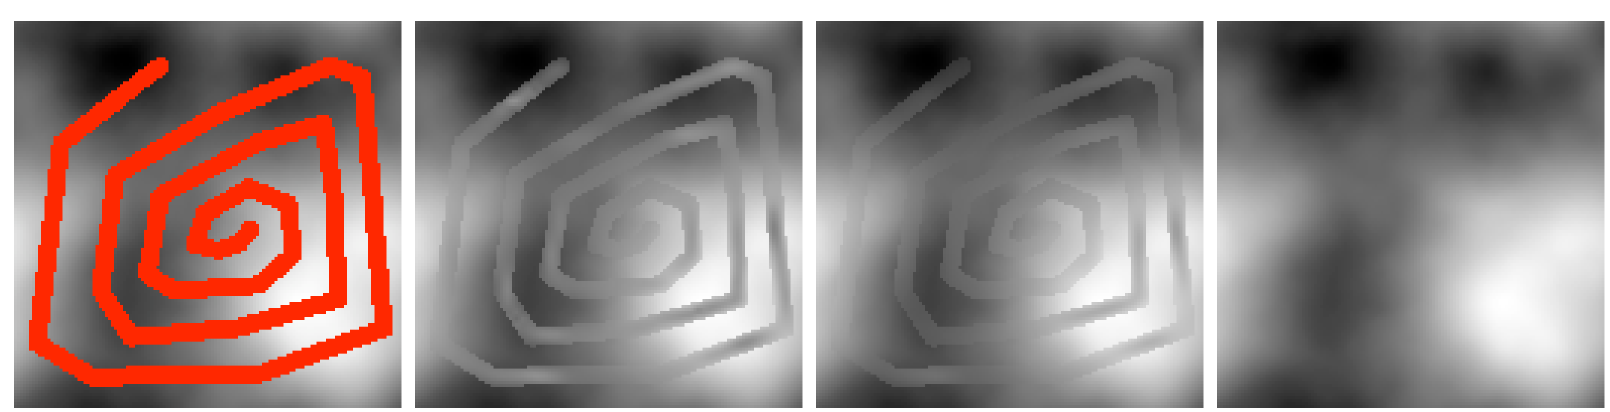

Measurements y | Iter. #1 | Iter. #3 | Iter. #50

Cartoon Manifold Model

Manifold of affine edges: $\mathcal{M} = \{a + bP_{\theta,\delta} \setminus a, b, \theta, \delta\}$

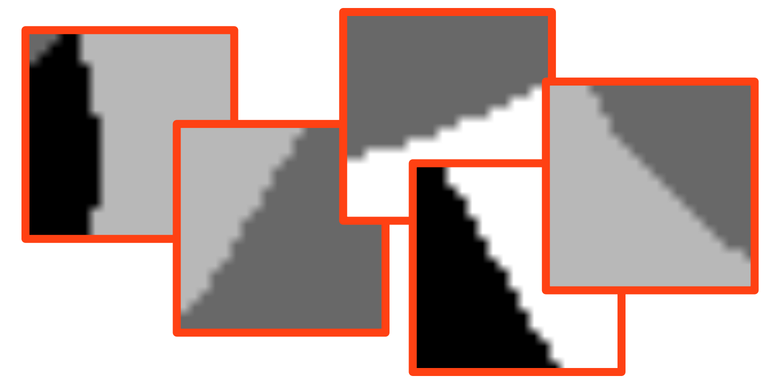

where $\begin{cases} P_{\theta,\delta}(t) = P_{0,0}(R_\theta(t-\delta)) \\ P_{0,0}(x) = 1_{x_1 \geqslant 0}(x) \end{cases}$

Inpainting: $y = 1_\Omega \cdot f, \quad \Omega^c$ =missing pixels.

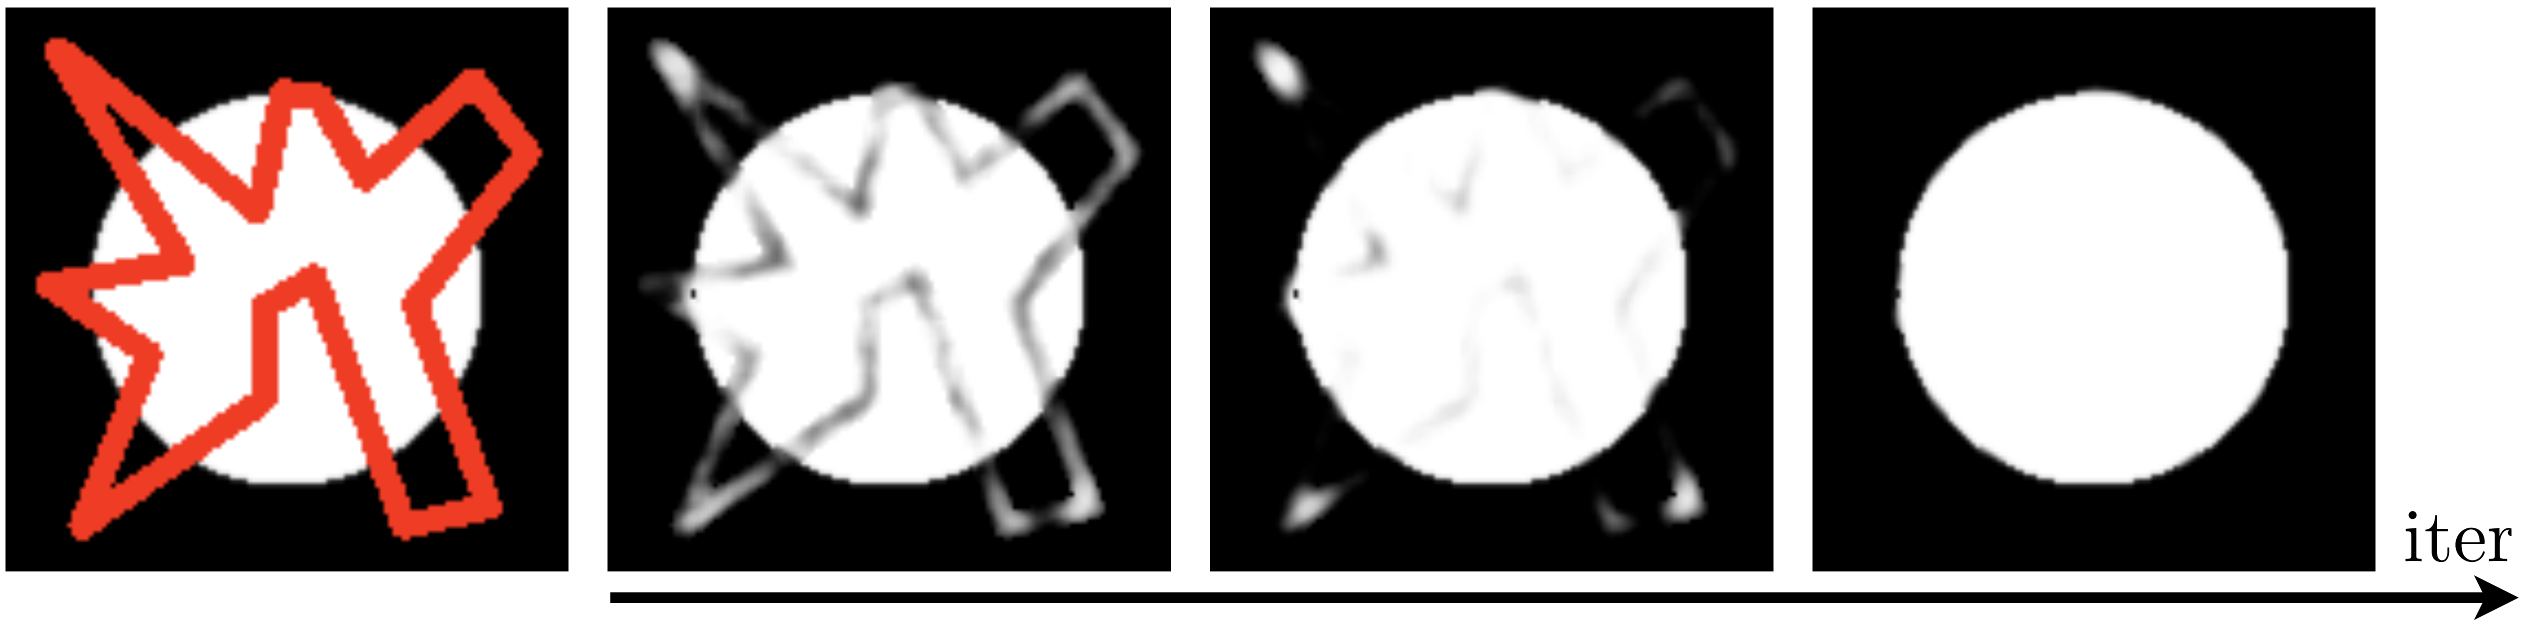

Compressed Sensing recovery: $y = \Phi f,\ \Phi \in \mathbb{R}^{N \times P},\ P = N/4$ random measures.

Wavelets, SNR=25.7dB   Manifold, SNR=31.3dB

Oscilating Texture Manifold

Adapting the oscilating profile: $\mathcal{M}_\gamma = \left\{ A P^\gamma_{\rho,\theta,\delta} \setminus A, \rho, \theta \right\}$

where $P^\gamma_{\rho,\theta,\delta}(t) = \cos^\gamma(\rho\langle t,\, \theta\rangle + \delta)$

$\gamma = 1$

$\gamma = 0.5$

$\gamma = 0.1$

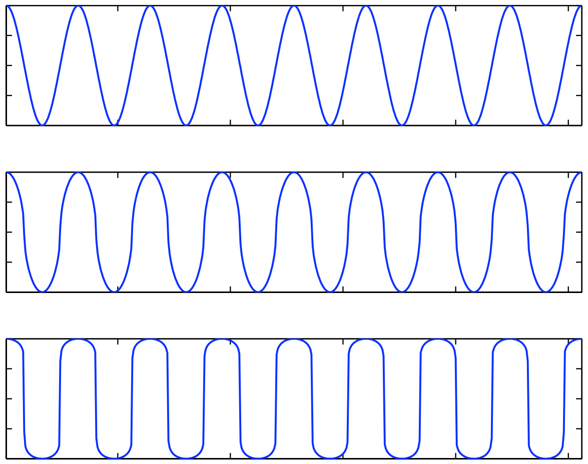

Inpainting with oscillating manifold:

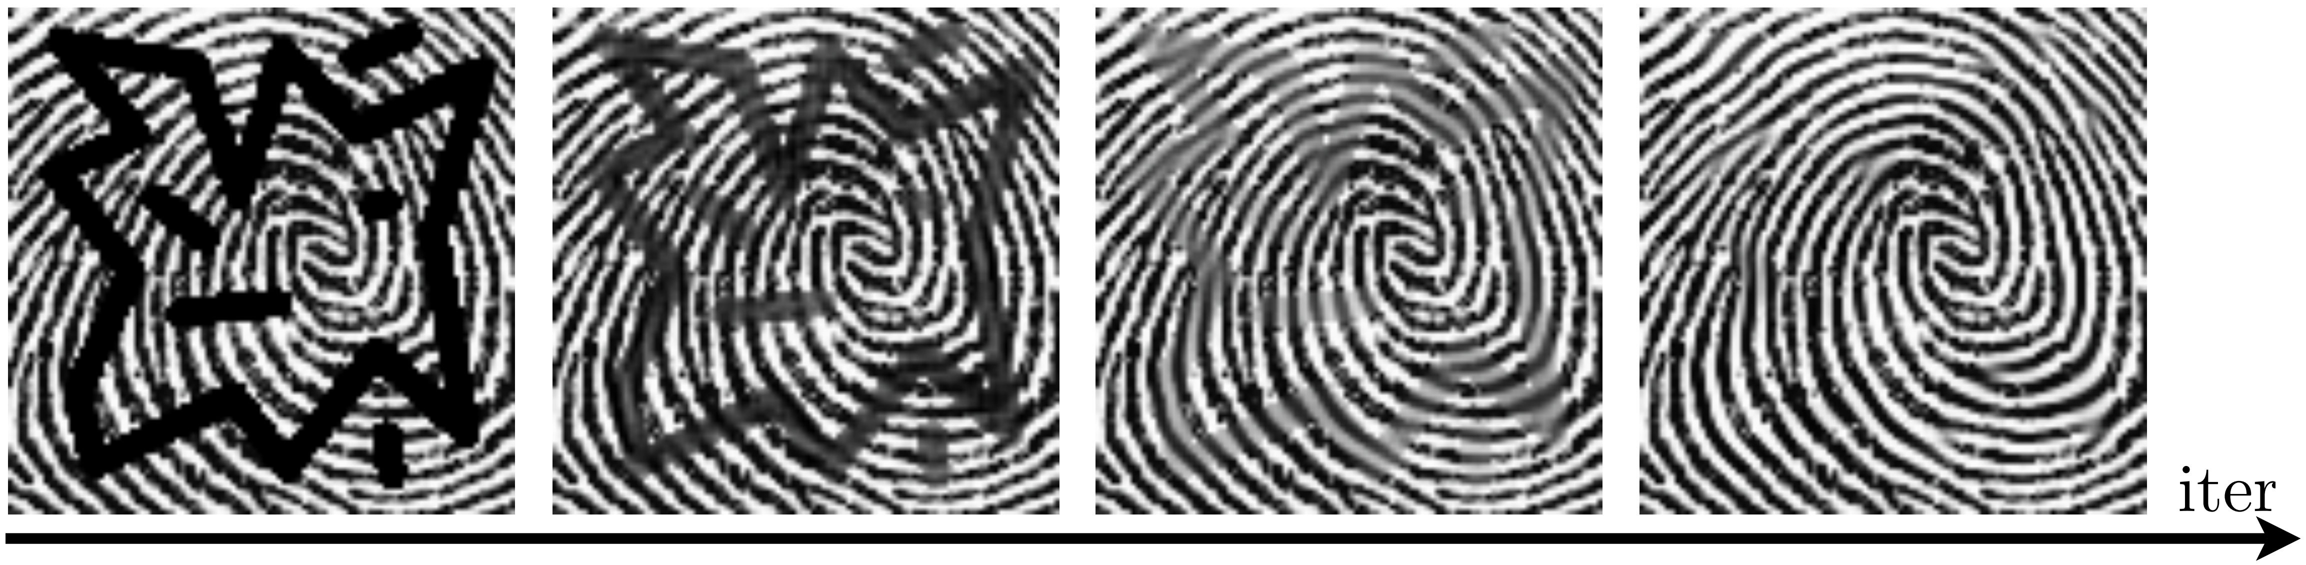

Compressed Sensing recovery: $y = \Phi f$, $\Phi \in \mathbb{R}^{N \times P}$, $P = N/8$ random measures.

Gabor, SNR=17.9dB

Manifold, SNR=19.5dB

Overview

- Manifolds: Image Libraries vs. Patches
- Examples of Patch Manifolds
- Manifold Energies for Inverse Problems
- Non-adaptive Manifold Models
- Texture Synthesis with Manifold Models
- Adaptive Manifold Model

Texture Synthesis with Manifold Model

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

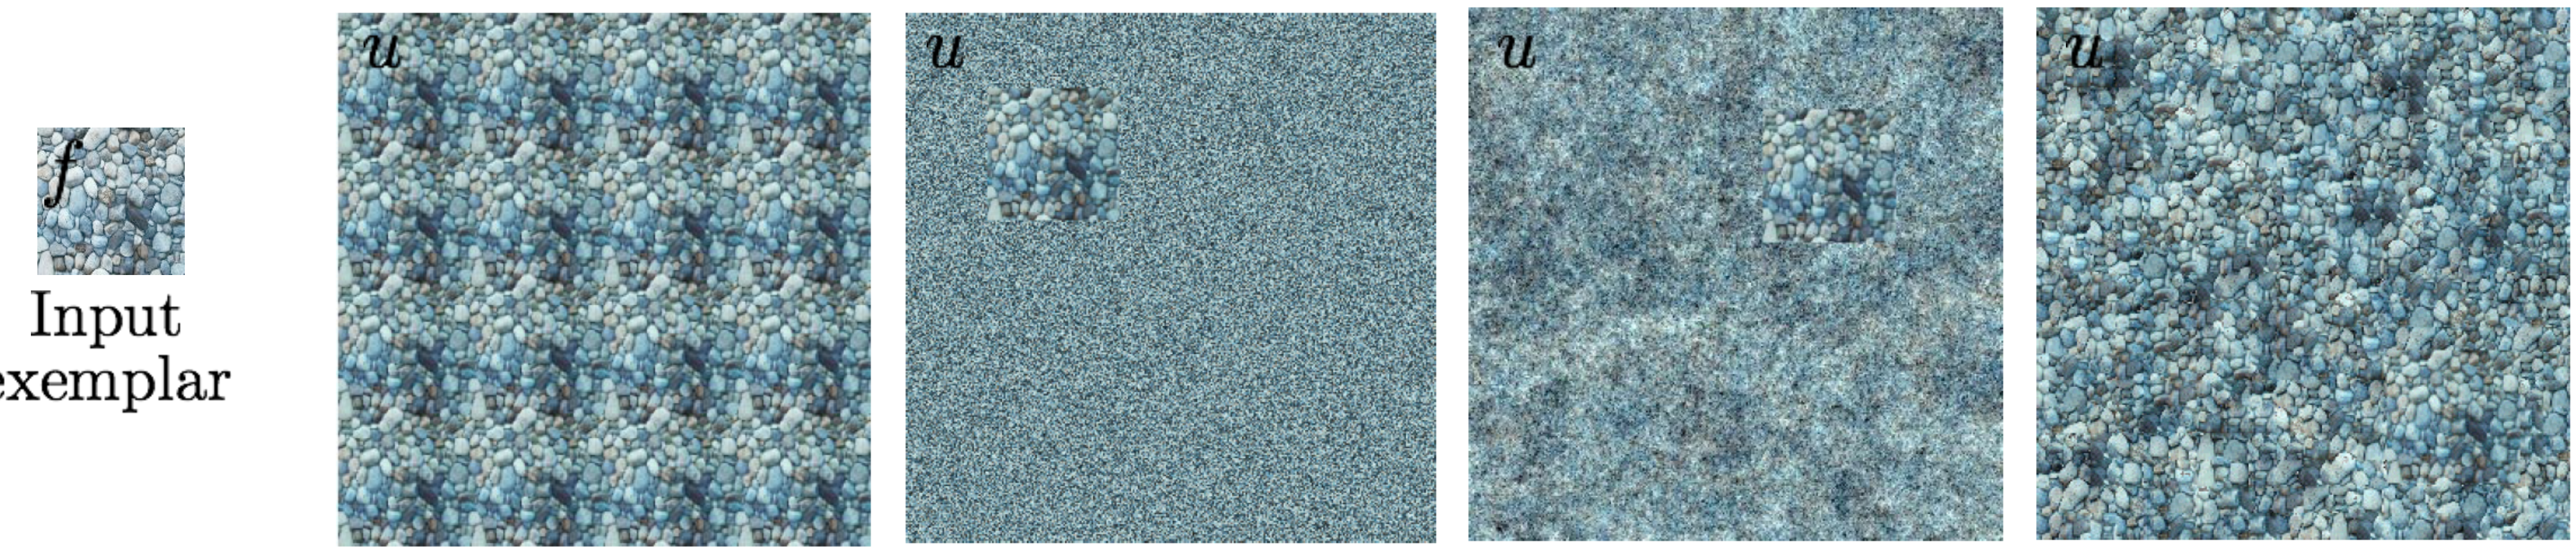

Input exemplar | Periodic copy | Spatial matching | Wavelet matching | Clever copy

Texture Synthesis with Manifold Model

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

Input exemplar

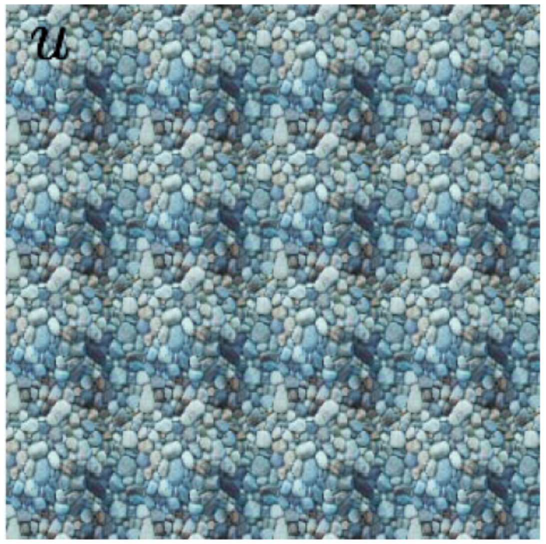

Periodic copy

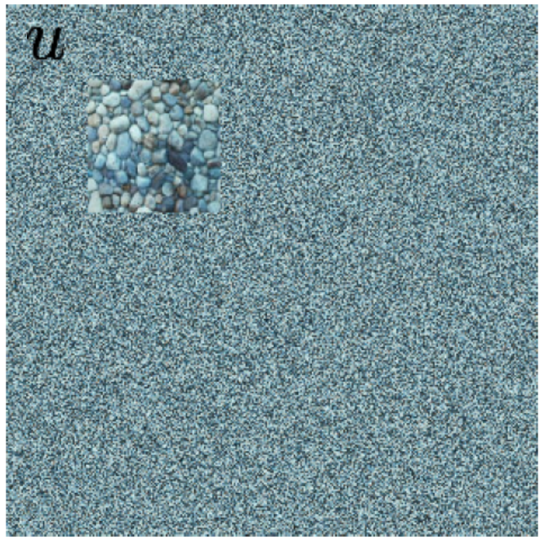

Spatial matching

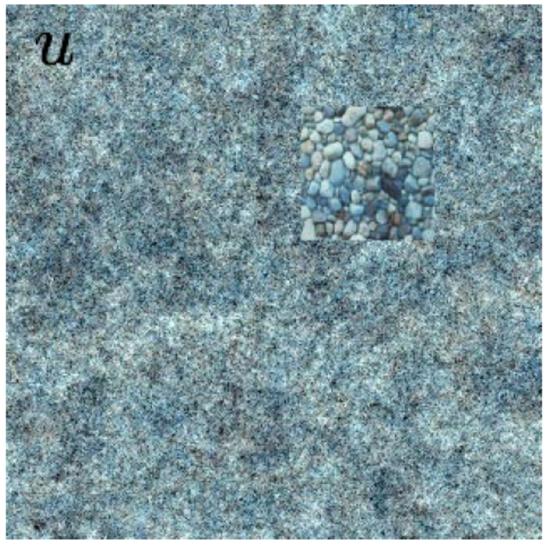

Wavelet matching

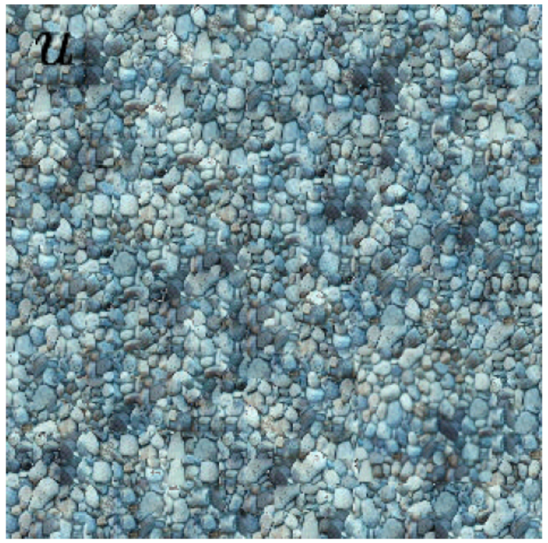

Clever copy

Manifold synthesis: draw $f^\star$ at random in $\Theta \cap \mathcal{C}$.

$$\Theta = \{ f \setminus \forall\, x,\ p_x(f) \in \mathcal{M} \}$$

$\mathcal{C}$ is an additional set of constraints (energy $\|f^\star\| = c$, histogram, etc).

