Manifold Models for Signals and Images

Gabriel Peyré

► To cite this version:

Gabriel Peyré. Manifold Models for Signals and Images. Computer Vision and Image Understanding, 2009, 113 (2), pp.249-260. hal-00359729

HAL Id: hal-00359729

<https://hal.science/hal-00359729>

Submitted on 9 Feb 2009

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Manifold Models for Signals and Images

Gabriel Peyré^{*}

CNRS and CEREMADE, Université Paris-Dauphine

Abstract

This article proposes a new class of models for natural signals and images. The set of patches extracted from the data to analyze is constrained to be close to a low dimensional manifold. This manifold structure is detailed for various ensembles suitable for natural signals, images and textures modeling. These manifolds provide a low-dimensional parameterization of the local geometry of these datasets. These manifold models can be used to regularize inverse problems in signal and image processing. The restored signal is represented as a smooth curve or surface traced on the manifold that matches the forward measurements. A manifold pursuit algorithm computes iteratively a solution of the manifold regularization problem. Numerical simulations on inpainting and compressive sensing inversion show that manifolds models bring an improvement for the recovery of data with geometrical features.

Key words: signal processing, image modeling, texture, manifold.

PACS: code, code

Capturing the complex geometry of signals and images is at the core of recent advances in sound and natural image processing. Edges and texture patterns create complex non-local interactions. This paper studies these geometries for several sounds, images and textures models. The set of local patches in the dataset is modeled using smooth manifolds. These local features trace a continuous curve (resp. surface) on the manifold, which is a prior that can be used to solve inverse problems.

^{*} Place du Maréchal De Lattre De Tassigny, 75775 Paris Cedex 16, France
Email address: gabriel.peyre@ceremade.dauphine.fr (Gabriel Peyré).
URL: <http://www.ceremade.dauphine.fr/> (Gabriel Peyré).

1 Introduction

1.1 Previous Works

Global manifold models for image libraries. Dimensionality reduction methods such as Isomap [1], eigenmaps [2], LLE [3] or diffusion geometries [4] have been used to study the manifold structure of a library of images. The manifold regularity of certain images ensembles is emphasized by Donoho and Grimes [5] and Wakin et al. [6]. These global non-linear models have applications to image synthesis in computer vision [7]. Manifold valued functions are introduced in [8] and can be processed using multiscale methods. A global manifold model is used by Baraniuk and Wakin [9] to reconstruct an image from compressive sensing measurements.

Local edge manifold and cartoon images. The study of set of patches of 3×3 pixels extracted from natural images has been carried over by Lee et al. [10]. They report statistical evidences showing that the set of high contrast patches is located around the manifold of edges. These results have been refined by Carlsson et al. [12] that perform a simplicial approximation of the manifold.

Manifold parameterizations of local structures such as edges and corners is used in computer vision to detect salient features in images. Baker et al. [13] propose fast algorithms to search in a feature manifold. Huggins and Zucker [14] propose a local principal components analysis over a feature manifold to speed up computations.

Images with contours contain sharp variations along regular curves that make wavelets sub-optimal because of their square support [15]. Total variation methods [16] cannot make use of the regularity of the edge curves since they only constraint the overall length of the edges. A simple image model to describe geometric images is the cartoon model introduced by Donoho [17]. A cartoon function is regular outside a set of edge curves which are themselves regular. Several tools from harmonic analysis give optimal representations for such cartoon functions, including wedgelets [17], curvelets [18] and bandlets [19,20].

Cartoon images and the edge manifold is studied in Section 3.2.

Locally parallel textures. Some natural textures are composed of nearly parallel stripes that are modeled as local oscillations. This model of locally parallel textures is the extension to images of the model of locally stationary sounds. This model is studied Ben-Shahar and Zucker [21] who emphasize the role of the regularity of the underlying flow in image perception. Spatially varying orientations has been used in psychophysics as the simplest model for geometric textures, see for instance the model of second order edges of Landy and co-workers [22].

Demanet and Ying [23] propose a waveatom basis to capture efficiently the anisotropic

regularity of such textures. Adaptive decomposition such as the grouplets bases of Mallat [24] capture efficiently turbulent textures. The idea of characterizing textures as highly oscillating function has been introduced by Meyer [25]. Efficient algorithms to process such textures have been proposed by Aujol et al. [26] among others.

Section 3.3 studies locally parallel textures and shows how a patch manifold can be optimized to capture specific texture patterns.

Non-local and sparse patch processing. The most successful texture synthesis algorithms perform a consistent recopy of pixels and patches [27,28]. Local manifolds of patches have been used for texture synthesis and modification [29,30]. Non-local filtering based on patch comparison has been proposed by Buades et al. [31] to perform denoising, and adapted orthogonal basis can be constructed from non-local graphs [32]. Non local regularization is able to solve inverse problems such as super-resolution [33] or compressive sensing [34]

Wavelets and more recent tools from harmonic analysis [15] leads to efficient image compression, but fail to capture the geometry of natural images and textures [35]. The weakness of fixed representations can be alleviated by learning a dictionary from examples. Olshausen and Field [36] obtained an optimized set of oriented filters trained on patches extracted from natural images. Other algorithms have then been proposed that minimizes various sparsity enforcing criterions, see for instance [37–40]. These dictionary on small patches can be used to perform image denoising [41], facial image compression [42] and to solve inverse problems such as inpainting [43] or image separation [44].

Section 3.5 details a sparse model for patches originally introduced by Peyré [45] for texture synthesis. It fits into our patch manifold framework, and can thus be used to regularize inverse problems, as shown in Section 4.

1.2 Contributions

Section 3 details several patch manifolds for signals and images. This extends previous studies that consider mainly local edge models and makes the connection with non-local filtering and sparse patch decompositions.

Section 4 proposes a manifold regularization to solve inverse problems, which relies on the local manifold structure of images. Numerical results are shown on inpainting and compressive sensing reconstruction. This extends global manifold regularization that works for library of images and regularization with sparsity prior in patch dictionaries.

2 Manifolds of Signals and Images Patches

2.1 Local Manifold Model for Patches

A patch $p_x(f)$ of width $\tau > 0$ extracted from a signal ($d = 1$) or image ($d=2$) $f \in L^2([0, 1]^d)$ around $x \in [0, 1]^d$ is

$$\forall t \in [-\tau/2, \tau/2]^d, \quad p_x(f)(t) = f(x + t). \quad (1)$$

The point x is the center of the patch $p_x(f)$ and the width τ controls the size of typical features one wants to analyze.

A signal ensemble $\Theta \subset L^2([0, 1]^d)$ gathers typical data one is interested in. The patch manifold associated to this ensemble is

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

Section 3 details the structure of $\mathcal{M}$ for various signal and image ensembles Θ . Section 4 uses a fixed manifold $\mathcal{M}$ to regularize inverse problems.