Texture Synthesis with Manifold Model

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

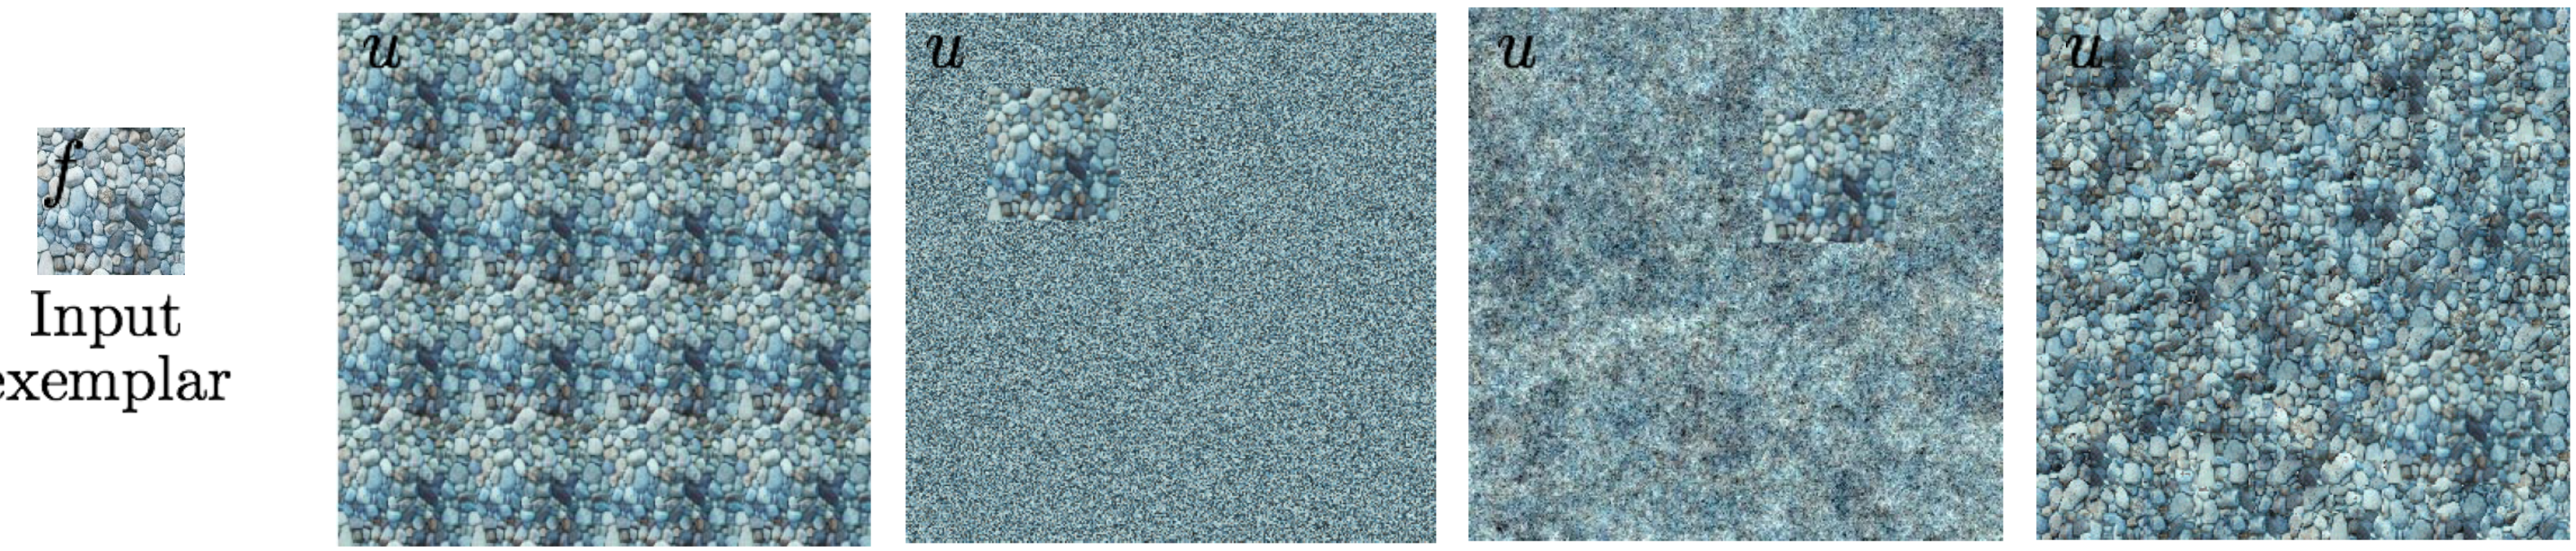

Input exemplar | Periodic copy | Spatial matching | Wavelet matching | Clever copy

Manifold synthesis: draw $f^\star$ at random in $\Theta \cap \mathcal{C}$.

$$\Theta = \{ f \setminus \forall\, x,\ p_x(f) \in \mathcal{M} \}$$

$\mathcal{C}$ is an additional set of constraints (energy $\|f^\star\| = c$, histogram, etc).

Iterative projection algorithm:

Step #1: the image $f^\star$ is fixed, $\quad p_x^\star \leftarrow \mathrm{Proj}_{\mathcal{M}}(p_x(f^\star))$.

Step #2: $(p_x^\star)_x$ fixed, $f^\star$ computed by averaging

$$f^\star(x) = \frac{1}{\tau^2} \sum_{|x-y| \leqslant \tau/2} p_y^\star(x-y)$$

Step 3: impose additional constraints: $f^\star \leftarrow \mathrm{Proj}_{\mathcal{C}}(f^\star)$.

Examples of Manifold Synthesis

Synthesis with edge manifold:

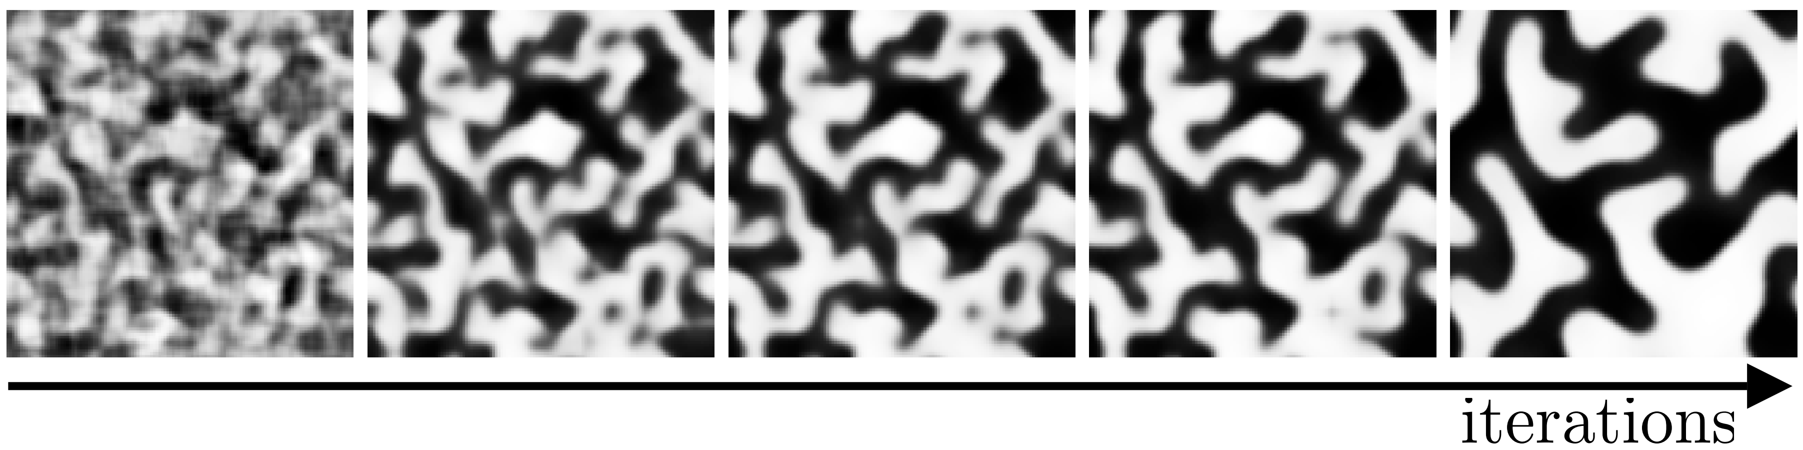

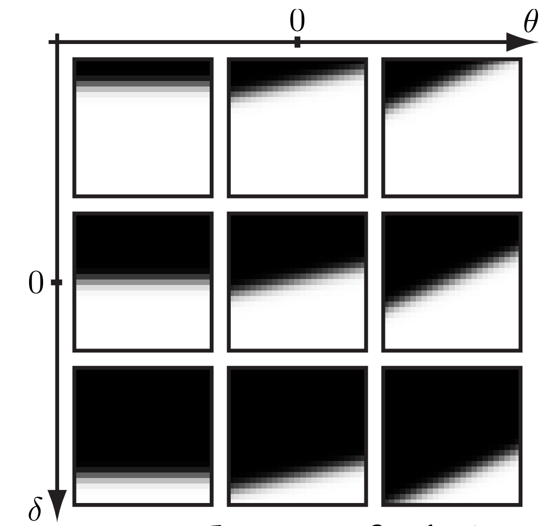

patches of $\mathcal{M}$

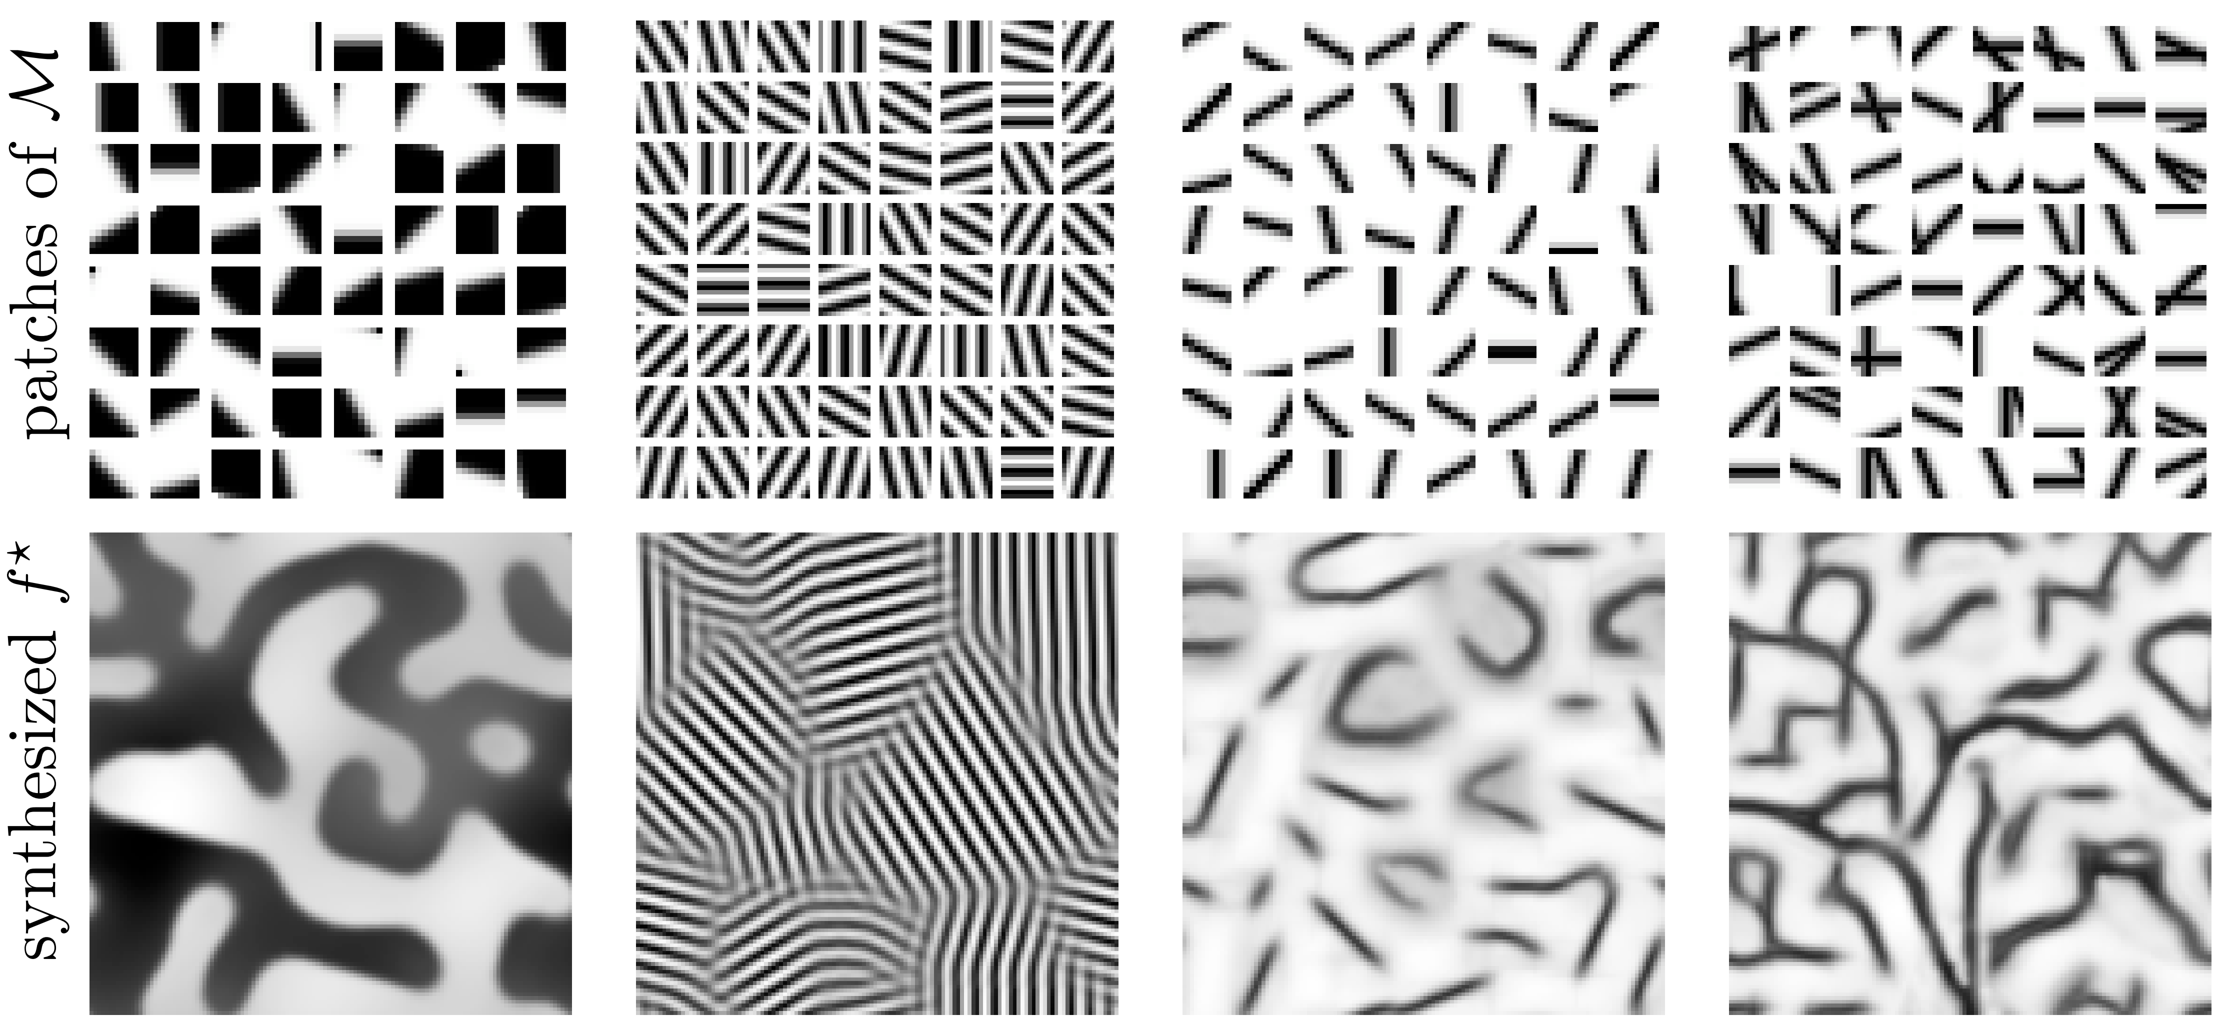

Synthesis with Sparse Manifold

Dictionary $\Psi = (\psi_j)_j$ learned from input texture f.

$$\mathcal{M} = \{\Psi s \setminus \|s\|_{\ell^0} \leqslant k\} \qquad \Theta = \{f \setminus \forall\, x,\; p_x(f) = \Psi s_x \quad \text{with} \quad \|s_x\|_{\ell^0} \leqslant k\}$$

$$\mathrm{Proj}_{\mathcal{M}}(p_x) = \Psi s_x^\star \quad \text{where} \quad s_x^\star = \underset{s}{\mathrm{argmin}}\, \|p_x - \Psi s\| \quad \text{subject to} \quad \|s\|_{\ell^0} \leqslant k.$$

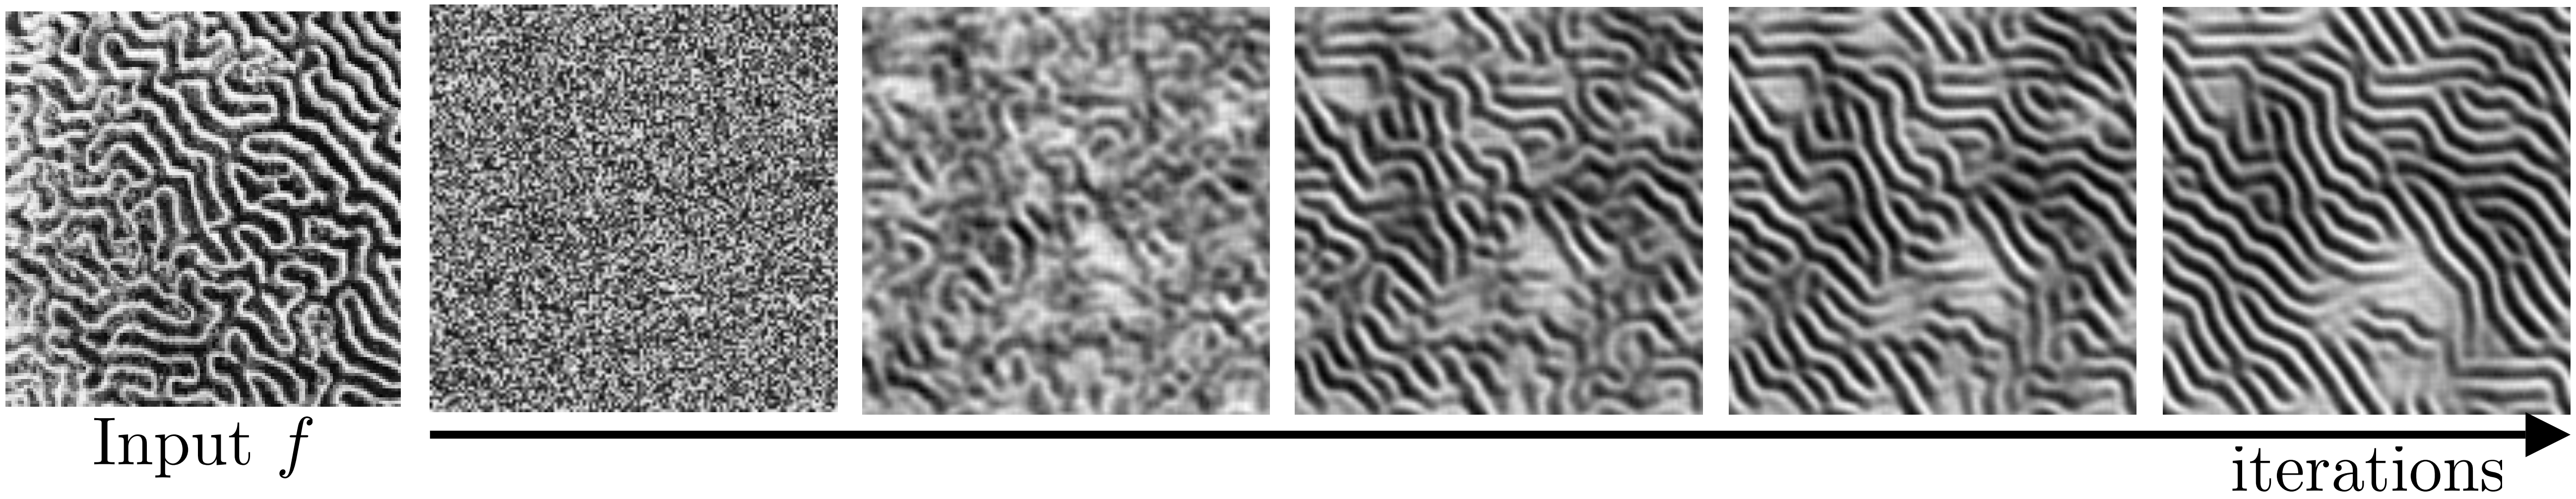

Synthesis with Sparse Manifold

Dictionary $\Psi = (\psi_j)_j$ learned from input texture f.