Figure 1 shows two examples of image and texture ensembles Θ . The patches extracted from these images are parameterized by a small number of variables (for instance the direction of the edge, the orientation and frequency of the texture) so that the set $\mathcal{M}$ has the structure of a smooth manifold.

Fig. 1. Examples of patches $p_x(f) \in \mathcal{M}$ extracted from images $f \in \Theta$ in the cartoon images ensemble (left, see section 3.2) and in the oscillating textures ensemble (right, see section 3.4).

2.2 Manifold Parameterization

The set $\mathcal{M}$ is assumed to have a smooth manifold structure and is thus locally parameterized around a feature $p \in \mathcal{M}$ using a smooth mapping

$$\varphi : \Omega \subset \mathbb{R}^m \longrightarrow \varphi(\Omega) \subset \mathcal{M},$$

where $p \in \varphi(\Omega)$. The dimension m gives the number of parameters that describe the local geometry of images in Θ .

All the manifold studied in the sequel are globally described using a small number m of parameters. A global parameter domain $\Omega \subset \mathbb{R}^{m-k} \times (S^1)^k$ is used to parameterize the whole manifold. The circle $S^1 = \{\theta \mid |\theta| = 1\} \simeq [0, 2\pi)$ is used to account for directional features and can be replaced occasionally by the set of orientations (half circle)

$$\tilde{S}^1 = (S^1 / \sim) \simeq [0, \pi) \quad \text{where} \quad (\theta \sim \theta' \Leftrightarrow \theta = \theta' \bmod \pi).$$

Thanks to this global parameterization, any signal $f \in \Theta$ is represented as a curve (for signals when $d = 1$) or a surface (for images when $d = 2$) traced on the manifold $\mathcal{M}$ and equivalently over the parameter domain Ω

$$\begin{cases} c_f : x \mapsto p_x(f) \in \mathcal{M}, \\ \mathcal{C}_f : x \mapsto \varphi^{-1}(p_x(f)) \in \Omega. \end{cases} \quad (3)$$

For signals (resp. images) with periodic boundary conditions, the corresponding curve (resp. surface) is closed (resp. is topologically equivalent to a torus).

The mapping c_f is a manifold valued function and in the following, one denotes by $\mathcal{V}(\mathcal{M})$ the set of such manifold valued functions from $[0, 1]^d$ to $\mathcal{M}$. Such functions have been studied in [8].

2.3 Global Manifold Model for Signals and Images

The local manifold $\mathcal{M}$ introduced in (2) defines a patch manifold model that assigns to any patch $p \in L^2([-\tau/2, \tau/2]^d)$ its distance to $\mathcal{M}$

$$d(p, \mathcal{M}) = \|p - \text{Proj}_{\mathcal{M}}(p)\| \quad \text{where} \quad \text{Proj}_{\mathcal{M}}(p) = \underset{q \in \mathcal{M}}{\text{argmin}} \|p - q\| \quad (4)$$

The goodness of fit $d(p, \mathcal{M})$ to the manifold model is extended to any signal or image f by averaging the distances of all the patches

$$E_{\mathcal{M}}(f) = \int_{[0,1]^d} d(p_x(f), \mathcal{M})^2 dx = \int_{[0,1]^d} \|p_x(f) - \text{Proj}_{\mathcal{M}}(p_x(f))\|^2 dx. \quad (5)$$

A signal or an image f with a low energy $E_{\mathcal{M}}(f)$ traces a curve (resp. a surfaces) $c_f = \{p_x(f)\}_{x \in [0,1]^d}$ close to the low dimensional manifold $\mathcal{M}$. This curve is projected on the manifold $\mathcal{M}$ and on the parameter domain Ω as follow

$$\tilde{c}_f(x) = \text{Proj}_{\mathcal{M}}(p_x(f)) \in \mathcal{M} \quad \text{and} \quad \tilde{\mathcal{C}}_f(x) = \varphi^{-1}(\tilde{c}_f(x)) \in \Omega. \quad (6)$$

The projection $\text{Proj}_{\mathcal{M}}(f)$ of a signal f on the patch manifold model generated by $\mathcal{M}$ is achieved by computing the set of patches $c_f(x) = p_x(f)$, projecting these patches on the manifold $\tilde{c}_f(x) = \text{Proj}_{\mathcal{M}}(c_f(x))$ and then reconstructing a signal

$\text{Proj}_{\mathcal{M}}(f) = \text{Aver}(\tilde{c}_f)$ from $\tilde{c}_f$ using the averaging operator

$$\text{Aver}(c)(x) = \frac{1}{\tau^d} \int_{|x-z| \leq \tau/2} p_z(x-z) dz \quad \text{with} \quad p_z = c(z). \quad (7)$$

Section 3 studies for several signal and image ensembles Θ the geometry of $\tilde{c}_f$. Section 4 uses the manifold energy $E_{\mathcal{M}}$ to regularize inverse problems.

3 Examples of Patch Manifolds

This section studies the manifold model introduced in the previous section for various image and signal ensembles Θ .

3.1 Manifold of Smooth Variations

Smooth signals and images f belongs to the set

$$\Theta = \left\{ f \in C^1([0, 1]^d) \setminus \|f\|_{\infty} \leq a_{\max} \quad \text{and} \quad \left\| \frac{\partial f}{\partial x_i} \right\|_{\infty} \leq b_{\max} \right\}.$$

Typical patches from functions $f \in \Theta$ are well approximated by affine functions

$$\forall |t| \leq \tau/2, \quad p_x(f)(t) \approx a(x) + \langle b(x), t \rangle \quad \text{where} \quad \begin{cases} a(x) = f(x), \\ b(x) = \nabla_x f. \end{cases}$$

This leads us to consider the patch manifold of affine functions

$$\tilde{\mathcal{M}} = \{p_{a,b} \mid |a| \leq a_{\max} \text{ and } |b_i| \leq b_{\max}\}, \quad \text{where} \quad p_{a,b} : t \mapsto a + \langle b, t \rangle \quad (8)$$

which is close to the true manifold $\mathcal{M}$ generated by patches of Θ as defined in 2. In the following, for simplicity, we do not make the distinction between $\tilde{\mathcal{M}}$ and $\mathcal{M}$.

This affine manifold is parameterized globally as follow

$$\varphi : \begin{cases} [-a_{\max}, a_{\max}] \times [-b_{\max}, b_{\max}]^d \longrightarrow & \mathcal{M} \\ (a, b) & \longmapsto (t \mapsto a + \langle b, t \rangle). \end{cases}$$

This shows that $\mathcal{M}$ is a flat Euclidean manifold of dimension $d+1$. Geodesic distances along $\mathcal{M}$ are equal, up to a constant, to the Euclidean distance over the parameter domain. This analysis can be carried over for functions of higher smoothness or for bandlimited functions. These smooth signals ensembles are processed optimally using linear Fourier decomposition and linear filtering.

The projection $\text{Proj}_{\mathcal{M}}(p)$ of $p \in L^2([-\tau/2, \tau/2]^d)$ is computed by a linear regression.

Uniformly regular 1D signals. Smooth periodic signals in 1D are locally described using the 2D manifold $\mathcal{M}$ of affine 1D functions. Each signal f defines a 1D submanifold $\tilde{\mathcal{C}}_f \subset \mathcal{M}$, see equation (6). This submanifold corresponds to a curve over the parametric domain

$$\forall x \in [0, 1], \quad \tilde{\mathcal{C}}_f(x) = \varphi^{-1}(\text{Proj}_{\mathcal{M}}(p_x(f))) \approx (f(x), f'(x)) \in \mathbb{R}^2.$$

Figure 2, shows an example of a smooth function f together with the corresponding embedding $\tilde{\mathcal{C}}_f$ in (b). This shows that $\tilde{\mathcal{C}}_f$ exhibits a finite number of self intersections, which corresponds to points $x \neq y$ where $f(x) = f(y)$ and $f'(x) = f'(y)$.

Fig. 2. *Manifold of smooth signals.*

Uniformly regular 2D images. Smooth images f with periodic boundary conditions are described using the 3D manifold $\mathcal{M}$ of affine images. The manifold $\mathcal{M}$ thus defines a surface

$$\forall x \in [0, 1]^2, \quad \tilde{\mathcal{C}}_f(x_1, x_2) = \varphi^{-1}(\text{Proj}_{\mathcal{M}}(p_x(f))) \approx \left(f(x), \frac{\partial f}{\partial x_1}, \frac{\partial f}{\partial x_2} \right) \in \mathbb{R}^3.$$

Figure 3, left, shows an example of a regular image, that is generated as a realization of a Gaussian random field X whose power spectrum satisfies $|\hat{X}(\omega)| = (1 + |\omega|)^{-r}$ for a regularity exponent $r > 1$ (for the figure we have set $r = 3$).

Figure 3, right, shows the corresponding 2D surface $\tilde{\mathcal{C}}_f$ embedded in the 3D volume $\Omega = \varphi^{-1}(\mathcal{M})$. This surface is topologically equivalent to a torus but exhibits self intersections along 1D curves.

3.2 Manifold of Cartoon Images

Uniformly smooth signals and images are of little practical interest to model natural datasets. Signals with step discontinuities can be efficiently processed with wavelets that avoid ringing artifacts of a linear Fourier approximation. Image with contours

Image f Surface $\tilde{c}_f$
 Fig. 3. *Manifold of smooth images.*

contains sharp variations along regular curves that makes wavelets sub-optimal because of their square support [15].

A simple model of binary images is defined as

$$\Theta = \left\{ 1_B * h(x) \mid B \subset [0, 1]^2 \text{ with } \partial B \text{ regular} \right\}.$$

where h is a regular kernel. The compact set B represents the object of interest in the scene. It is supposed to be connected with ∂B of bounded curvature. This model can be extended to multiple objects as long as their boundaries are separated by a distance larger than τ .

Locally a patch of f is well approximated by a single straight edge

$$p_x(f)(t) \approx P_{(\theta(x), \delta(x))}(t) \quad \text{where} \quad P_{(\theta, \delta)}(t) = P(R_\theta(t - (\delta, 0))),$$

where R_θ is the planar rotation of angle θ . The step is $P = h * \tilde{P}$ where $\tilde{P}(t) = 0$ if $t_1 < 0$ and $\tilde{P}(t) = 1$ otherwise. Figure 4 shows some examples of typical edge patches.

Fig. 4. *Left: an example of cartoon image. Right: parameterization of the manifold of edge patches and some examples.*

This leads to the following 2D parameterization of the manifold of binary edge patches

$$\varphi : \begin{cases} S^1 \times \mathbb{R}^+ \longrightarrow \mathcal{M} \\ (\theta, \delta) \longmapsto P_{(\theta, \delta)}. \end{cases} \quad (9)$$

The manifold $\mathcal{M}$ is thus topologically equivalent to a cylinder, however, due to the lack of translation invariance when the edge approaches the boundary of $[-\tau/2, \tau/2]^2$, the manifold is not flat. Indeed the two constant functions equal to 0 and 1 play a special role of poles. Figure 5 shows a 3D display of the corresponding embedding. We note that for this edge manifold, one generally has $\mathcal{M} = \tilde{\mathcal{C}}_f$ since the boundary ∂B covers all the orientation $\theta \in S^1$. However the surface $\mathcal{C}_f$ traced by f on the parameter domain $S^1 \times \mathbb{R}^+$ can be complex when ∂B exhibits concavities that lead to self intersections of $\mathcal{C}_f$. Figure 5 shows curves that correspond to a 1D sections in $\mathcal{C}_f$. These curves trace closed loops on the manifold $\mathcal{M}$.

Fig. 5. *Left: a cartoon image. Right: 3D representation of the edge manifold $\mathcal{M}$ (depicted in 3D as a cylinder). The two curves on the manifold corresponds to patches extracted along the two lines in the image.*

3.3 Manifold of Locally Stationary Sounds

Natural sounds are usually modeled as highly oscillating signals with a phase that is slowly varying. Such a signal can be written as

$$f(x) = A(x) \cos(\Psi(x)),$$

where $A(x) \geq 0$ is the local amplitude, and $\Psi'(x) \geq 0$ the local phase of the oscillations. Such a decomposition is however non uniquely defined and one usually assumes that A and $\Psi'(x)$ are slowly varying with respect to the signal sampling so that they can be reliably estimated. This leads to the following signals ensemble

$$\Theta = \{x \mapsto f(x) = A(x) \cos(\Psi(x)) \mid \|A'\|_\infty \leq A_{\max} \quad \text{and} \quad \|\Psi''\|_\infty \leq \Psi_{\max}\}$$

This model of locally stationary signals leads to the manifold of constant oscillations

$$\mathcal{M} = \{P_{(A,\rho,\delta)} \mid A \geq 0 \quad \text{and} \quad \rho \geq 0 \quad \text{and} \quad \delta \in S^1\}$$

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

The parameterization $(A, \rho, \delta) \mapsto P_{(A, \rho, \delta)}$ shows that $\mathcal{M}$ is equivalent to $\Omega = \mathbb{R}^+ \times \mathbb{R}^+ \times \mathbb{S}^1$.