$$\mathcal{M} = \{\Psi s \setminus \|s\|_{\ell^0} \leqslant k\} \qquad \Theta = \{f \setminus \forall\, x,\ p_x(f) = \Psi s_x \quad \text{with} \quad \|s_x\|_{\ell^0} \leqslant k\}$$

$$\mathrm{Proj}_{\mathcal{M}}(p_x) = \Psi s_x^\star \quad \text{where} \quad s_x^\star = \underset{s}{\mathrm{argmin}}\ \|p_x - \Psi s\| \quad \text{subject to} \quad \|s\|_{\ell^0} \leqslant k.$$

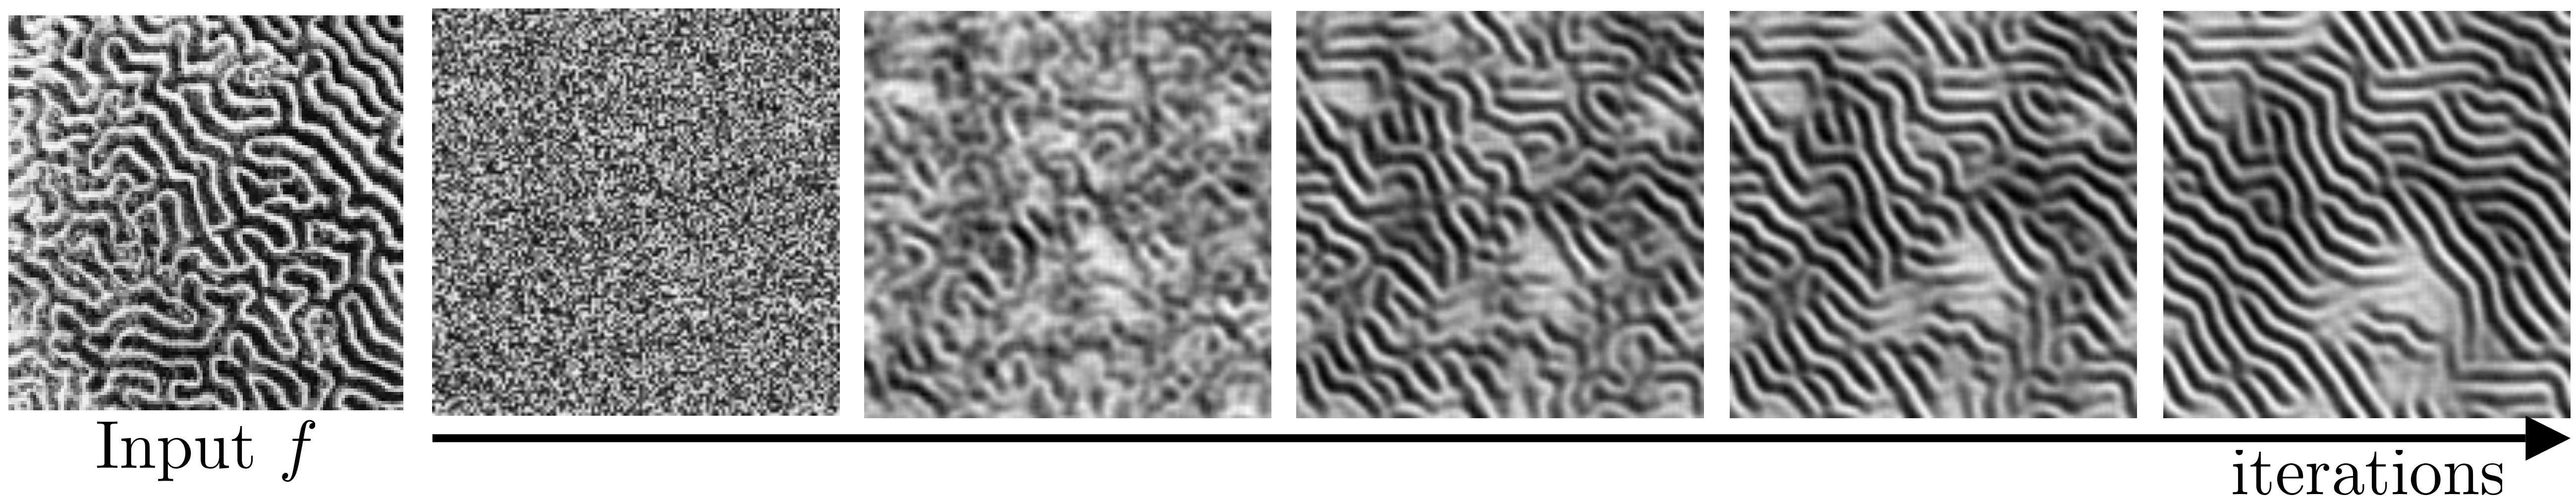

Parameterization of the texture model:

- redundancy $q/n > 1$ of the dictionary.
- sparsity $k > 1$ of the patch expansion.

See [Peyré, “Sparse modelling of textures”, 2008]

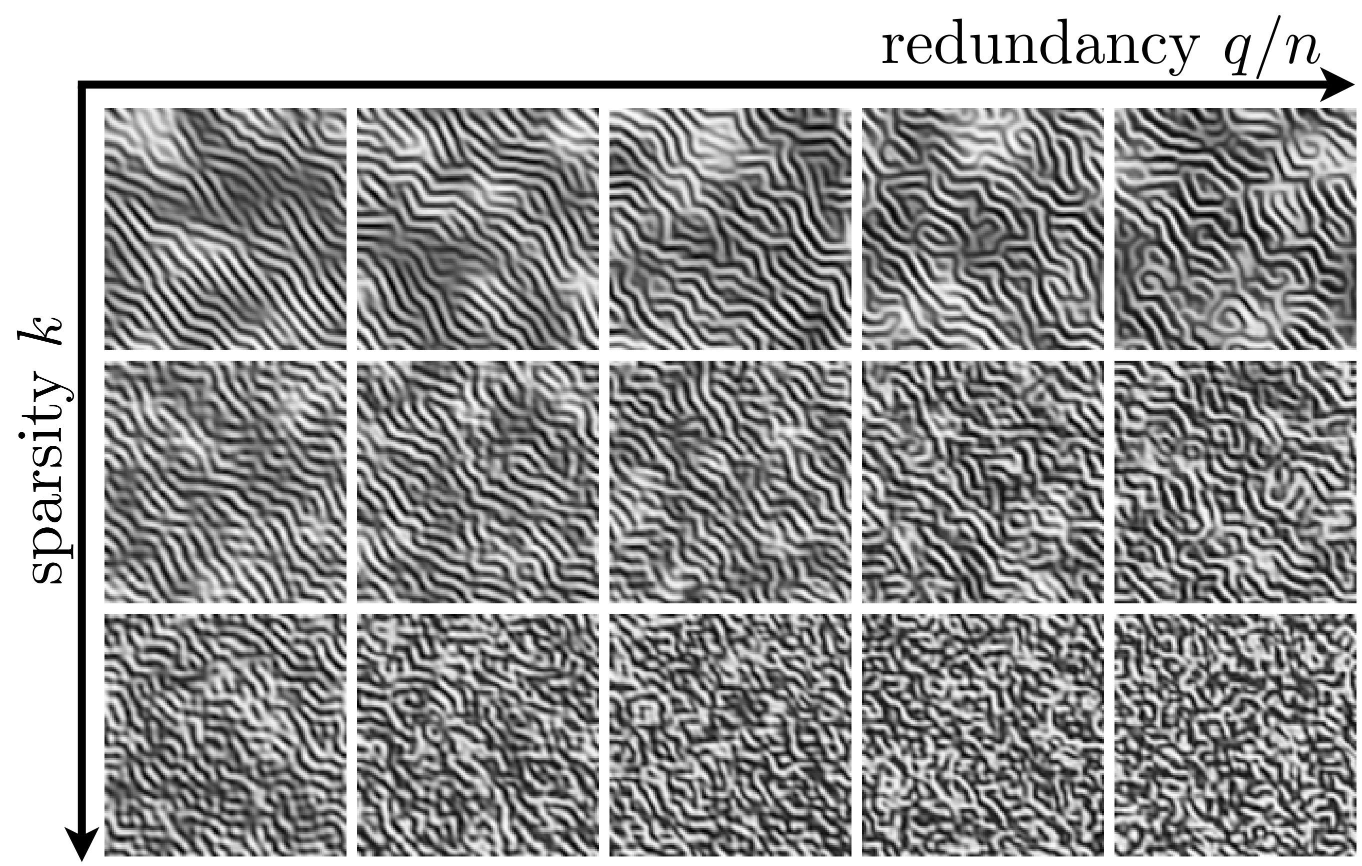

Computer Graphics Approach

f $\{p_x(f)\}_x$

Dictionary: all patches $\Psi = (p_x(f))_x$.

Sparsity $k = 1$: patch recopy.

$$p_x(f^\star) = p_{\varphi(x)}(f)$$

$$\text{where} \quad \varphi(x) = \underset{y}{\text{argmin}} \, \|p_x(f^\star) - p_y(f)\|$$

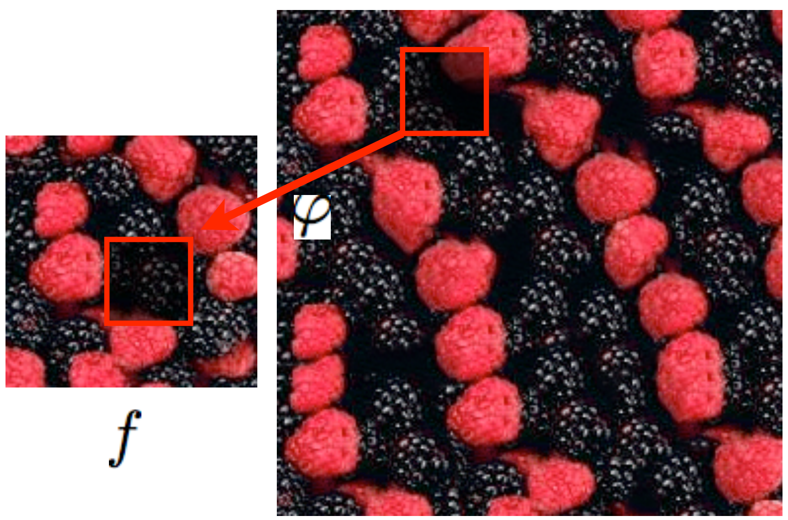

f synthesized $f^\star$

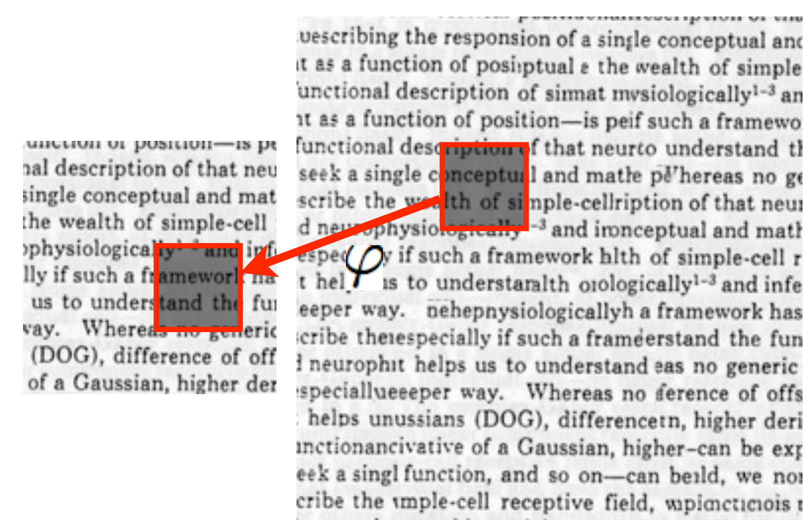

Similar to [Efros, Leung, 1999] and others, but parallel update of all p_x.

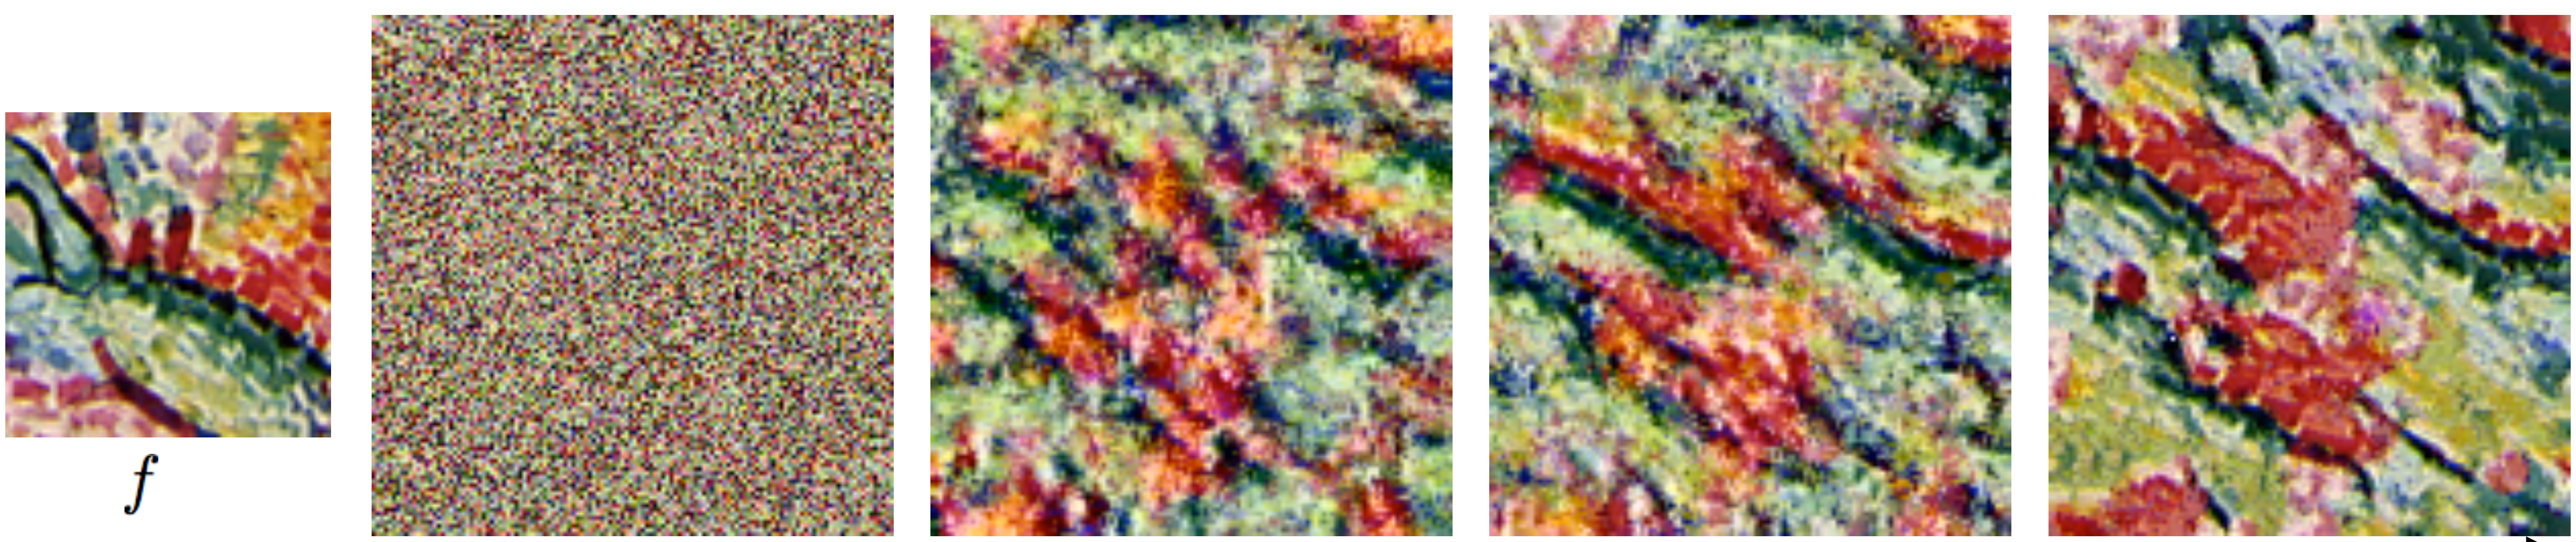

f

iterations

Texture Inpainting

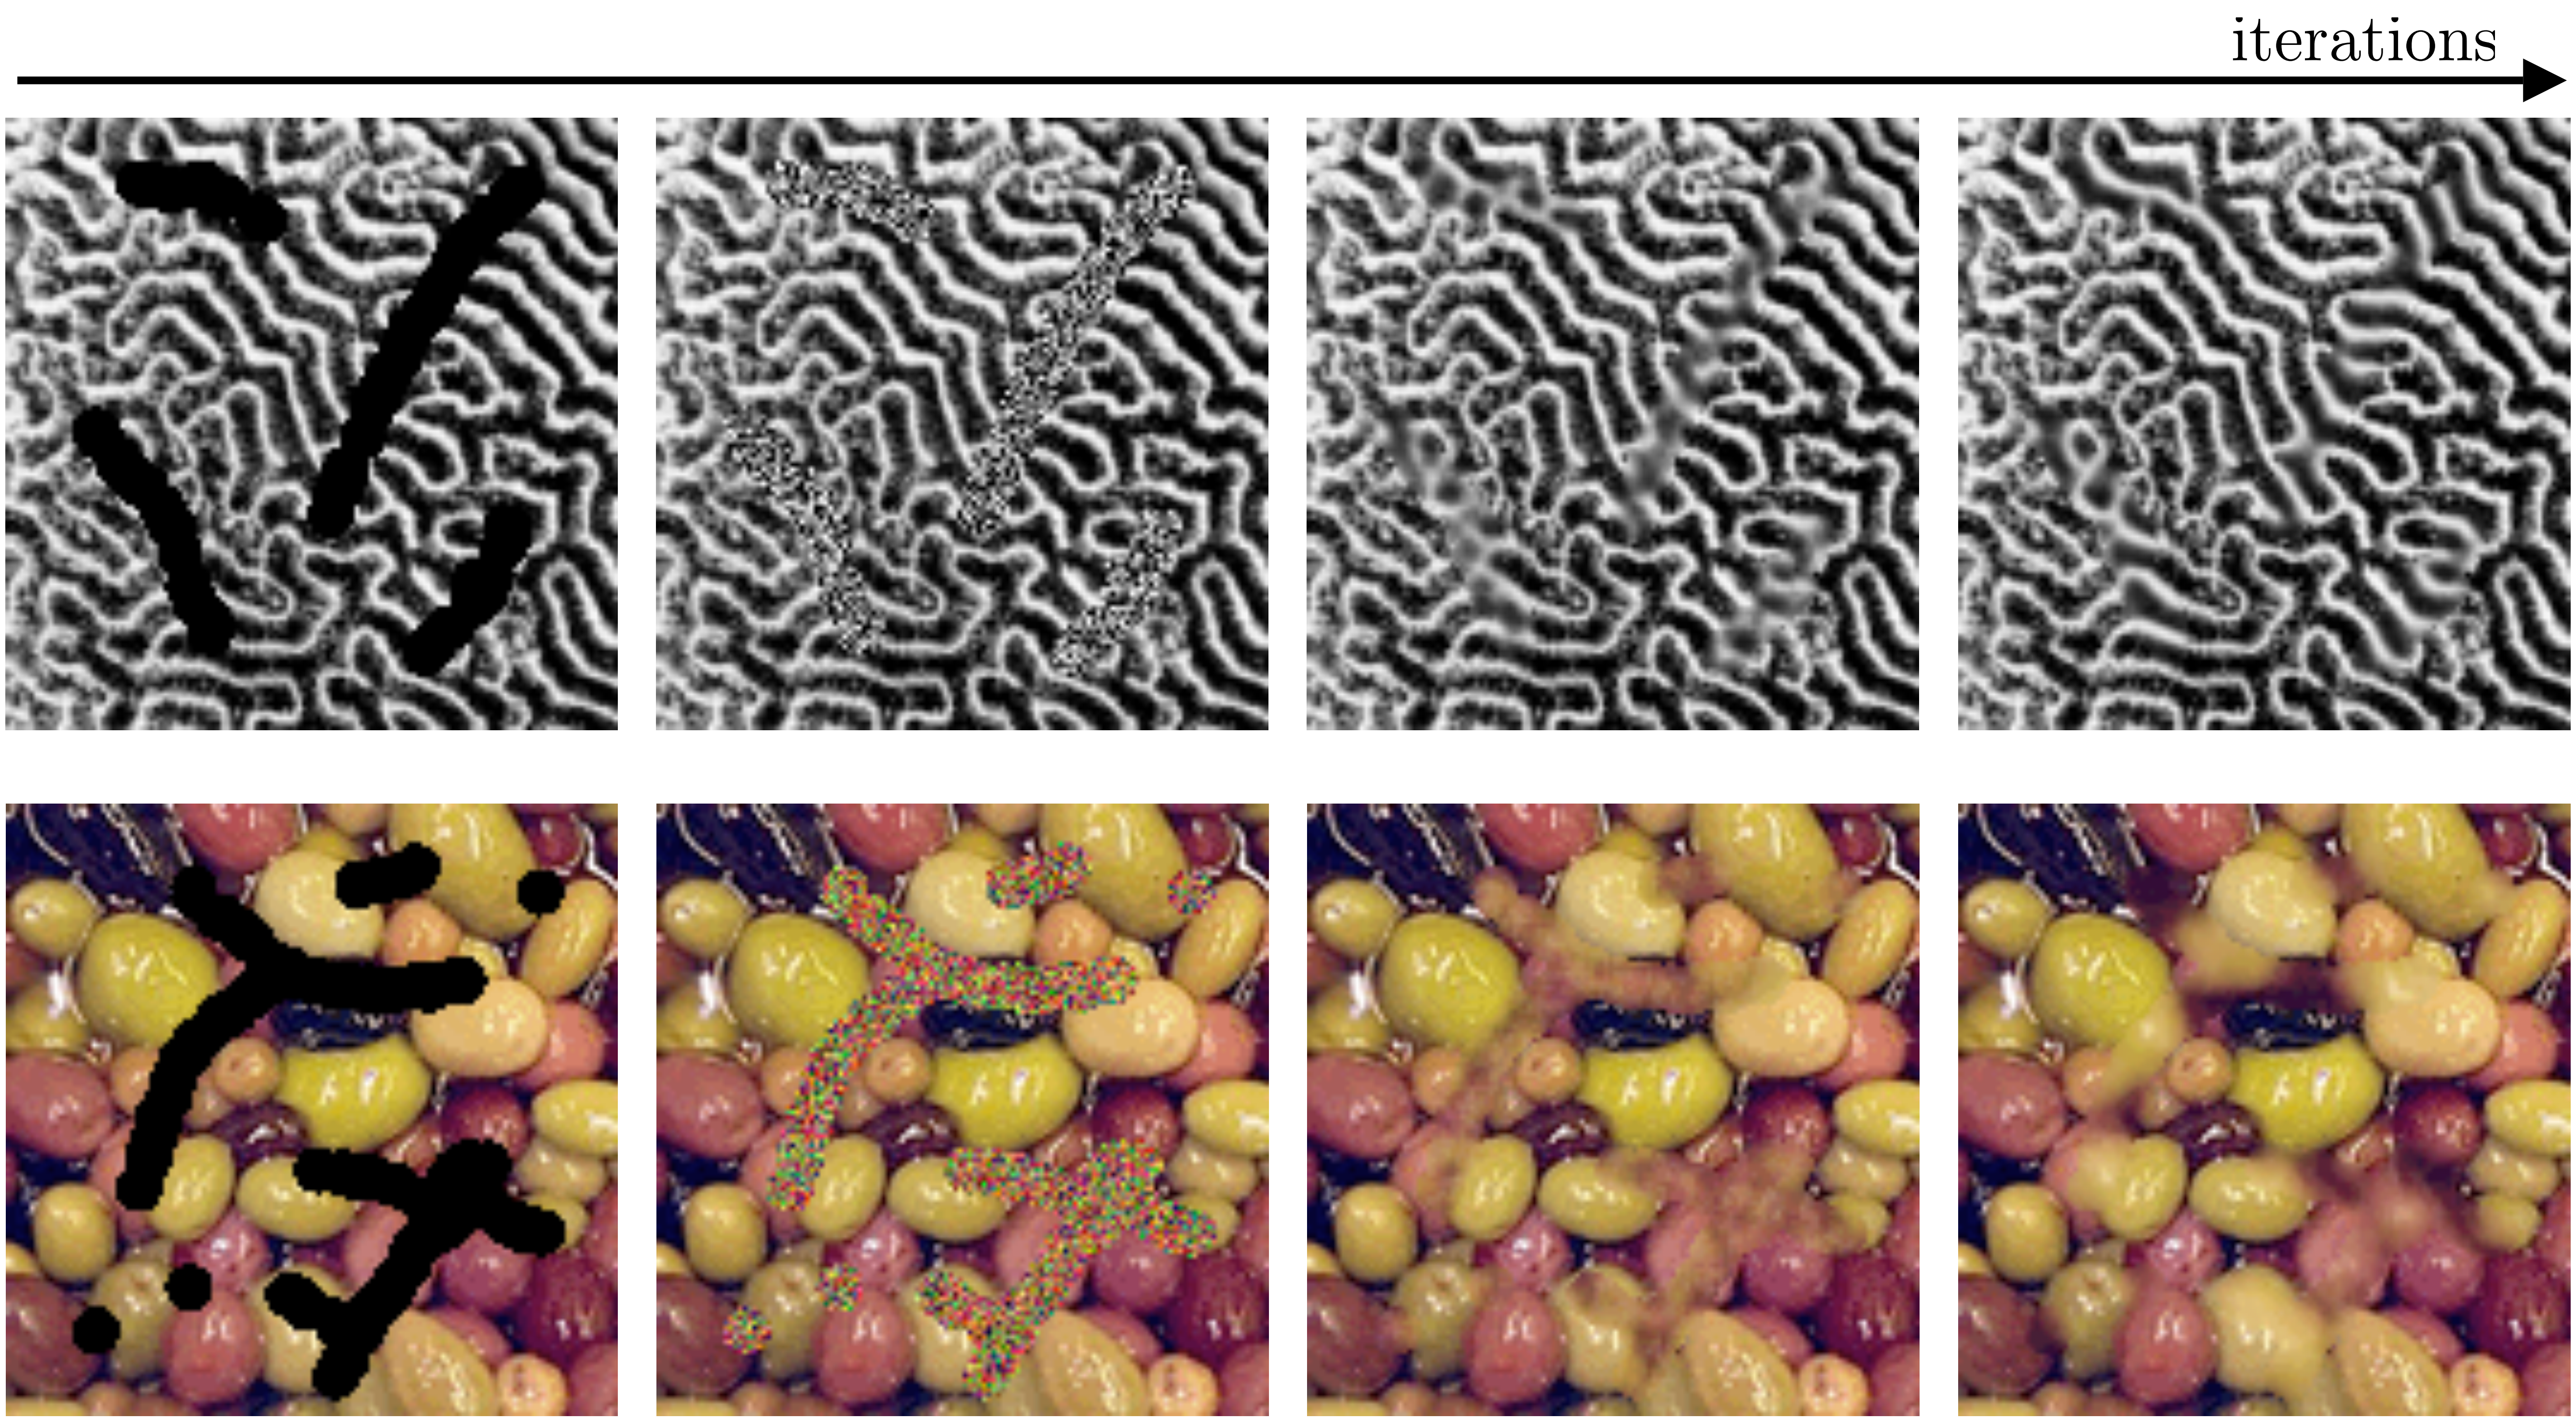

Overview

- Manifolds: Image Libraries vs. Patches
- Examples of Patch Manifolds
- Manifold Energies for Inverse Problems
- Non-adaptive Manifold Models
- Texture Synthesis with Manifold Models
- Adaptive Manifold Models

Graph Representation of Point Clouds

Point cloud $\{p_i\}_i$, each $p_i \in \mathbb{R}^n$.

Example: p_i an image or $p_i = p_{x_i}(f)$ a patch of n pixels.

Weighted graph: $w(i,j) \geqslant 0$ measures "similarity" $i \sim j$.

Example: ε-nearest neighbor graph $\quad w(i,j) = \begin{cases} 1 & \text{if} \quad \|p_i - p_j\| \leqslant \varepsilon, \\ 0 & \text{otherwise.} \end{cases}$

Example: Gaussian kernel $\quad w(i,j) = \exp\left(-\frac{\|p_i - p_j\|^2}{2\varepsilon^2} \right)$

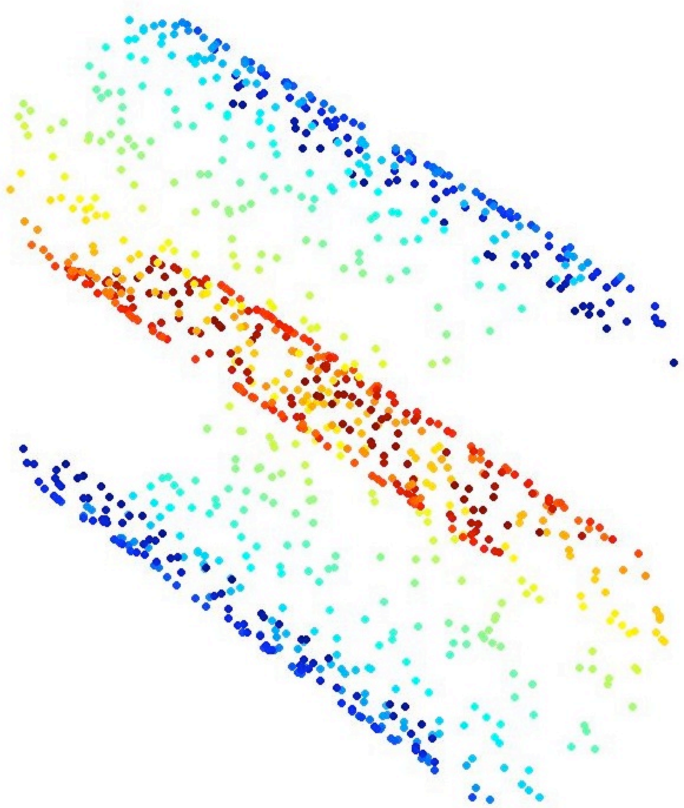

Point cloud $\{p_i\}_i$

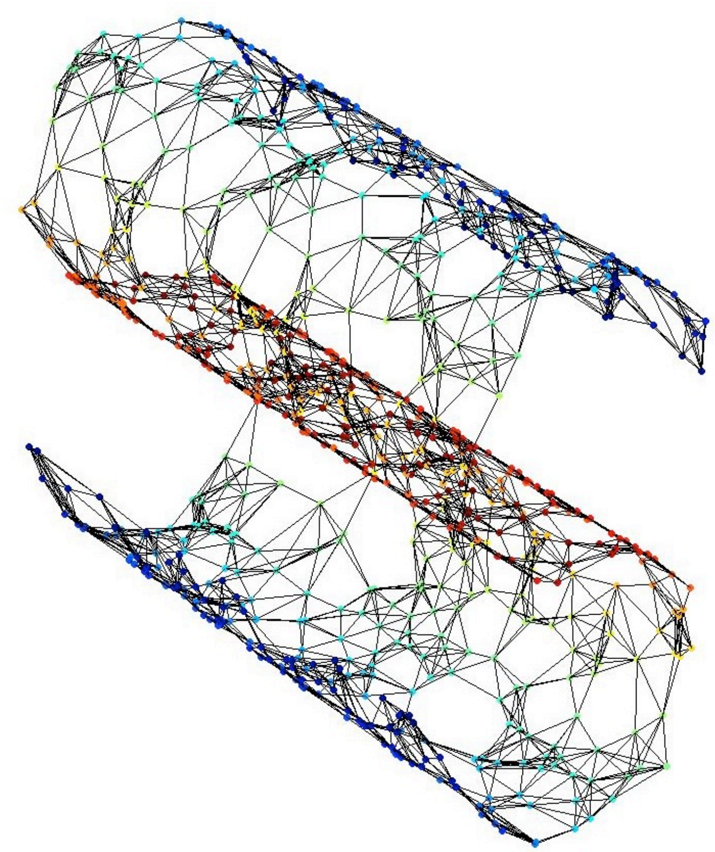

Graph $\{(p_i, p_j)\}_{w(i,j)>0}$

Weights for Image Patches

Weights for a patch manifolds estimated from an image f:

$$w_f(x,y) = w(p_x(f), p_y(f)) = \exp\left(-\frac{\|p_x(f) - p_y(f)\|^2}{2\varepsilon^2} \right)$$

Non-local means [Buades, Coll, Morel, 2005]

Image filtering W_f associated to $w_f(x,y)$

$$W_f g(x) = \frac{1}{Z_x} \sum_y w_f(x,y) g(x) \quad \text{where} \quad Z_x = \sum_y w_f(x,y)$$

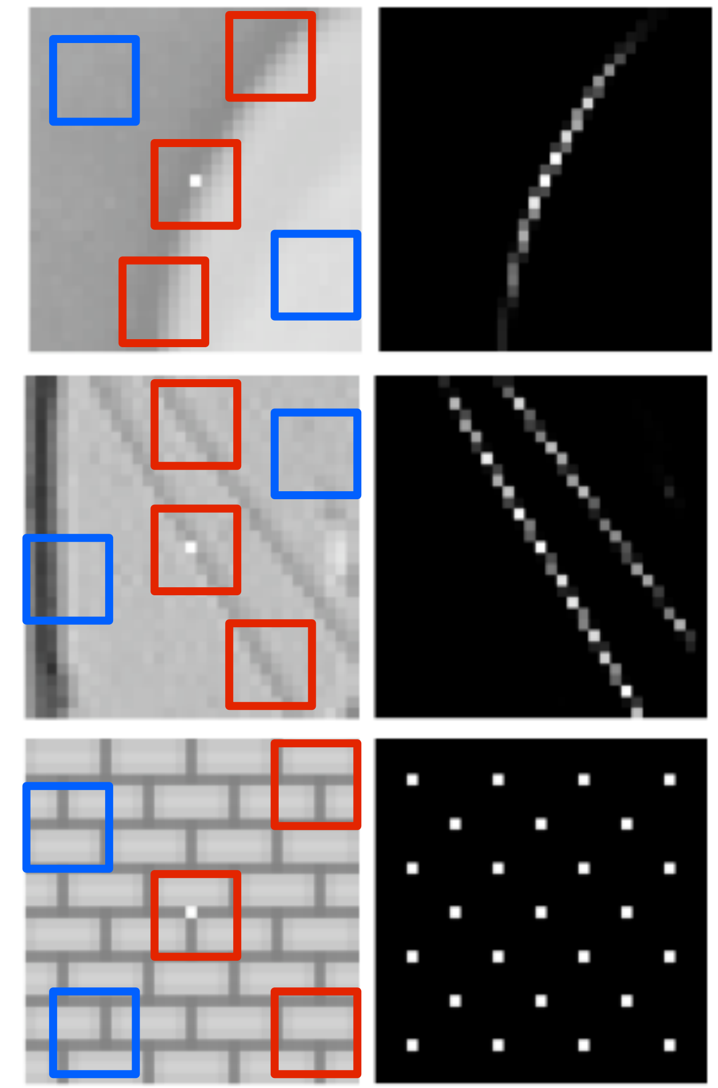

Image f Weights $w(x,\cdot)$

Weights for Image Patches

Weights for a patch manifolds estimated from an image f:

$$w_f(x,y) = w(p_x(f), p_y(f)) = \exp\left(-\frac{\|p_x(f) - p_y(f)\|^2}{2\varepsilon^2} \right)$$

Non-local means [Buades, Coll, Morel, 2005]

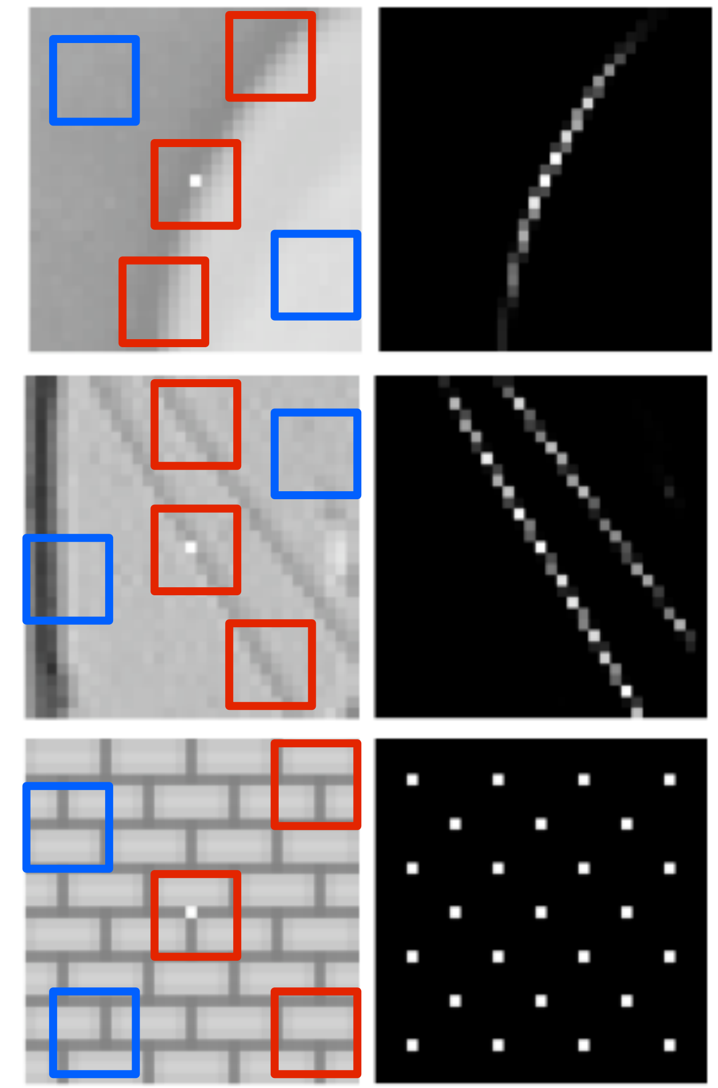

Image f Weights $w(x,\cdot)$

Image filtering W_f associated to $w_f(x,y)$

$$W_f g(x) = \frac{1}{Z_x} \sum_y w_f(x,y) g(x) \quad \text{where} \quad Z_x = \sum_y w_f(x,y)$$

Non-local means: apply W_f to f itself!

$$\tilde{f} = W_f f$$

$\longrightarrow$ adaptive filtering

Noisy f Gaussian blurring   NL-means $\tilde{f}$

Adaptive Manifold Energies

Setting #2: $\mathcal{M} = \mathcal{M}_f = (p_x(f))_x$ is computed from some image f.

Weighted graph $w_f(p_x, p_y) = \exp\left(-\frac{\|p_x - p_y\|^2}{2\varepsilon^2}\right)$

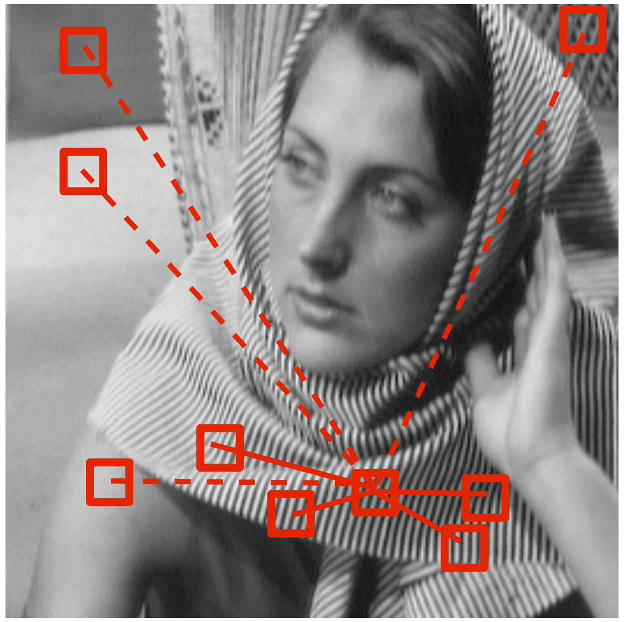

Weight $w_f(x, y)$ on image.

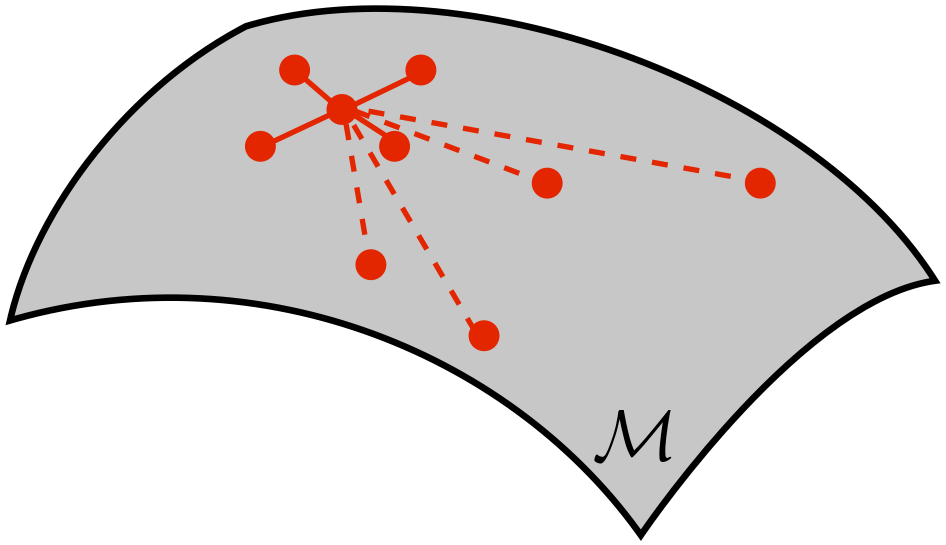

Weight $w_f(p_x, p_y)$ on manifold.

Adaptive Manifold Energies

Setting #2: $\mathcal{M} = \mathcal{M}_f = (p_x(f))_x$ is computed from some image f.

Weighted graph $w_f(p_x, p_y) = \exp\left(-\frac{\|p_x - p_y\|^2}{2\varepsilon^2}\right)$

Weight $w_f(x, y)$ on image.

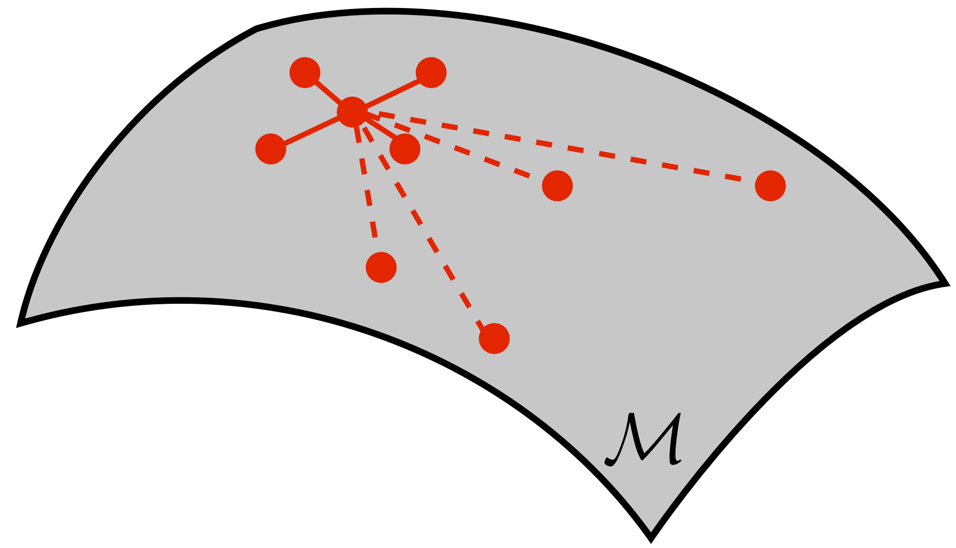

Weight $w_f(p_x, p_y)$ on manifold.