The projection of a patch $p \in L^2([-\tau/2, \tau/2])$ on $\mathcal{M}$ can be carried over approximately using a windowed Fourier transform. One uses a smooth window function h supported on $[-\tau/2, \tau/2]$ and defines the windowed Fourier transform of p

$$\hat{p}(\omega) = \int h(t)p(t) \exp(-i\omega t) dt.$$

Following Delprat et al. [46] (see also [15]), the projection of p is then given as

$$\text{Proj}_{\mathcal{M}}(p) \approx P_{(A, \rho, \delta)} \quad \text{where} \quad \begin{cases} \rho = \underset{\omega \geq 0}{\text{argmax}} |\hat{p}(\omega)| \\ \hat{p}(\rho) = A \exp(i\delta). \end{cases} \quad (10)$$

A 1D signal f defines a 1D curve $\tilde{c}_f \subset \mathcal{M}$ traced on the manifold and a 1D curve $\tilde{C}_f$ in 3D parameter space

$$\tilde{C}_f = \{(A(x), \rho(x), \delta(x))\}_{x \in [0, 1]} \quad \text{where} \quad P_{(A(x), \rho(x), \delta(x))} = \text{Proj}_{\mathcal{M}}(p_x(f)).$$

Figure 6 shows examples of a locally stationary oscillating signal together with its spectrogram and the corresponding curve $\tilde{C}_f$ over the parametric space.

Fig. 6. *Upper-left: 1D signal of a bird singing, bottom left: the corresponding log-spectrogram $\log(|\hat{p}_x(f)(\omega)|)$ (the blue curve is the maxima curve $(x, \rho(x))$), right: the 2D curve $(A(x), \rho(x))$. For the display, the curve has been disconnected in areas where the bird stop singing (characterized by a low value of $A(x)$).*

3.4 Manifold of Locally Parallel Textures

Some natural textures are composed of nearly parallel stripes that are modeled as local oscillations. A locally parallel texture is defined as $f(x) = A(x)h(\Psi(x)) + B(x)$ where $\nabla_x \Psi$ controls the direction and frequency of the oscillations around the point x , $A(x)$ is a local amplitude, $B(x)$ a local shift, and $h : \mathbb{R} \rightarrow \mathbb{R}$ is the 1D profile of the oscillations. This model of locally parallel textures is the extension of locally stationary sounds to images, if one considers the case $h(x) = \cos(x)$. Figure 1, right, shows examples of locally parallel textures.

A local texture patch is approximated as

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

The parameter $\theta \in \tilde{S}^1$ is the direction of the oscillations, ρ their frequency and δ is a phase shift. The manifold $\mathcal{M}$ is parameterized as

$$\varphi : \begin{cases} \mathbb{R}^+ \times \mathbb{R} \times \mathbb{R}^+ \times \tilde{S}^1 \times S^1 \longrightarrow & \mathcal{M} \\ (A, B, \rho, \theta, \delta) & \longmapsto P_{(A, B, \rho, \theta, \delta)}. \end{cases} \quad (11)$$

An approximate estimation of $(\rho(x), \theta(x), \delta(x))$ from a texture f can be carried over using local Fourier expansions over windows of size τ as defined in the 1D case in equation (10).

Image f Subsampling $\text{Proj}_{\mathcal{M}}(f)$
 Fig. 7. An image and its projection $\text{Proj}_{\mathcal{M}}(f)$ on the manifold model of oscillating textures. The center image displays a projection computed with a subset of non-overlapping patches $p_x(f)$, to better see the importance of translation invariance.

The feature manifold depends on the specific choice of the profile h

$$\mathcal{M} = \mathcal{M}_h = \left\{ P_{(A, B, \rho, \theta, \delta)} \mid (A, B, \rho, \theta, \delta) \in \Omega \right\}.$$

An adaptation of the manifold to an exemplar $f \in L^2([0, 1]^d)$ to process is achieved by optimizing the profile h so that the patches of f are as close as possible from the manifold $\mathcal{M}_h$. This amounts to minimizing the energy $E_{\mathcal{M}_h}$ defined in (5) to select

an adapted profile h^*

$$h^* = \underset{h \in H}{\operatorname{argmin}} E_{\mathcal{M}_h}(f), \quad (12)$$

where $H \subset L^2(\mathbb{R})$ is a set of profiles, that might for instance contains some smoothness constraints on the texture profile.

We restrict the computation to a simple set of 1D profiles parameterized by a contrast $\gamma \in (0, 1]$

$$h_\gamma(x) = \operatorname{sign}(\sin(x)) |\sin(x)|^\gamma. \quad (13)$$

The optimal γ minimizing $E_{\mathcal{M}_{h_\gamma}}(f)$ is computed by testing a discrete set of values in $[0, 1]$. Figure 7 shows examples of projection with $\gamma = 0.35$, which is the value minimizing $E_{\mathcal{M}_{h_\gamma}}(f)$ for the fingerprint image f . More elaborated adaptation strategies could be used to fit arbitrary texture profiles.

3.5 Manifold of Sparse Patches

Sparse patch expansion. The sparse patch model assumes that each patch $p_x(f)$ extracted from a signal or image f has a sparse expansion

$$p_x(f) = \sum_i s_x(i) \varphi_i = D s_x$$

using a dictionary $D = \{\varphi_i\}_{i=0}^{P-1}$ of P atoms. Each $\varphi_i \in L^2([- \tau/2, \tau/2]^d)$ is an atomic template and $s_x(i) \in \mathbb{R}$ is the corresponding coefficient. The sparsity of the decomposition is ensured by constraining the ℓ^0 pseudo-norm of the coefficients to be smaller than $S \in \mathbb{N}^*$

$$\|s_x\|_{\ell^0} = \# \{i \mid s_x(i) \neq 0\} \leq S.$$

This leads to the manifold of sparse patches

$$\mathcal{M} = \mathcal{M}_D = \left\{ \sum_i s(i) \varphi_i \mid \|s\|_{\ell^0} \leq S \right\}.$$

This set $\mathcal{M}$ is not a smooth manifold but rather a non-linear union of S -dimensional linear spaces. It is parameterized by the dictionary D and the sparsity level S .

The corresponding signal ensemble Θ with sparse patches is

$$\Theta = \Theta_D = \{f \mid \forall x, p_x(f) = D s_x \text{ with } \|s_x\|_{\ell^0} \leq S\} \quad (14)$$

This model has been introduced by Peyré [45] for texture synthesis and is re-casted in our manifold patch model.

The projection on the sparse manifold $\mathcal{M}$ requires the computation of

$$\operatorname{Proj}_{\mathcal{M}}(p) = \underset{s \in \mathbb{R}^m}{\operatorname{argmin}} \|p - Ds\|_{\ell^2} \quad \text{subject to} \quad \|s\|_{\ell^0} \leq S.$$

The exact computation of this projection is NP-hard for an arbitrary dictionary D , but it can be solved approximately using for instance S steps of orthogonal matching pursuit algorithm, see [15].

Dictionary learning. The manifold $\mathcal{M}_D$ is optimized so that a given exemplar f is as close as possible to Θ_D . This corresponds to the learning of the dictionary D , that is optimized in order to minimize the goodness of fit of f to the model generated by D

$$D^* = \underset{D}{\operatorname{argmin}} E_{\mathcal{M}_D}(f) \quad (15)$$

where the energy $E_{\mathcal{M}_D}$ is defined in (5). This is similar to the adaptation of the profile to optimize the oscillating texture model (12). In the optimization (15), the atoms φ_i of D are only constrained to be of unit norm, $\|\varphi_i\| = 1$.