Manifold Sobolev energy: $J_w^{\text{sob}}(g) = \sum_{x,y} w_f(x,y)|g(x) - g(y)|^2$.

Manifold TV energy: $J_w^{\text{tv}}(g) = \sum_{x,y} w_f(x,y)|g(x) - g(y)|$.

$\forall\,(x,y), \quad g(x) \approx g(y)$ for points $(p_x(f), p_y(f))$ close on the manifold $\mathcal{M}_f$.

Adaptive Manifold Energies

Setting #2: $\mathcal{M} = \mathcal{M}_f = (p_x(f))_x$ is computed from some image f.

Weighted graph $w_f(p_x, p_y) = \exp\left(-\frac{\|p_x - p_y\|^2}{2\varepsilon^2}\right)$

Weight $w_f(x, y)$ on image.

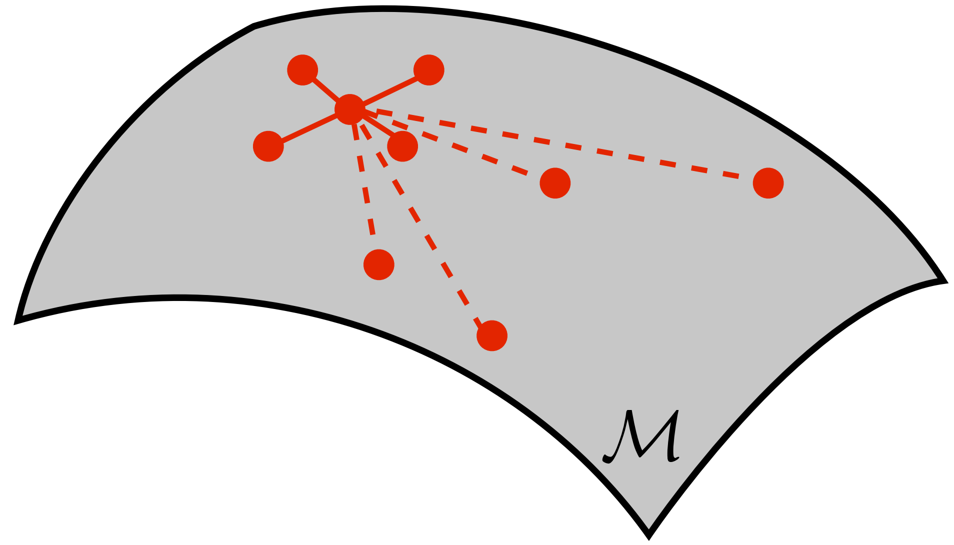

Weight $w_f(p_x, p_y)$ on manifold.

Manifold Sobolev energy: $J_w^{\text{sob}}(g) = \sum_{x,y} w_f(x,y)|g(x) - g(y)|^2$.

Manifold TV energy: $J_w^{\text{tv}}(g) = \sum_{x,y} w_f(x,y)|g(x) - g(y)|$.

$\forall\,(x, y), \quad g(x) \approx g(y)$ for points $(p_x(f), p_y(f))$ close on the manifold $\mathcal{M}_f$.

Optimize w to the geometry of the solution.

$\longrightarrow$ denoising: easy, adapt w to the noisy observation f+noise.

[Coifman, Lafon et al. 2005] [Gilboa et al. 2007] $\cdots$

$\longrightarrow$ inverse problems: difficult, needs to find both w and $f^\star$.

Adaptive Manifold Regularization

Find both solution $f^\star$ and adapted weights $w^\star$:

$$(f^\star, w^\star) = \underset{(g,w)}{\text{argmin}} \; \frac{1}{2}\|y - \Phi g\|^2 + \lambda J_w(g)$$

Iterative minimization algorithm for $J_w = J_w^{\text{sob}}$:

Step 1: $w^\star$ fixed, gradient descent with step τ

$$f^\star \leftarrow f^\star + \tau \Phi^*(\Phi f^\star - y) - \tau\lambda\Delta^{w^\star} f^\star$$

Step 2: $f^\star$ fixed, estimate the graph $w^\star$

$$w^\star(x, y) \leftarrow \exp\left(-\frac{\|p_x(f^\star) - p_y(f^\star)\|^2}{2\varepsilon^2}\right)$$

For non-smooth $J_w = J_w^{\text{tv}}$ replace gradient descent by proximal iterations.

See [Peyré, Bougleux, Cohen, ECCV'08]

Inpainting Results

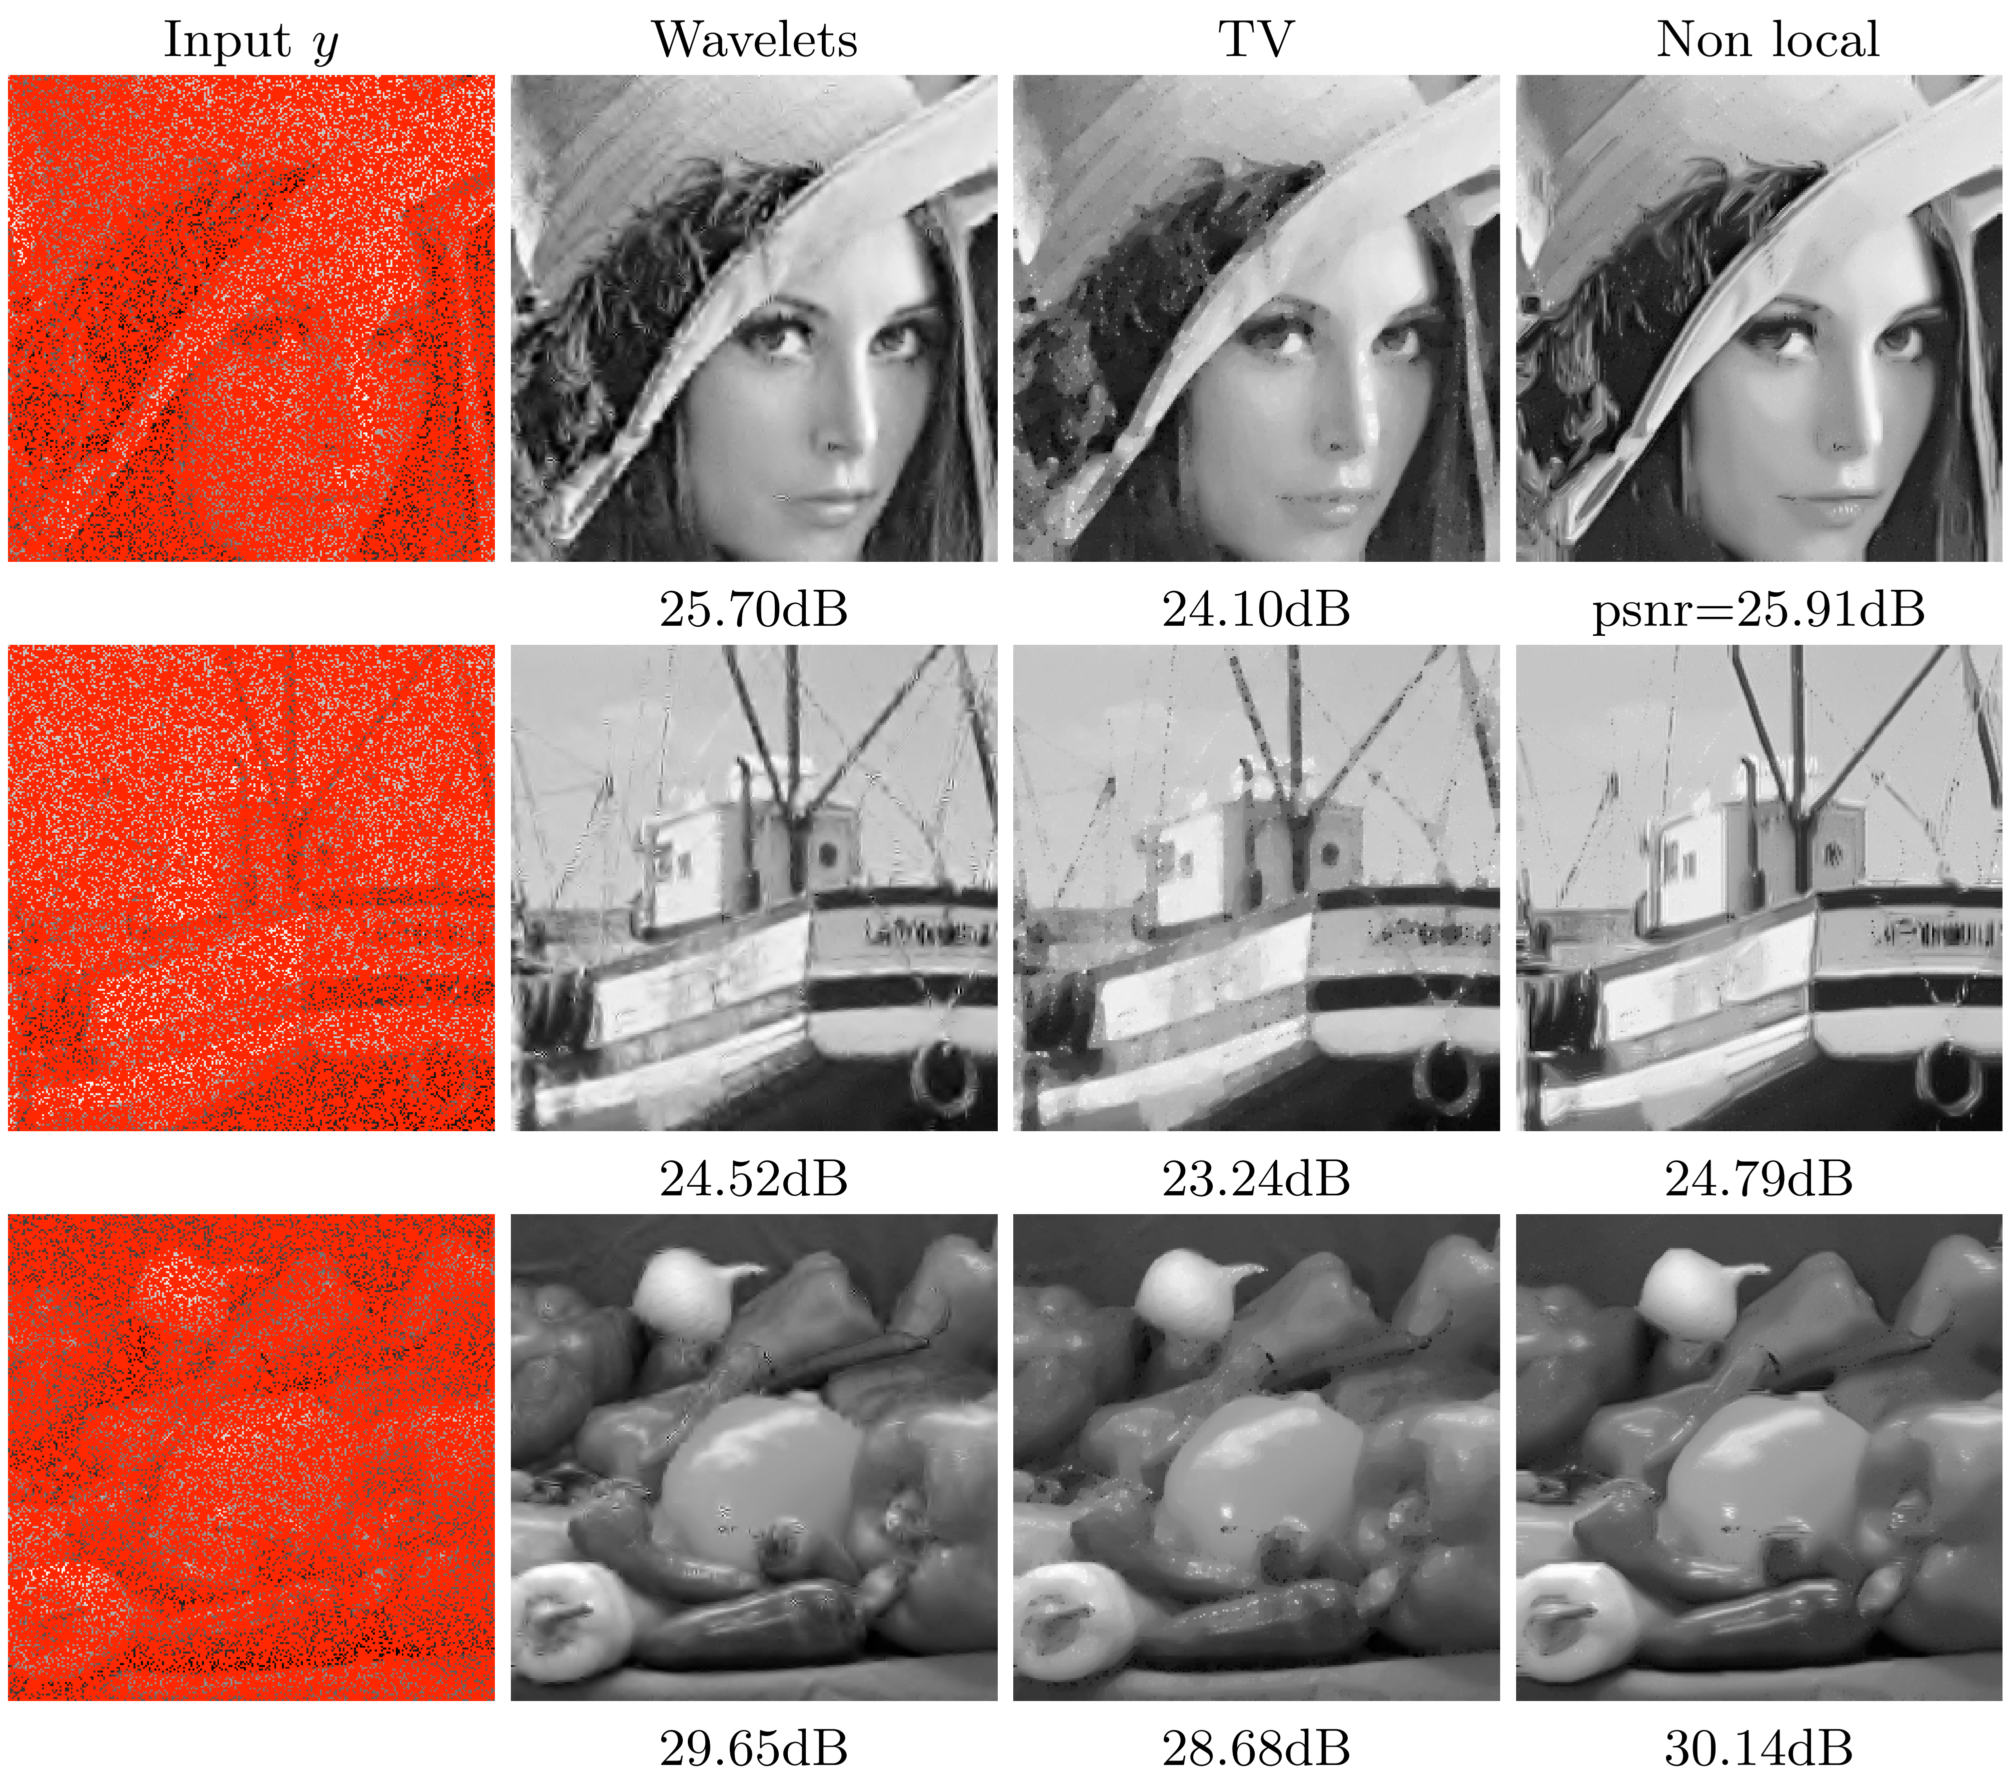

Super-resolution Results

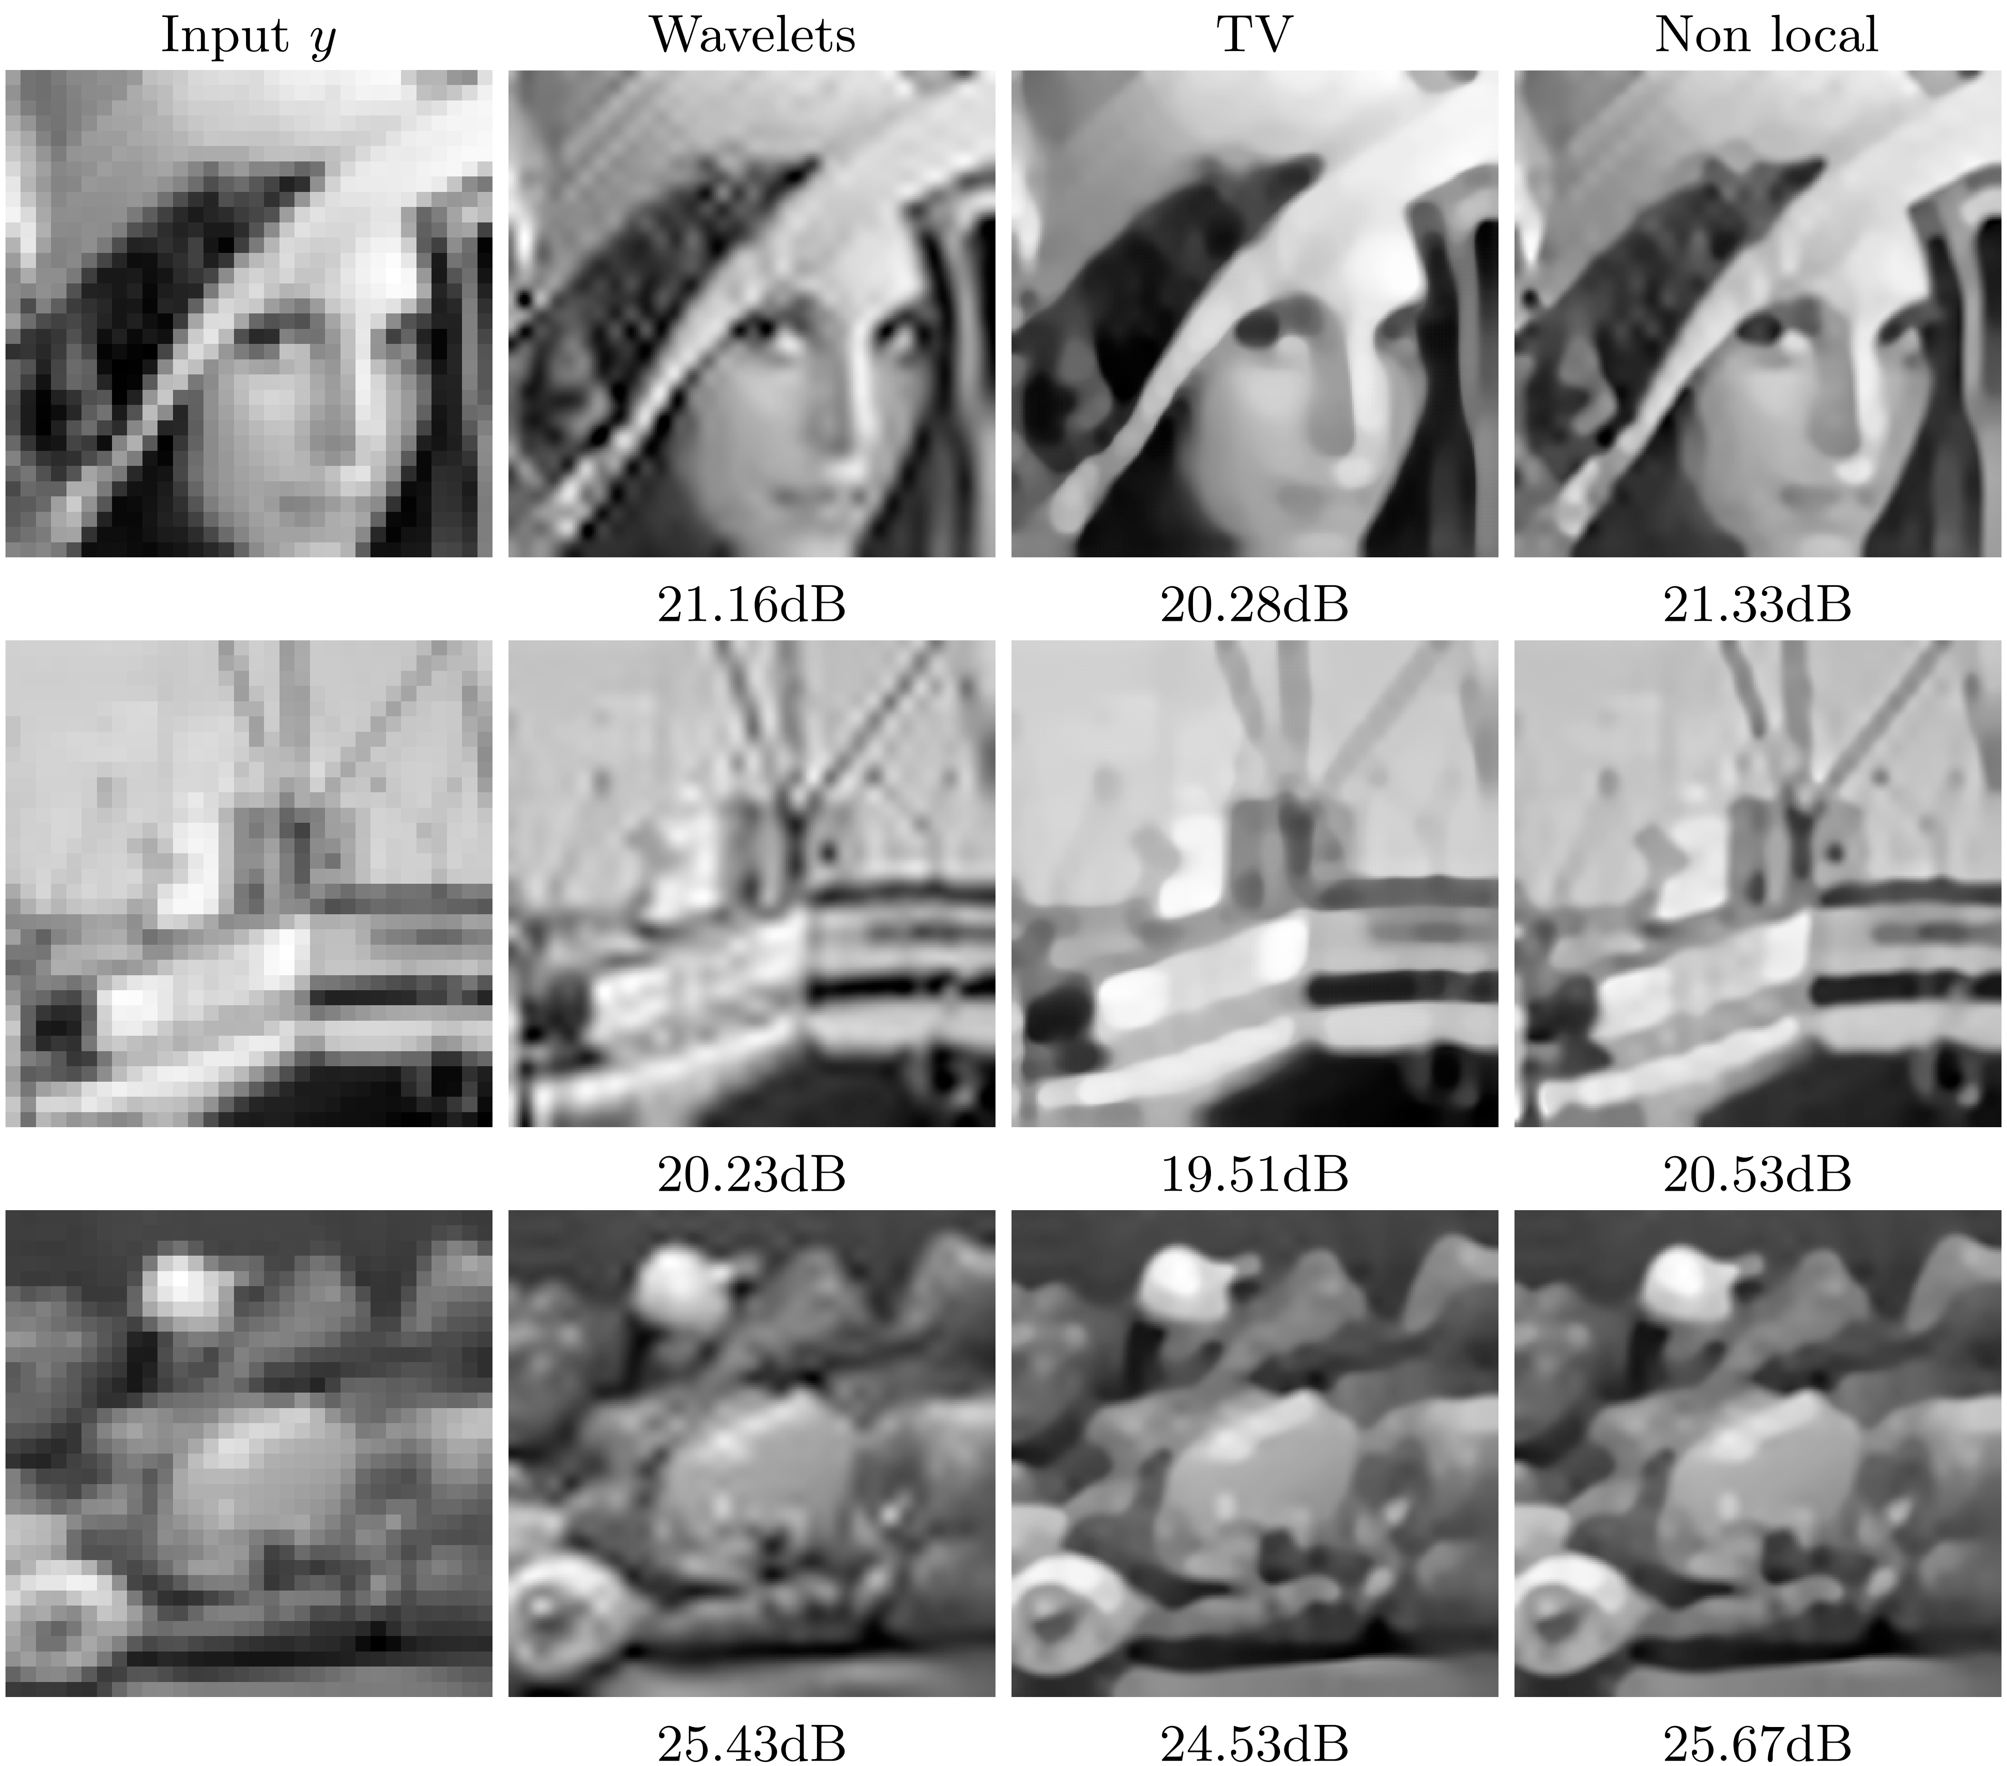

Compressed Sensing Results

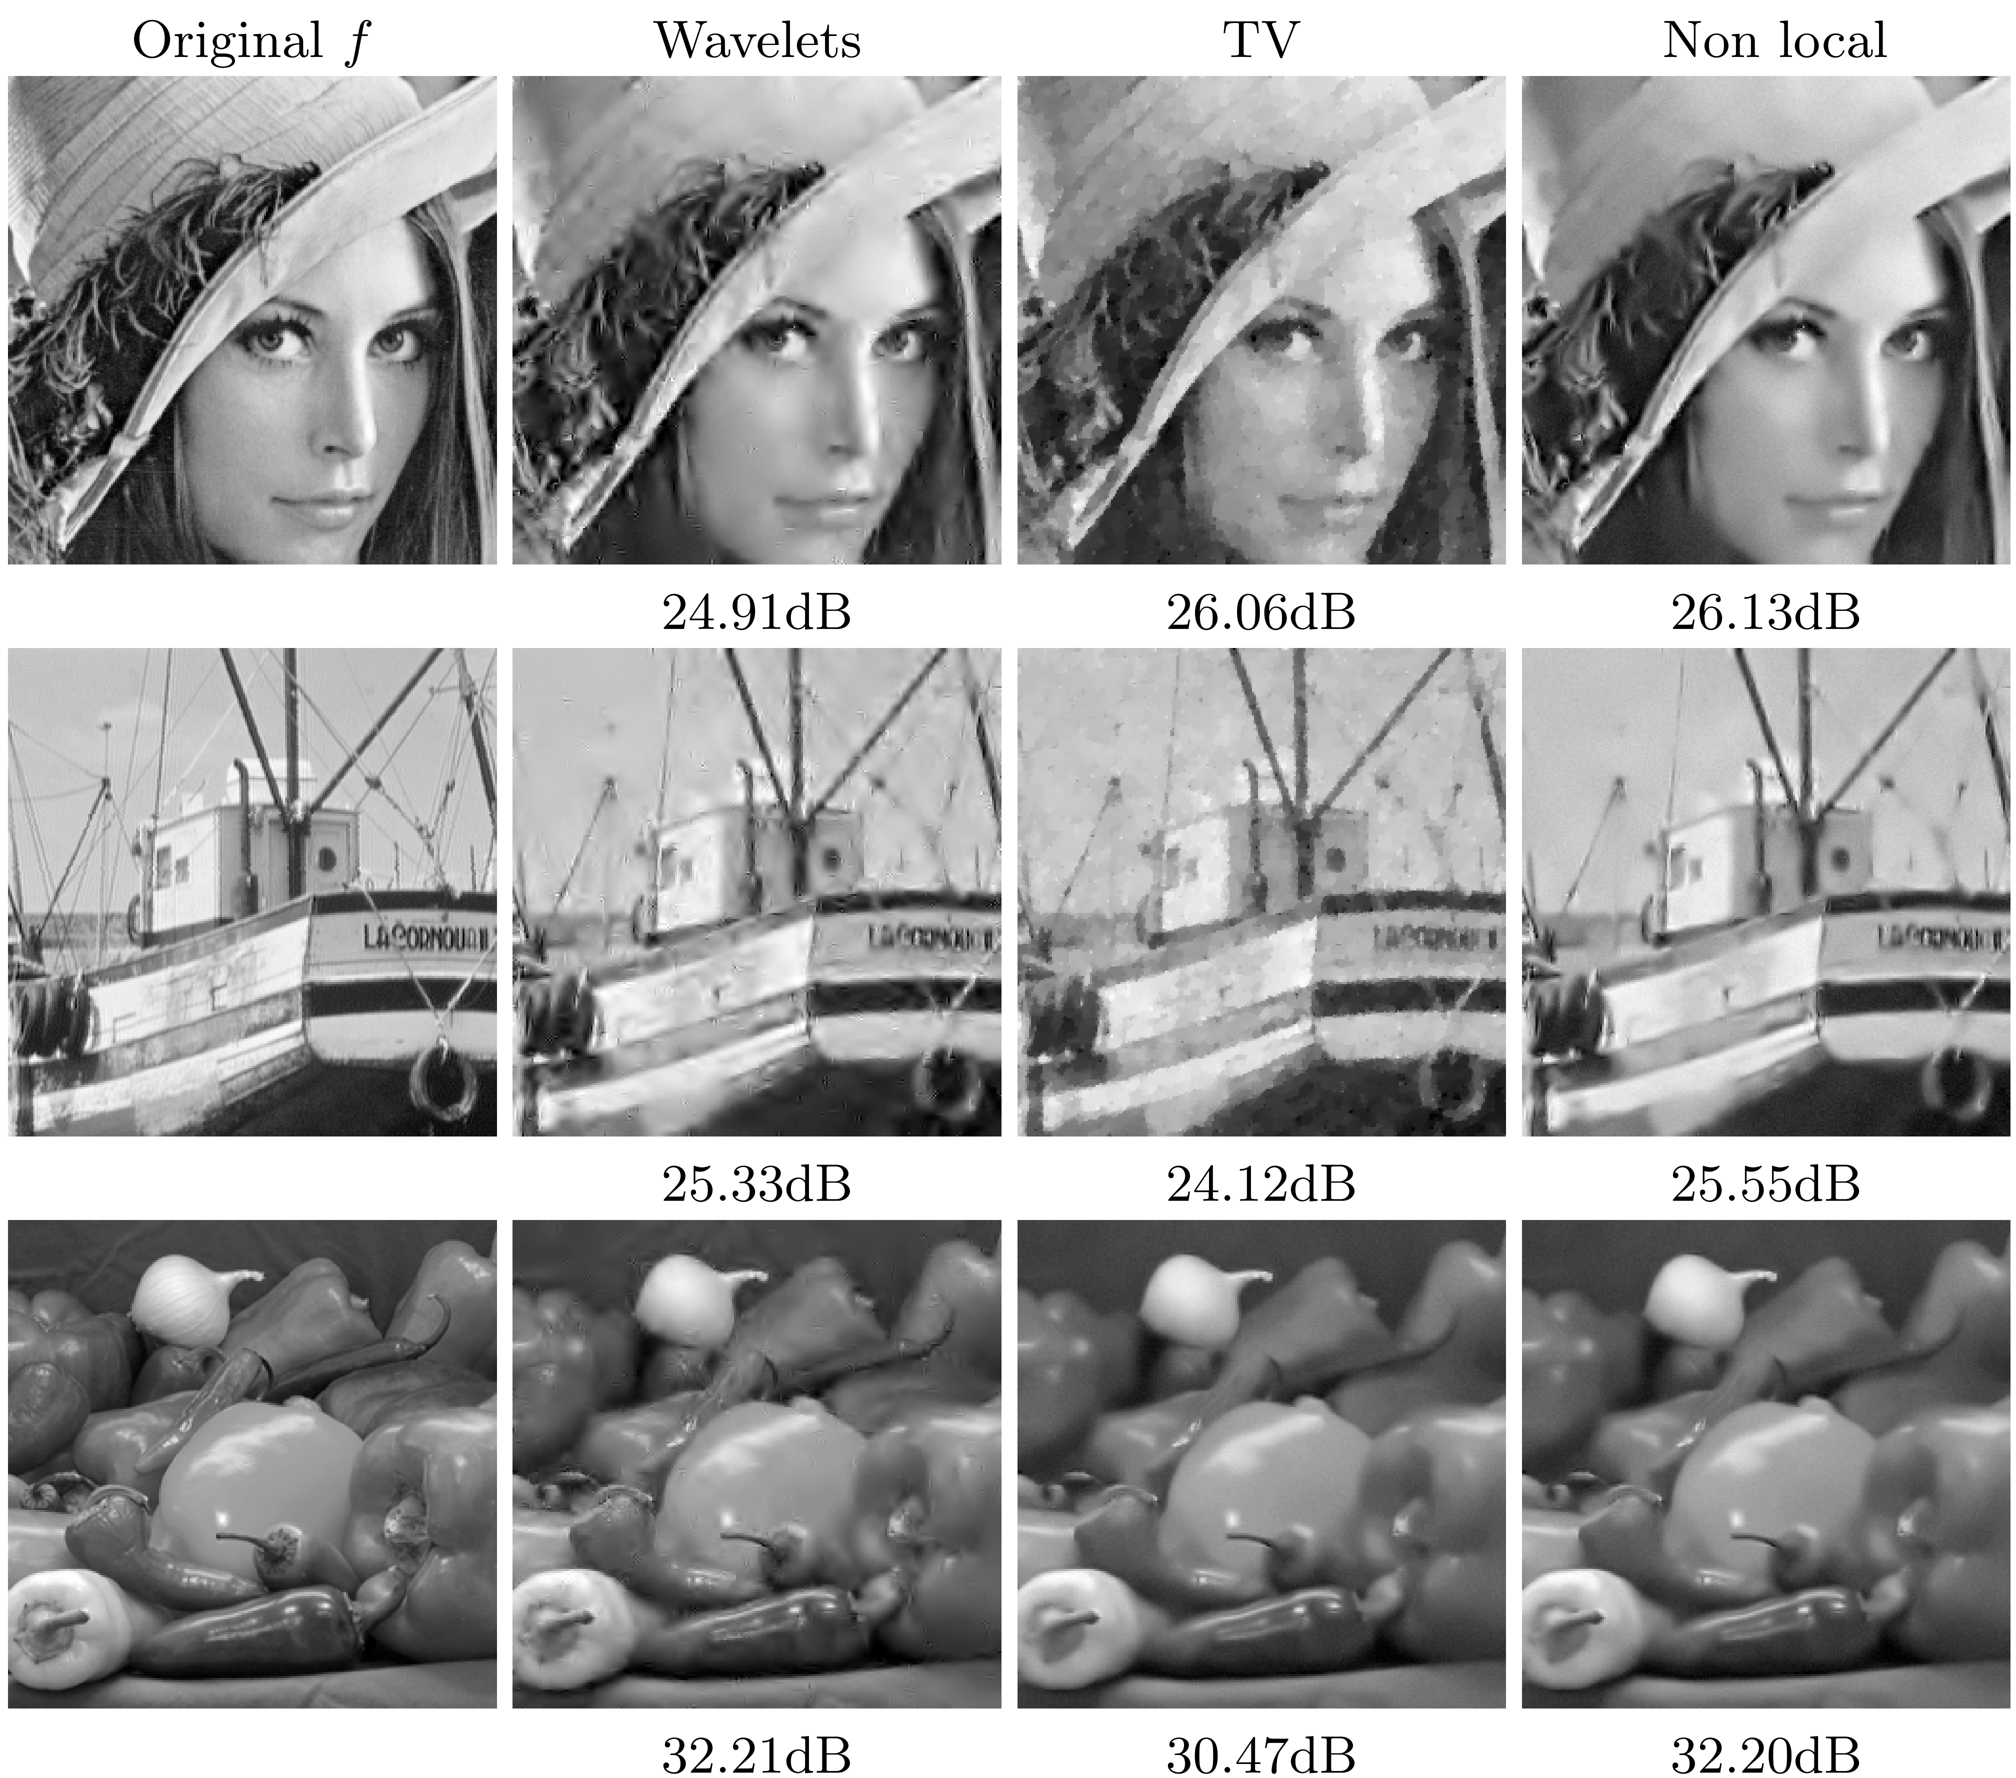

Conclusion

The local geometry of images can sometimes be captured by a manifold $\mathcal{M}$.

$\longrightarrow$ low dimensional parameterization of the features.

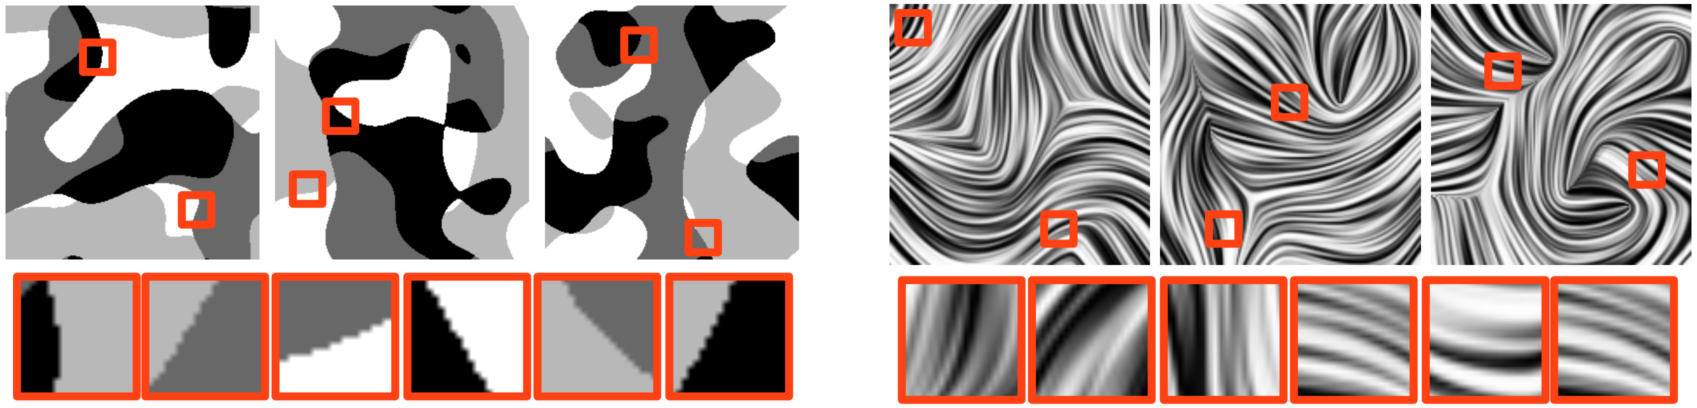

Conclusion

The local geometry of images can sometimes be captured by a manifold $\mathcal{M}$.

$\longrightarrow$ low dimensional parameterization of the features.

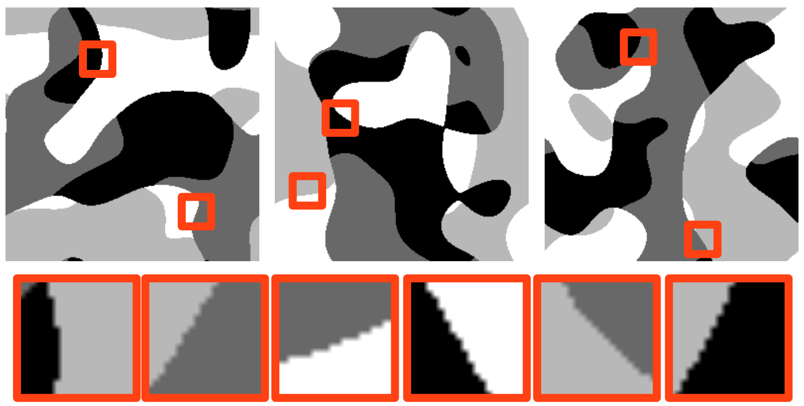

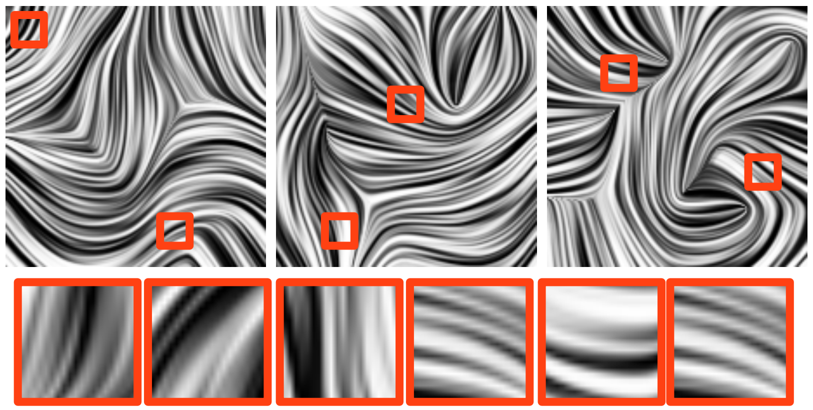

For complex images, the manifold can be learned from the data.

$\longrightarrow$ computing non-local connexions between pixels.

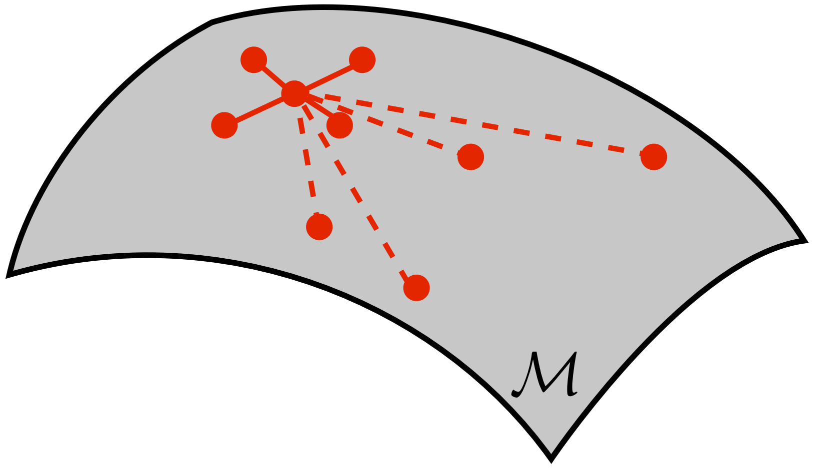

Conclusion

The local geometry of images can sometimes be captured by a manifold $\mathcal{M}$.

$\longrightarrow$ low dimensional parameterization of the features.