The optimization (15) is re-written by optimizing over both the dictionary D and the coefficients s_x of the patches $p_x(f)$

$$(D^*, \{s_x^*\}_x) = \underset{D, \{s_x\}_x}{\operatorname{argmin}} \sum_x \|p_x(f) - Ds_x\|^2 \quad \text{subject to} \quad \|s_x\|_{\ell^0} \leq S. \quad (16)$$

This optimization problem is highly non-convex and a stationary point of $E_{\mathcal{M}_D}$ with respect to D can be computed using several iterative algorithms, for instance the MOD algorithm [37] or K-SVD [40]. Table 1 details the MOD optimization process, that alternates between the computation of the coefficients $\{s_x\}_x$ and the optimization of the dictionary atoms φ_i .

Table 1 MOD dictionary learning algorithm.

- (1) *Initialization:* set each φ_i as a realization of a white noise, normalized so that $\|\varphi_i\| = 1$.
- (2) *Coefficient update:* compute the projection Ds_x of each patch $p_x(f)$, on the manifold $\mathcal{M}$

$$s_x = \underset{s}{\operatorname{argmin}} \|p_x(f) - Ds\| \quad \text{subject to} \quad \|s\|_{\ell^0} \leq S.$$

This is computed approximately with S steps of orthogonal matching pursuit applied to $p_x(f)$.

- (3) *Dictionary update:* the dictionary D is computed by minimizing

$$\min_D \sum_x \|p_x(f) - Ds_x\|^2,$$

whose solution is given as

$$D \leftarrow PS^+ \quad \text{where} \quad S^+ = (S^T S)^{-1} S^T$$

where $S = \{s_x\}_x$ is the matrix whose columns are the coefficients s_x and $P = \{p_x(f)\}_x$ is the matrix whose columns are the (discretized) patches $p_x(f)$.

- (4) *Normalization:* for all i , $\varphi_i \leftarrow \varphi_j / \|\varphi_i\|$.
 - (5) *Stop:* while not converged, go back to 2.
-

Figure 8 shows an example of a dictionary learned from an homogeneous texture.

Fig. 8. *Left: input texture f , right: dictionary D learned to sparsity the patches $p_x(f)$.*

4 Inverse Problem Regularization with Manifold Models

4.1 Inverse Problems

The manifold prior model introduced in Section 2.3 is used to regularize the inversion of an operator $\Phi : L^2([0, 1]^d) \rightarrow V$, where V is an Hilbert space of finite or infinite dimension. The mapping Φ is typically ill-posed and difficult to invert since Φf gathers only a limited amount of information from the original f to recover.

Forward measurements correspond to the computations of

$$y = \Phi f + \varepsilon \in V,$$

where f is the data to recover and ε is an additive noise.

Regularization theory assumes that f belongs to some functional space H such as a Sobolev space (linear regularization) or the space of bounded variations (non-linear regularization). The recovered signal f^* is the solution of an optimization problem

$$f^* = \operatorname{argmin}_{g \in H} \|y - \Phi g\|^2 + \lambda E(g), \quad (17)$$

where E should be small when g is close to the smoothness model. The weight λ should be adapted to match the amplitude of the noise ε , which might be a non-trivial task in practical situations.

Classical variational priors include

- *Total variation*: The bounded variation model imposes that f^* has a finite bounded variation and uses

$$E(g) = \|g\|_{\text{TV}} = \int |\nabla_x g| dx. \quad (18)$$

This prior has been introduced by Rudin, Osher and Fatemi [16] for denoising purpose.

- *Sparsity priors*: Given an orthogonal basis $\{\psi_k\}_k$ of $L^2([0, 1]^d)$, a sparsity enforcing prior is defined as

$$E(g) = \sum_k |\langle g, \psi_k \rangle|. \quad (19)$$

This prior has been introduced by Donoho and Johnstone [47] with the wavelet basis for denoising purpose. It has then been used to solve more general inverse problems, see for instance [48] and the references therein. It can also be used in conjunction with redundant frames instead of orthogonal bases, see for instance [49, 50].

4.2 Regularization with Manifold Model

This paper explores the global manifold energy $E = E_{\mathcal{M}}$ defined in equation (5) to perform the regularization. The optimization (17) with $E = E_{\mathcal{M}}$ is performed by introducing a manifold valued function $c^* \in \mathcal{V}(\mathcal{M})$. At each location x , the patch $c^*(x) \in \mathcal{M}$ is tracking a local feature of f^* and should be close to $p_x(f)$.

The optimization (17) is re-written using c^* as

$$(f^*, c^*) = \underset{(g, c) \in L^2([0, 1]^d) \times \mathcal{V}(\mathcal{M})}{\operatorname{argmin}} \mathcal{E}_{\mathcal{M}}(g, c), \quad (20)$$

$$\text{where } \mathcal{E}_{\mathcal{M}}(g, c) = \|y - \Phi g\|^2 + \lambda \int_{[0, 1]^d} \|p_x(g) - c(x)\|^2 dx \quad (21)$$

whose solution f^* also solves the original problem (17) for $E = E_{\mathcal{M}}$.

The manifold valued mapping c^* should be close to the original projected curve (or surface when $d = 2$) $\tilde{c}_f$ as defined in (6). The manifold regularization of inverse problems (20) computes a curve (or surface) traced on the manifold $\mathcal{M}$ that also matches the forward measurements y .

A stationary point of (20) is computed by alternatively minimizing over f^* and c^* :

- *The image is fixed*. The manifold valued function c_0 minimizing $\mathcal{E}_{\mathcal{M}}(g, c)$ with respect to $c \in \mathcal{V}(\mathcal{M})$ is computed with the projection defined in equation (4)

$$\forall x \in [0, 1]^d, \quad c_0(x) = \operatorname{Proj}_{\mathcal{M}}(p_x(g)). \quad (22)$$

- *The manifold-valued mapping is fixed*. The signal g_0 minimizing $\mathcal{E}_{\mathcal{M}}(g, c)$ with respect to $g \in L^2([0, 1]^d)$ is the solution of

$$(\Phi^T \Phi + \lambda \operatorname{Id})g_0 = \Phi^T y + \lambda \operatorname{Aver}(c)$$

where the averaged signal $\operatorname{Aver}(c) \in L^2([0, 1]^d)$ of c is defined in (7).

The iteration of these two steps corresponds to the manifold pursuit detailed in table 2.

Table 2 Manifold pursuit to minimize (20).

- (1) *Initialization*: set $f^{(0)} = \Phi^T y$ and $k \leftarrow 0$.
- (2) *Manifold closest point*: update the manifold valued function as

$$\forall x \in [0, 1]^d, \quad c^{(k+1)}(x) = \text{Proj}_{\mathcal{M}}(p_x(f^{(k)})),$$

where the manifold projection $\text{Proj}_{\mathcal{M}}$ is defined in (4).

- (3) *Least square fit*: update the current estimate as

$$f^{(k+1)} = (\Phi^T \Phi + \lambda \text{Id})^{-1} (\Phi^T y + \lambda \text{Aver}(c^{(k+1)})).$$

and where the averaging $\text{Aver}(c^{(k+1)})$ of $c^{(k+1)}$ is defined in (7).

- (4) *Stopping criterion*: while not converged, set $k \leftarrow k + 1$ and go back to 2.
-

The algorithm of table 2 shares some similarities with iterative thresholdings methods used to solve the non-linear regularized inversion (17) with a sparsity prior (19), see for instance [48]. To handle the noiseless case $\varepsilon = 0$, the value of the regularization parameter λ can be decreased toward 0 during the iterations.

The main computational burden in the manifold pursuit is the numerical computation of the projection $\text{Proj}_{\mathcal{M}}$ on the manifold. Depending on the image model considered, specific solvers can be used. For complex manifolds without any analytical description, a dense sampling of the manifold is used together with a fast closest-point algorithm.

Since the energy to optimize (20) is non-convex, the manifold pursuit might fail to converge to the global minimizer of the problem. For a smooth manifold $\mathcal{M}$, the iterates $(f^{(k)}, c^{(k)})$ of the algorithm 2 converge to a stationary point (f^*, c^*) of the energy $\mathcal{E}_{\mathcal{M}}$ since this energy satisfies the hypotheses of [51].

4.3 Numerical Experiments

Although any linear operator can be treated within our regularization framework, this paper focusses on the following examples.

- *Inpainting* corresponds to the operation of removing pixels from an input data

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

where $\Omega \subset [0, 1]^d$ is the region where the input data has been damaged. Classical methods for inpainting use partial differential equations that propagate the infor-

mation from the boundary of Ω to its interior, see for instance [52–55]. Sparsity promoting prior such as (19) in wavelets frames and local cosine bases have been used to solve inpainting as a inverse problem [49,50].

- *Compressive sensing* is a new sampling theory that uses a fixed set of linear measurements together with a non-linear reconstruction [56,57], see also the review [58]. The sensing operator computes the projection of the data on a finite set of k vectors

$$\Phi f = \{\langle f, u_i \rangle\}_{i=0}^{k-1} \in \mathbb{R}^k. \quad (23)$$

The signal are discretized finite dimensional vectors $f \in \mathbb{R}^n$. Compressive sensing states hypotheses on both the input signal f and the sensing vectors $\{u_i\}_i$ for this non-uniform sampling process to be invertible with ℓ^1 minimization. The sensing vectors $\{u_i\}_i$ must be incoherent, which is the case with high probability if they are drawn randomly from unit norms Gaussian white noise vectors. Under the additional condition that f is sparse in some orthogonal basis $\{\psi_k\}_k$

$$\#\{j \mid \langle f, \psi_k \rangle \neq 0\} \leq S,$$

the optimization of (17) with the sparsity prior (19) leads to a perfect recovery $f^* = f$ if $k = O(S \log(n/S))$, where n is the dimension of the sampled signal f . This result holds in the noiseless case $\varepsilon = 0, \lambda \rightarrow 0$ and can be extended to an approximate recovery in the noisy case $\varepsilon \neq 0, \lambda > 0$.

The following numerical experiments compare the efficiency of sparsity regularization with ℓ^1 prior (19) with the patch manifold energy $E_{\mathcal{M}}$ (5). Sparsity regularization leads to a convex optimization, which is an advantage over the non-convex manifold regularization that might be trapped in a local minimum.

For numerical computation, discretized signals and images are obtained by an uniform sampling at n points. The corresponding manifolds of patches $\mathcal{M} \subset \mathbb{R}^{\tau^{d_n}}$ are embedded in a finite dimensional space. The following numerical experiments are performed with patches of width $w = 10$ pixels. Compressive sensing experiments are performed with sensing vectors u_i that are random discrete Fourier vectors, which allows one to compute the sensing operator Φ with the FFT algorithm.

Smooth images. Figure 9 shows iterations of the algorithm 2 to solve the inpainting problem on a smooth image using a manifold prior with 2D linear patches, as defined in 8, with a low amplitude noise ε . This manifold regularization together with the overlapping of the patches performs a smooth interpolation of the missing pixels.

The iterations of the algorithm are similar to a linear diffusion that propagates the available information inside the set Ω of removed pixels. The performances of the algorithm are similar to linear methods such as inpainting with a Sobolev regularizer $E(f) = \int |\nabla f|^2$ that corresponds to a heat diffusion inside Ω .

Cartoon images. Figures 10 and 11 show iterations of the projection algorithm 2 with a manifold model of binary edges, as defined in equation (9). For this numerical

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

Fig. 9. *Iterations of the inpainting algorithm on an uniformly regular image.*

optimization, the manifold of edges is discretized as already done for the display of figure 5 and the projection $\text{Proj}_{\mathcal{M}}$ is computed with a fast nearest-neighbor search. For both inpainting and compressive sampling, the manifold of edges allows to reconstruct with good precision the boundary of a single smooth object (here a disk).

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

Fig. 10. *Iterations of the inpainting algorithm on a geometrical image with the binary edge model.*

Iter. #1 Iter. #2 Iter. #3 Iter. #50

Fig. 11. *Iterations of the compressive sensing reconstruction algorithm on a geometrical image with the binary edge model. The number of sensed vectors is $n_0 = n/10$ where n is the number of pixels.*

Figure 12 shows a more challenging compressive sensing problem where the image is composed of layers of occluding objects with smooth boundaries and varying intensities. In order to cope with a non-binary image, the manifold of affine edges is used

$$\mathcal{M} = \left\{ aP_{(\theta,\delta)} + b \mid (\theta,\delta) \in \mathbb{S}^1 \times \mathbb{R}^+ \text{ and } (a,b) \in \mathbb{R}^+ \times \mathbb{R} \right\}. \quad (24)$$

This manifold is four dimensional, but the projection on this manifold is computed efficiently by iteratively optimizing the projection over (θ,δ) parameters and then the (a,b) parameters. The patches extracted from figure 12, left, are however not

always close to this manifold because of crossings that occur when two singularity curves meet.

Figure 12 compares the compressive sensing reconstruction with a sparsity prior (19) in a translation invariant wavelet frame (center) and with a manifold prior in the affine edges manifold (24). The optimization of the sparsity energy (19) is performed with an iterative thresholding algorithm, see [48], whereas the optimization of the manifold energy (20) is performed with the algorithm 2. The reconstruction with a manifold prior is of better quality than the sparsity prior in a wavelet frame. This is because, in 2D, wavelets cannot take advantage of the regularity of the boundaries of the objects.