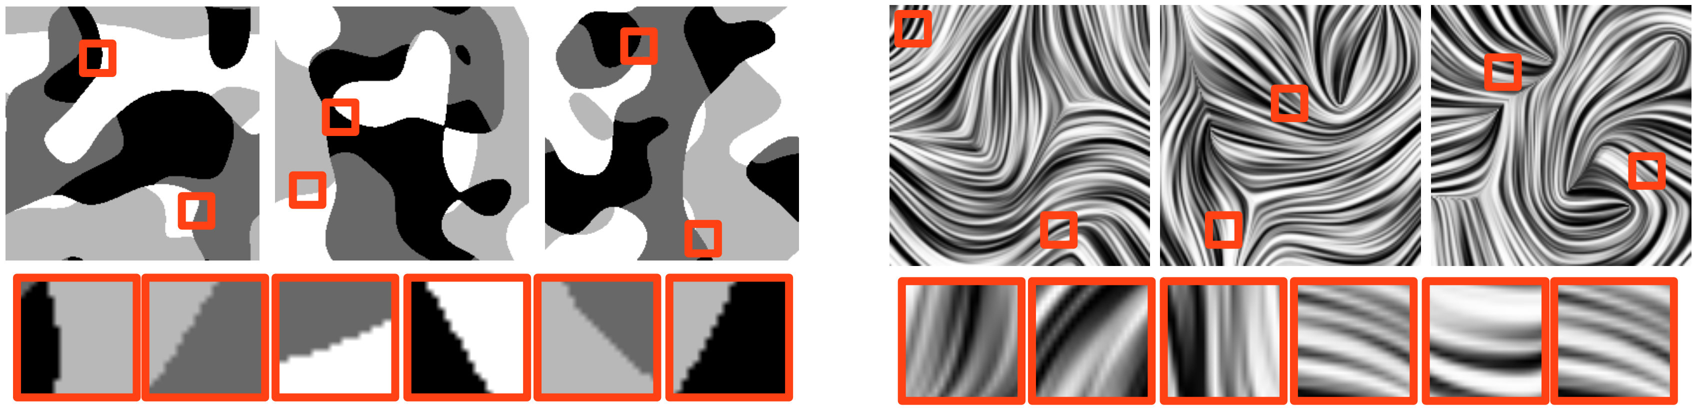

For complex images, the manifold can be learned from the data.

$\longrightarrow$ computing non-local connexions between pixels.

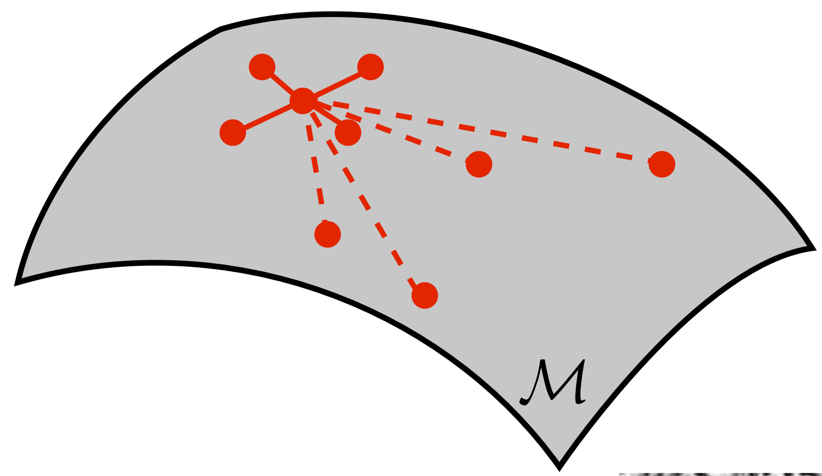

Inverse problem resolution: energy design and minimization.

$\longrightarrow$ fixed manifold $\mathcal{M}$: iterative projection.

$\longrightarrow$ adaptive manifold $\mathcal{M}_w$: optimizing the connexions w.

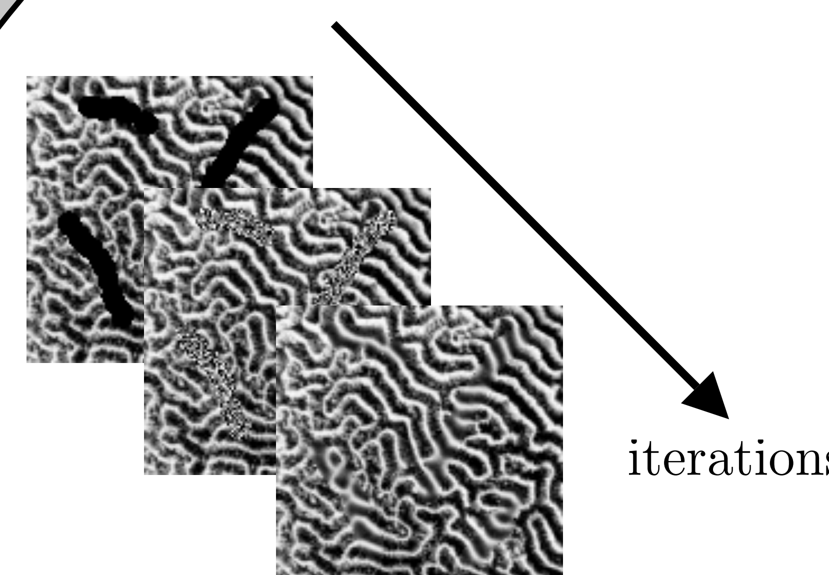

Additional Slides

Geodesics and Dimension Reduction

Hypothesis: existance of a global parameterization $\varphi : \Omega \subset \mathbb{R}^m \rightarrow \mathcal{M}$.

Isomap algorithm [Tenenbaum, de Silva, Langford, 2000]:

Find locations $\tilde{x}_i \in \mathbb{R}^d$ such that $\|\tilde{x}_i - \tilde{x}_j\| \approx d_{\mathcal{M}}(x_i, x_j)$.

Computation of geodesic distance $d_{\mathcal{M}}$: Dijkstra on the ε NN-graph w.

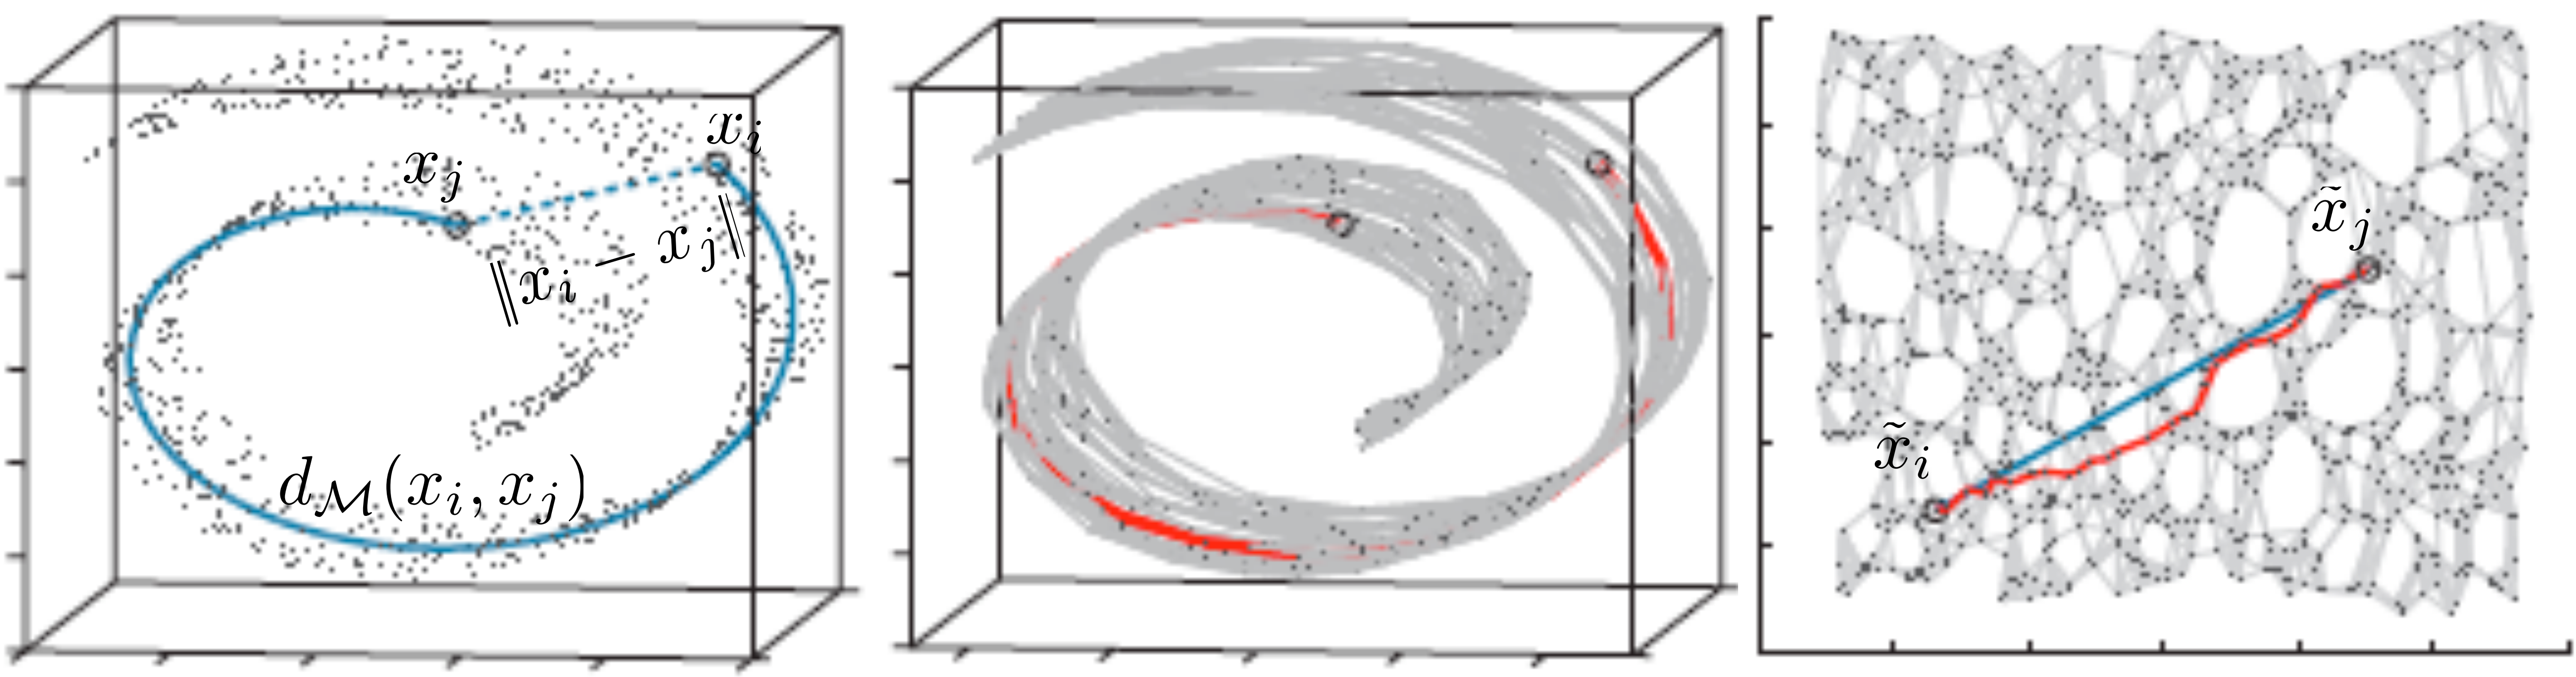

If $\mathcal{M}$ is isometric to Euclidean space, Isomap finds a valid parameterization.

Other methods: LLE, HLLE, Laplacian eigenmaps, geometric harmonics, $\ldots$.

Parameterization of Image Datasets

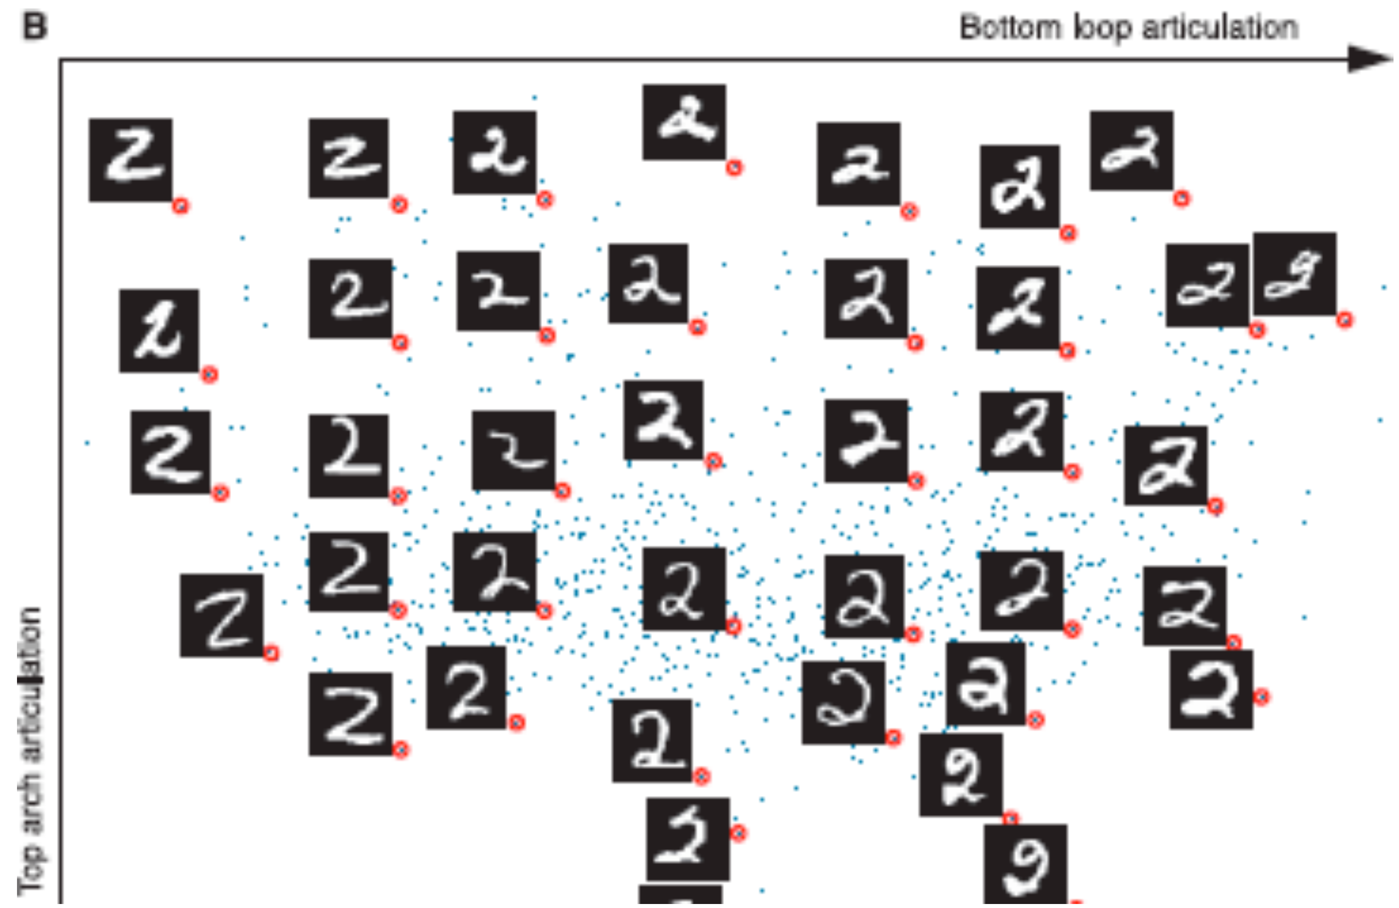

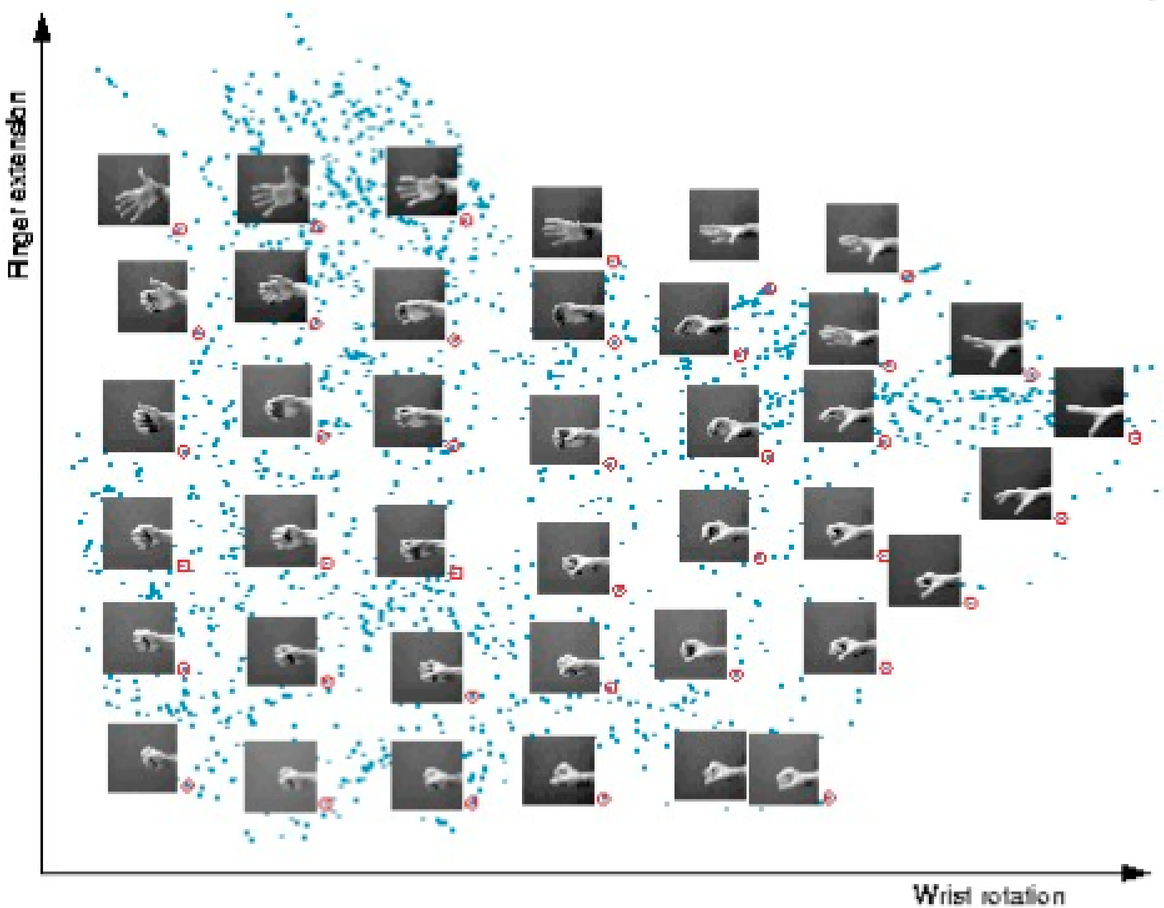

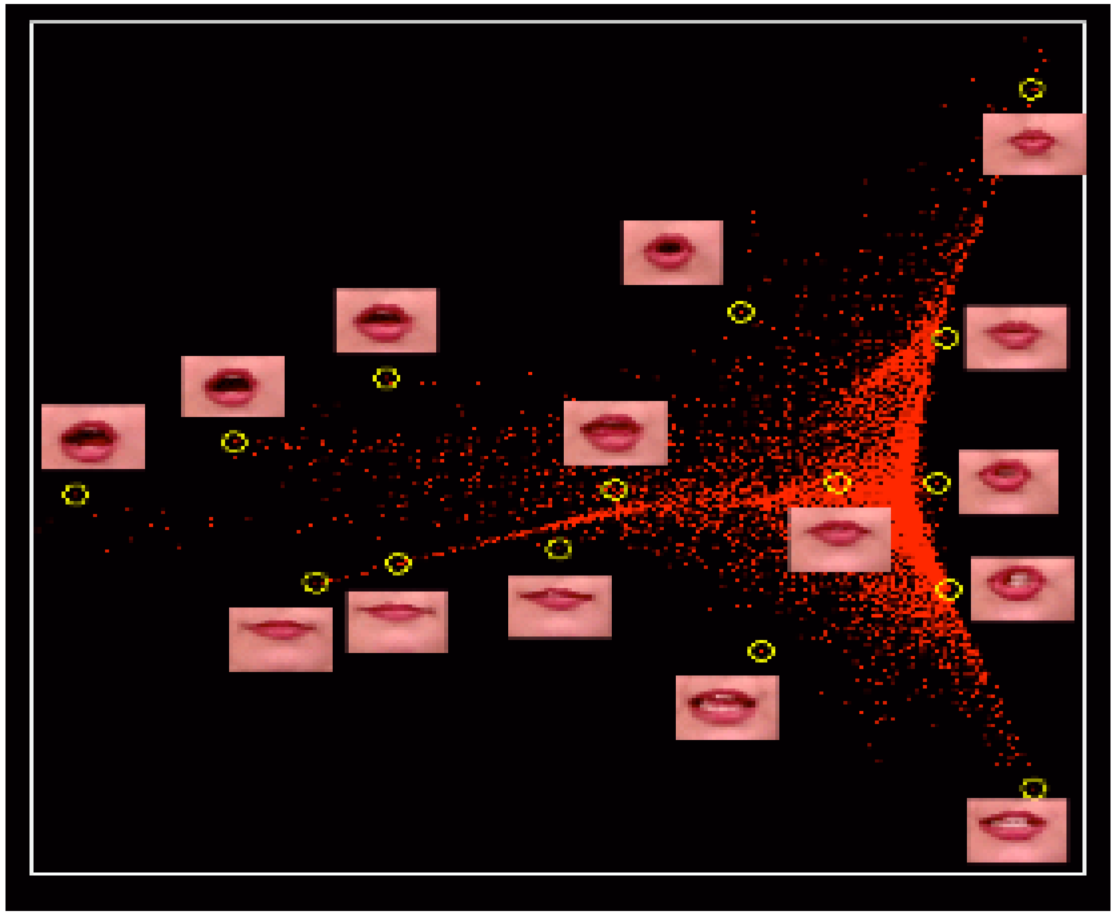

Differential Operators and Energies

Manifold Sobolev energy: $J_w^{\text{sob}}(g) = \sum_{x,y} w_f(x,y)|g(x)-g(y)|^2 = \langle g,\, \Delta^w g\rangle$.