Original f

Wavelets, SNR=25.7dB

Manifold, SNR=31.3dB

Fig. 12. *Compressive sensing reconstruction results on a geometrical image with sparsity prior in wavelets and with the manifold model of affine edges. The number of sensed vectors is $n_0 = n/8$ where n is the number of pixels.*

Locally parallel textures. Figure 13 shows an example of inpainting of a fingerprint image using the manifold of 2D oscillations (11). Figure 14 shows the compressive sensing reconstruction from the same image, where the manifold prior is compared to a sparsity prior (19) in a local Gabor redundant frame, see [15]. Such local oscillating atoms have been introduced with success for texture decomposition and inpainting in [49,50]. The manifold prior is better able to capture the geometric regularity of the texture than the sparsity prior that diffuses the orientation information over several Gabor coefficients.

Iter. #1

Iter. #2

Iter. #3

Iter. #50

Fig. 13. *Iterations of the inpainting reconstruction algorithm on a locally parallel texture.*

Original f

Gabor, SNR=7.9dB

Manifold, SNR=9.5dB

Fig. 14. *Compressive sensing reconstruction results on a locally parallel texture with sparsity prior in a redundant Gabor dictionary and with the manifold model. The number of sensed vectors is $n_0 = n/4$ where n is the number of pixels.*

Sparse patches. Figure 15 shows a reconstruction from compressive sensing using a manifold prior in the sparse texture ensemble (14). The dictionary D is learned by minimizing $E_{\mathcal{M}_D}(f_e)$ for an exemplar texture f_e that is close to the texture f to recover. In practice, both image are extracted from different localization of the same image. The central part of the texture is badly reconstructed, because this part is not similar to the exemplar f_e .

Original f

Iter. #1

Iter. #1

Iter. #20

Fig. 15. *Iterations of compressive sensing reconstruction results using a manifold model in a learned dictionary. The number of sensed vectors is $n_0 = n/6$ where n is the number of pixels.*

Figure 16 shows a comparison of inpainting with a sparse prior (19) in a translation invariant wavelet frame and a sparse patch manifold model prior. The dictionary D is learned from a set of patches $p_{x_i}(y)$ extracted from the observed data, where the point x_i are located at a distance larger than τ from the missing region to be inpainted. The resulting algorithm is similar to the inpainting algorithm of Mairal et al. [43], although they use a more complicated scheme to learn a dictionary with missing data.

Measurements y Wavelets, SNR=22.5dB Manifold, SNR=31.1dB
 Fig. 16. *Comparison of inpainting using a sparsity prior in a wavelet basis and the sparse manifold model with a dictionary learned from the observed data y .*

Conclusion

This paper has reviewed several manifold models for sounds, images and textures. These models constrain the set of patches extracted from the image and describe efficiently the non-linear geometry of some classes of natural signals and images. A new manifold pursuit algorithm is used to regularize ill-posed inverse problem while maintaining the manifold model constraints. Results on various images and textures show how this manifold-driven restoration enhances variational models based on sparse expansions for inpainting and compressive sensing reconstruction.

Acknowledgments

I would like to thank the anonymous reviewers for their comments that helped to improve the quality of the manuscript.

References

- [1] J. B. Tenenbaum, V. de Silva, J. C. Langford, A Global Geometric Framework for Nonlinear Dimensionality Reduction, *Science* 290 (5500) (2000) 2319–2323.
- [2] M. Belkin, P. Niyogi, Laplacian eigenmaps for dimensionality reduction and data representation, *Neural Computation* 15 (6) (2003) 1373–1396.
- [3] S. Roweis, L. Saul, Nonlinear Dimensionality Reduction by Locally Linear Embedding, *Science* 290 (5500) (2000) 2323–2326.
- [4] R. R. Coifman, S. Lafon, A. B. Lee, M. Maggioni, B. Nadler, F. Warner, S. W. Zucker, Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps, *Proceedings of the National Academy of Science* 102 (2005) 7426–7431.

- [5] D. Donoho, C. Grimes, Image manifolds isometric to Euclidean space, *J. Math. Imaging Computer Vision* 23 (1).
- [6] M. Wakin, D. Donoho, H. Choi, R. Baraniuk, The Multiscale Structure of Non-Differentiable Image Manifolds, in: *Optics and Photonics*, San Diego, CA, 2005.
- [7] M. Aharon, R. Kimmel, Representation analysis and synthesis of lip images using dimensionality reduction, Accepted to *International Journal of Computer Vision* 67 (3).
- [8] I. Ur-Raman, I. Drori, V. Stodden, P. Schroeder, Multiscale representations of manifold-valued data, *SIAM Multiscale Modeling and Simulation* 4 (4) (2005) 1201–1232.
- [9] R. G. Baraniuk, M. B. Wakin, Random projections of smooth manifolds, to appear in *Foundations of Computational Mathematics*.
- [10] A. Lee, K. Pedersen, D. Mumford, The nonlinear statistics of high-contrast patches in natural images, *International Journal of Computer Vision* 54 (1-3) (2003) 83–103.
- [11] G. Carlsson, T. Ishkhanov, V. de Silva, A. Zomorodian, On the local behavior of spaces of natural images, *Int. J. Comput. Vision* 76 (1) (2008) 1–12.
- [12] S. Baker, S. Nayar, H. Murase, Parametric feature detection, *International Journal of Computer Vision* 27 (1) (1998) 27–50.
- [13] P. S. Huggins, S. W. Zucker, Representing edge models via local principal component analysis, *Proc. of ECCV'02*, Springer LNCS 2350 (2002) 384–398.
- [14] S. Mallat, *A Wavelet Tour of Signal Processing*, Academic Press, San Diego, 1998.
- [15] L. I. Rudin, S. Osher, E. Fatemi, Nonlinear total variation based noise removal algorithms, *Phys. D* 60 (1-4) (1992) 259–268.
- [16] D. Donoho, Wedgelets: Nearly-minimax estimation of edges, *Ann. Statist* 27 (1999) 353–382.
- [17] E. Candès, D. Donoho, New tight frames of curvelets and optimal representations of objects with piecewise C^2 singularities, *Comm. Pure Appl. Math.* 57 (2) (2004) 219–266.
- [18] E. Le Pennec, S. Mallat, Bandelet Image Approximation and Compression, *SIAM Multiscale Modeling and Simulation* 4 (3) (2005) 992–1039.
- [19] S. Mallat, G. Peyré, Orthogonal bandlet bases for geometric images approximation, *Communications on Pure and Applied Mathematics* 61 (9) (2008) 1173–1212.
- [20] O. Ben-Shahar, S. Zucker, The perceptual organization of texture flow: A contextual inference approach, *IEEE Trans. Pattern Anal. Mach. Intell* 25 (4) (2003) 401–417.
- [21] M. S. Landy, I. Oruc, Properties of 2nd-order spatial frequency channels, *Vision Research* 42 (2002) 2311–2329.
- [22] L. Demanet, L. Ying, Wave atoms and sparsity of oscillatory patterns, *Applied and Computational Harmonic Analysis* 3 (23) (2007) 368–387.