Laplacian: $\Delta^w g(x) = \left(\sum_y w_f(x,y)\right) g(x) - \left(\sum_y w_f(x,y) g(y)\right)$

Gradient descent: non-local heat equation $\dfrac{\partial^2 g_t}{\partial t^2} = -\Delta^{w_f} g_t$ and $g_0 = g$

Denoise by heat diffusion $t \mapsto f_t$ with weights w_f and $f_0 = f$.

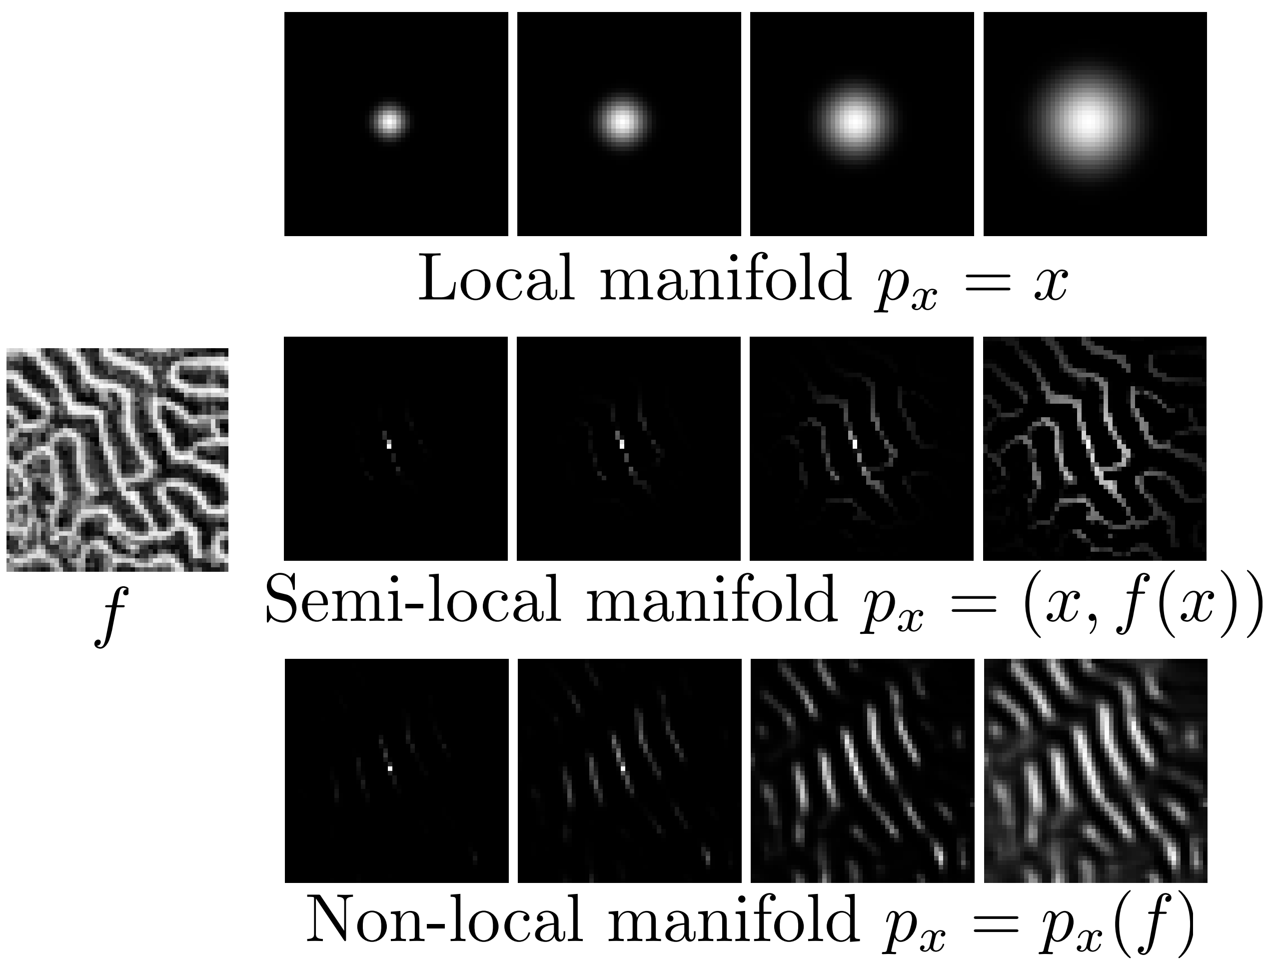

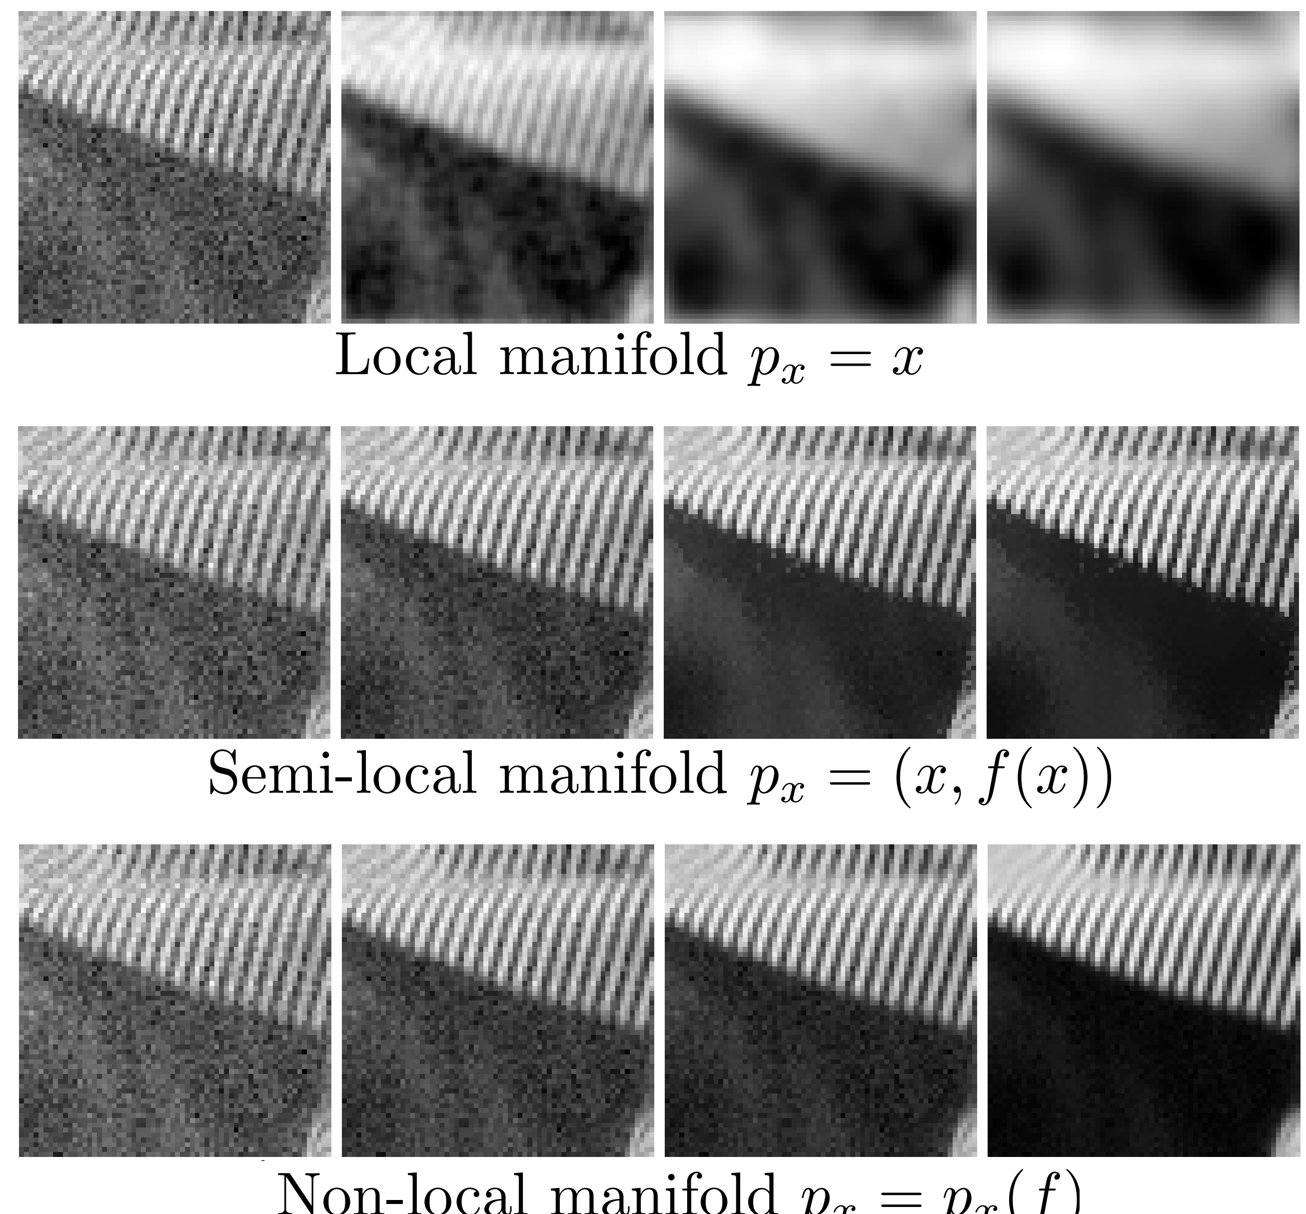

Manifold Spectral Basis

Eigenvectors of the Laplacian Δ^w: $\mathcal{B}(w) = \{\psi_j^w\}_j$ ortho-basis of $\mathbb{R}^n$.

$$\Delta^w \psi_j^w = \lambda_j \psi_k^w \qquad \lambda_j \simeq \text{frequency.}$$

$$J_w^{\text{sob}}(g) = \langle g,\, \Delta^w g \rangle = \sum_j \lambda_j |\langle f,\, \psi_j^w \rangle|^2 \qquad\qquad J_w^{\text{spars}}(g) = \sum_j |\langle f,\, \psi_j^w \rangle|$$

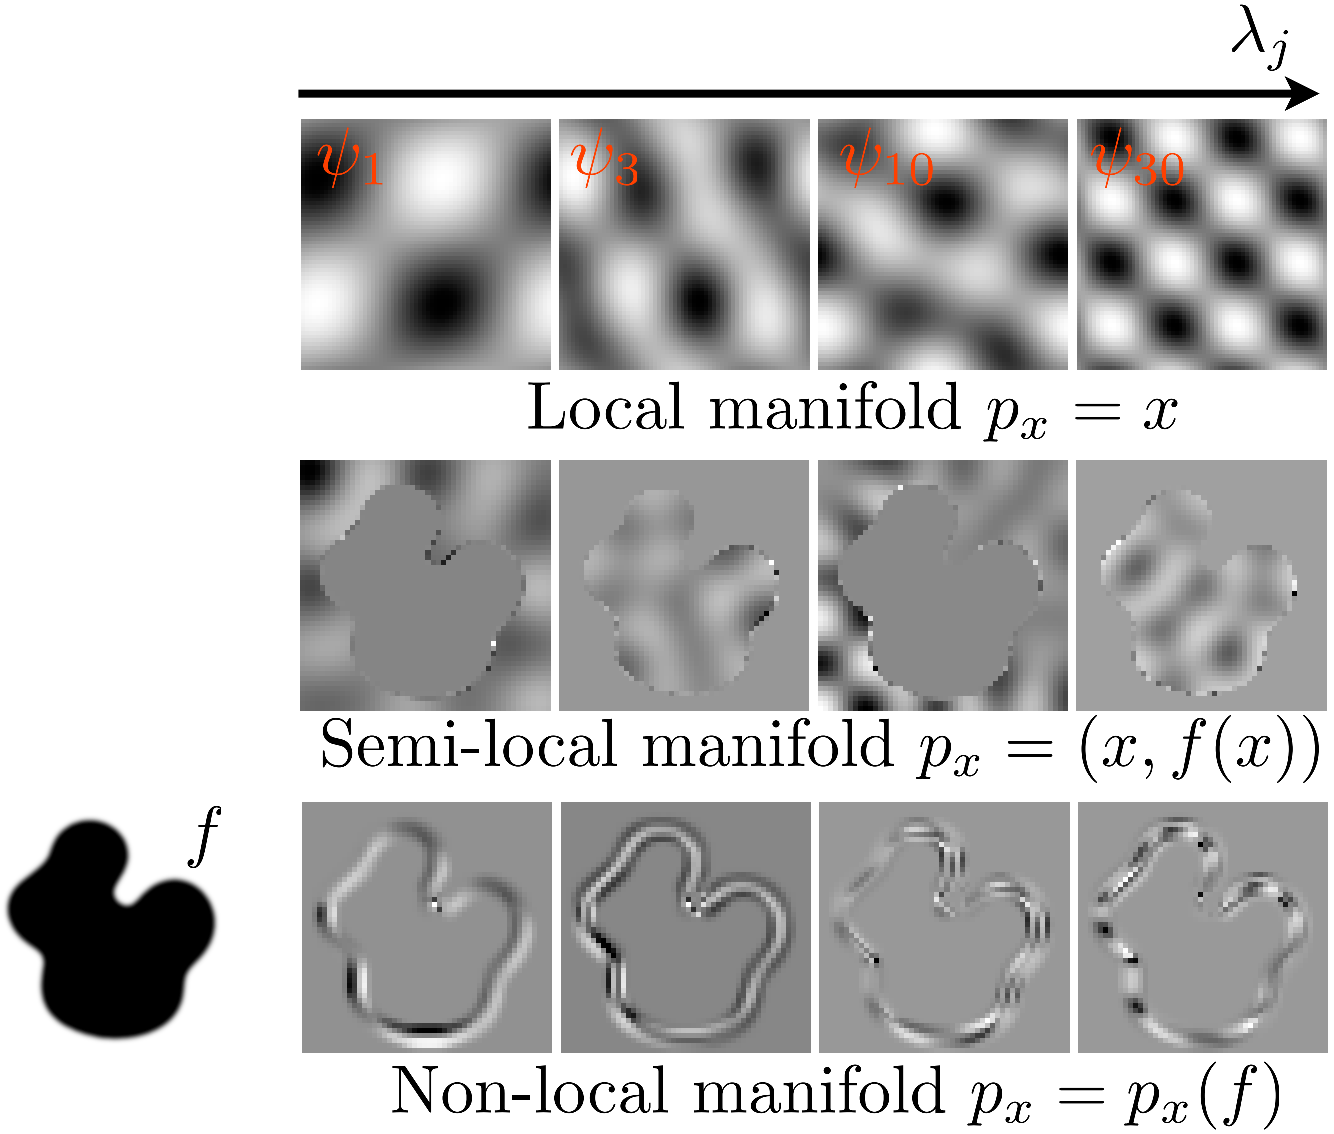

Manifold Spectral Basis

Eigenvectors of the Laplacian Δ^w: $\mathcal{B}(w) = \{\psi_j^w\}_j$ ortho-basis of $\mathbb{R}^n$.

$$\Delta^w \psi_j^w = \lambda_j \psi_k^w \qquad \lambda_j \simeq \text{frequency.}$$

$$J_w^{\text{sob}}(g) = \langle g,\, \Delta^w g\rangle = \sum_j \lambda_j |\langle f,\, \psi_j^w\rangle|^2 \qquad\qquad J_w^{\text{spars}}(g) = \sum_j |\langle f,\, \psi_j^w\rangle|$$

$$\underset{g}{\text{argmin}}\; \|f-g\|^2 + \lambda J_w^{\text{sob}}(g) = \sum_j \frac{\langle f,\, \psi_j^w\rangle}{1+\lambda\lambda_j}\psi_j^w$$

$$\underset{g}{\text{argmin}}\; \frac{1}{2}\|f-g\|^2 + \lambda J_w^{\text{spars}}(g) = \sum_j S_\lambda(\langle f,\, \psi_j^w\rangle)\psi_j^w$$

Soft thresholding operator

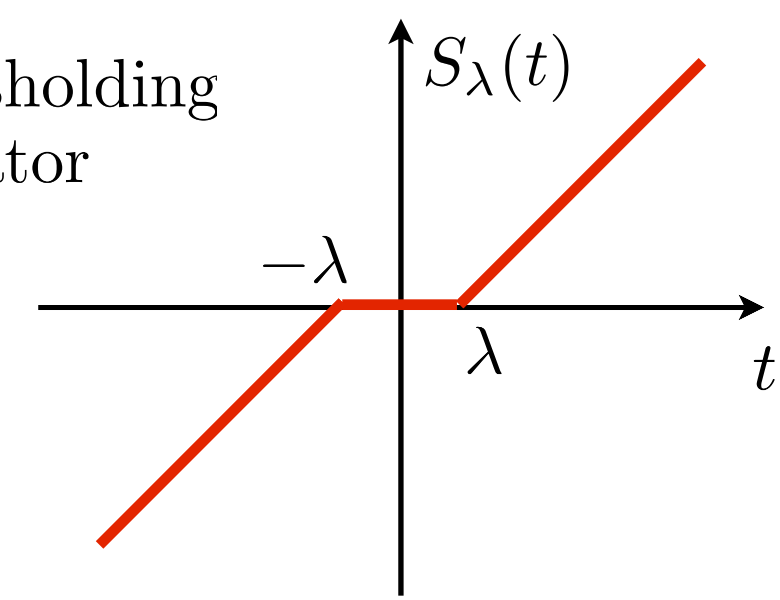

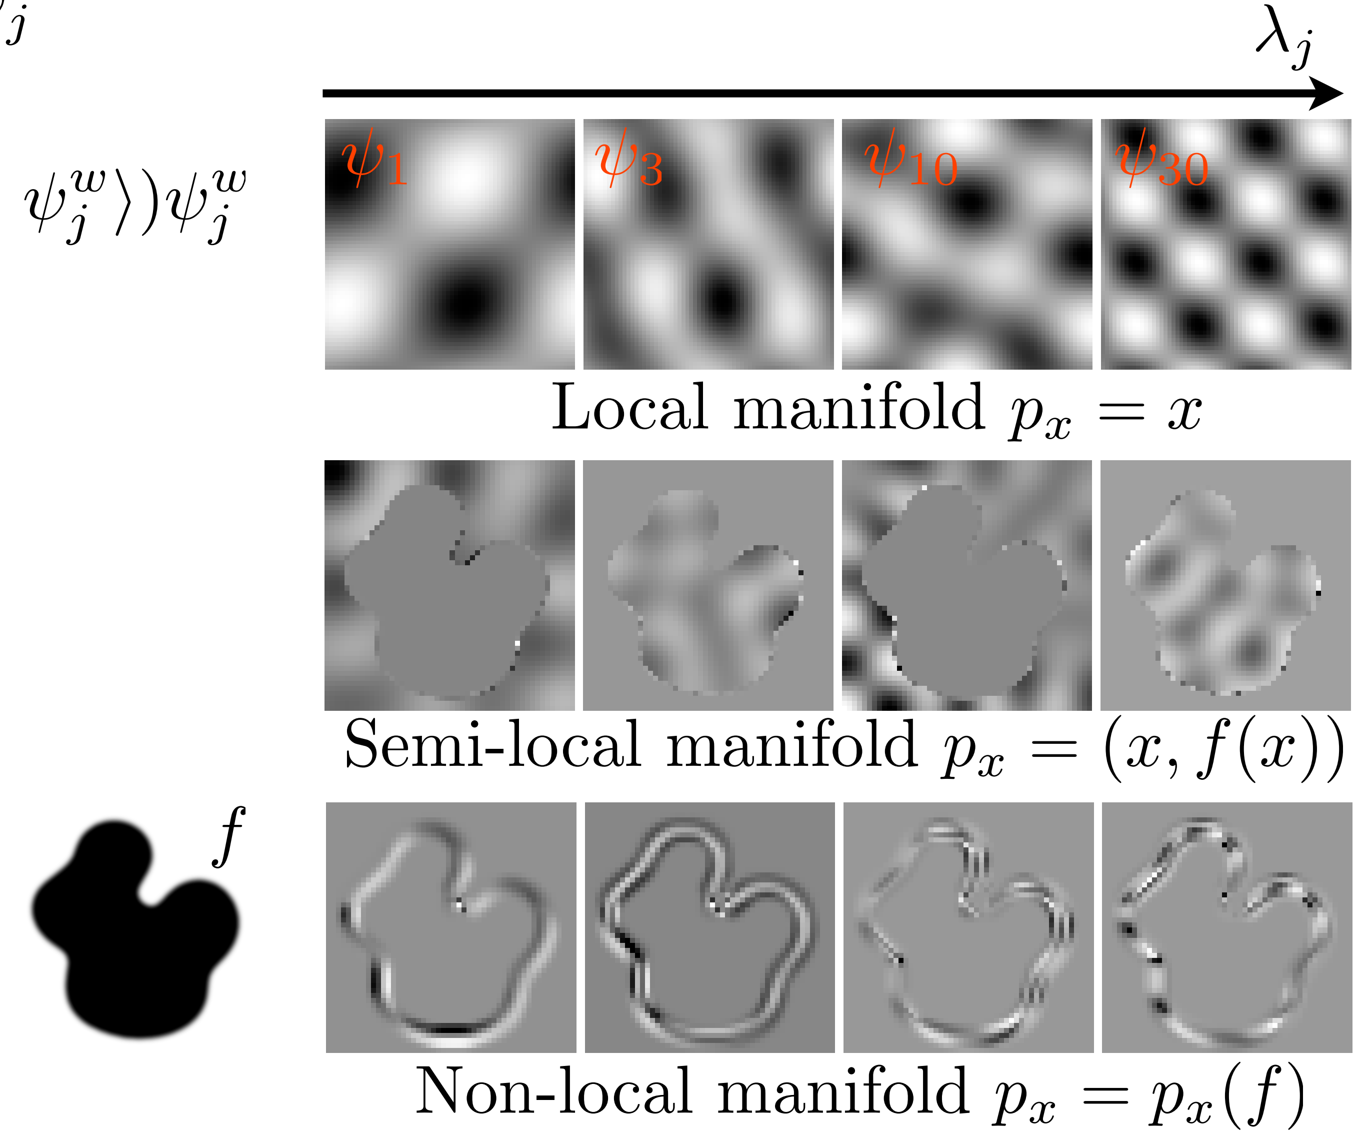

See [Peyré, SIAM MMS 2008]