- [23] S. Mallat, Geometrical grouplets, Submitted.
- [24] Y. Meyer, Oscillating Patterns in Image Processing and Nonlinear Evolution Equations, American Mathematical Society, Boston, MA, USA, 2001.
- [25] J. F. Aujol, G. Aubert, L. Blanc-Feraud, A. Chambolle, Image decomposition into a bounded variation component and an oscillating component, *Journal of Mathematical Imaging and Vision* 22 (1) (2005) 71–88.
- [26] A. A. Efros, T. K. Leung, Texture synthesis by non-parametric sampling, in: *ICCV '99: Proceedings of the International Conference on Computer Vision-Volume 2*, IEEE Computer Society, 1999, p. 1033.
- [27] L.-Y. Wei, M. Levoy, Fast texture synthesis using tree-structured vector quantization, in: *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, ACM Press/Addison-Wesley Publishing Co., 2000, pp. 479–488.
- [28] J. Wang, X. Tong, S. Lin, M. Pan, C. Wang, H. Bao, B. Guo, H.-Y. Shum, Appearance manifolds for modeling time-variant appearance of materials, in: *Proc. of Siggraph'06*, ACM Press, New York, NY, USA, 2006, pp. 754–761.
- [29] W. Matusik, M. Zwicker, F. Durand, Texture design using a simplicial complex of morphable textures, *ACM Trans. Graph.* 24 (3) (2005) 787–794.
- [30] A. Buades, B. Coll, J. M. Morel, On image denoising methods, *SIAM Multiscale Modeling and Simulation* 4 (2) (2005) 490–530.
- [31] G. Peyré, Manifold spectral bases for the non-local processing of images, *SIAM Multiscale Modeling and Simulation* 7 (2) (2008) 703–730.
- [32] M. Protter, M. Elad, H. Takeda, P. Milanfar, Generalizing the non-local-means to super-resolution reconstruction, to appear in *IEEE Transactions on Image Processing*.
- [33] G. Peyré, S. Bogleux, L. Cohen, Non-local regularization of inverse problems, in: *Proc. of ECCV 2008*, 2008.
- [34] E. P. Simoncelli, B. A. Olshausen, Natural image statistics and neural representation, 2001 24 (*Annual Review of Neuroscience*) 1193–1215.
- [35] B. A. Olshausen, D. J. Field, Emergence of simple-cell receptive-field properties by learning a sparse code for natural images, *Nature* 381 (6583) (1996) 607–609.
- [36] K. Engan, S. O. Aase, J. H. Husoy, Method of optimal directions for frame design, in: *Proc. ICASSP '99*, IEEE Computer Society, Washington, DC, USA, 1999, pp. 2443–2446.
- [37] K. Kreutz-Delgado, J. F. Murray, B. D. Rao, K. Engan, T.-W. Lee, T. Sejnowski, Dictionary learning algorithms for sparse representation, *Neural Comput.* 15 (2) (2003) 349–396.
- [38] M. S. Lewicki, T. J. Sejnowski, Learning overcomplete representations, *Neural Comput.* 12 (2) (2000) 337–365.

- [39] M. Aharon, M. Elad, A. Bruckstein, The k-svd: An algorithm for designing of overcomplete dictionaries for sparse representation, *IEEE Trans. On Signal Processing* (to appear) 54 (11) (2006) 4311–4322.
- [40] M. Elad, M. Aharon, Image denoising via sparse and redundant representations over learned dictionaries, *IEEE Trans. on Image Processing* 15 (12) (2006) 3736–3745.
- [41] O. Bryt, M. Elad, Compression of facial images using the K-SVD algorithm, *Journal of Visual Communication and Image Representation* 19 (4) (2008) 270–282.
- [42] J. Mairal, M. Elad, G. Sapiro, Sparse representation for color image restoration, *IEEE Trans. on Image Processing* 17 (1) (2008) 53–69.
- [43] G. Peyré, J. Fadili, J. Starck, Learning dictionaries for geometry and texture separation, *Proc. of SPIE Wavelet XII*.
- [44] G. Peyré, Sparse modeling of textures, to appear in *Journal of Mathematical Imaging and Vision*.
- [45] N. Delprat, B. Escudié, P. Guillemain, R. Kronland-Martinet, P. Tchamitchian, B. Torrèsani, Asymptotic wavelet and gabor analysis: extraction of instantaneous frequencies, *IEEE Trans. Inf. Th.* 38 (1992) 644–664.
- [46] D. Donoho, I. Johnstone, Ideal spatial adaptation via wavelet shrinkage, *Biometrika* 81 (1994) 425–455.
- [47] I. Daubechies, M. Defrise, C. D. Mol, An iterative thresholding algorithm for linear inverse problems with a sparsity constraint, *Comm. Pure Appl. Math* 57 (2004) 1413–1541.
- [48] M. Elad, J.-L. Starck, D. Donoho, P. Querre, Simultaneous cartoon and texture image inpainting using morphological component analysis (mca), *Journal on Applied and Computational Harmonic Analysis* 19 (2005) 340–358.
- [49] M. Fadili, J.-L. Starck, F. Murtagh, Inpainting and zooming using sparse representations, *The Computer Journal* (to appear).
- [50] P. Tseng, Convergence of a block coordinate descent method for nondifferentiable minimization, *Journal of Optimization Theory and Applications* 109 (3) (2001) 475–494.
- [51] S. Masnou, Disocclusion: a variational approach using level lines, *IEEE Trans. Image Processing* 11 (2) (2002) 68–76.
- [52] C. Ballester, M. Bertalmio, V. Caselles, G. Sapiro, J. Verdera, Filling-in by joint interpolation of vector fields and gray levels, *IEEE Trans. Image Processing* 10 (8) (2001) 1200–1211.
- [53] M. Bertalmio, G. Sapiro, V. Caselles, C. Ballester, Image inpainting, in: *Siggraph 2000*, 2000, pp. 417–424.
- [54] D. Tschumperlé, R. Deriche, Vector-valued image regularization with PDEs: A common framework for different applications, *IEEE Trans. Pattern Anal. Mach. Intell* 27 (4) (2005) 506–517.

- [55] E. Candès, T. Tao, Near-optimal signal recovery from random projections: Universal encoding strategies?, *IEEE Transactions on Information Theory* 52 (12) (2006) 5406–5425.
- [56] D. Donoho, Compressed sensing, *IEEE Transactions on Information Theory* 52 (4) (2006) 1289–1306.
- [57] E. Candès, Compressive sampling, *Proceedings of the International Congress of Mathematicians, Madrid, Spain* (3) (2006) 1433–1452.