IMAGE DENOISING AND REGISTRATION BY PDE'S ON THE SPACE OF PATCHES

David Tschumperlé *Luc Brun*

Image Team, GREYC (CNRS UMR 6072), ENSICAEN/Université de Caen,
6 Bd du Maréchal Juin, 14050 Caen, FRANCE,
{*David.Tschumperle, Luc.Brun*}*@greyc.ensicaen.fr*

ABSTRACT

We define a simple transform able to map any multi-valued image into a space of patches, such that each existing image patch is mapped into a single high-dimensional point. We show that solving variational problems on this particular space is an elegant way of finding the natural patch-based counterparts of classical image processing techniques, such as the *Tikhonov* regularization and *Lucas-Kanade* registration methods. We end up with interesting variants of already known (non-variational) patch-based algorithms, namely the *Non Local Means* and *Block Matching* techniques. The interest of considering variational approaches on patch spaces is discussed and illustrated by comparison results with corresponding non-variational and non-patch methods.

1. INTRODUCTION

In the fields of Image Analysis, Processing and Synthesis, patch-based techniques generally meet with success. Defined as local square neighborhoods of image pixels, patches are very simple objects to work with, but they have the intrinsic ability to catch large-scale structures and textures present in natural images. Patches provide one of the simplest way to analyze, compare and copy textures, as soon as the considered patch sizes are higher than the so-called *texel* sizes (texture element, seen as the smallest significant unit of a texture). Moreover, patch-based methods are somehow intuitive : They indirectly reproduce the way humans are accomplishing some vision tasks by comparing semi-local image neighborhoods together. Proposed algorithms are often simple to implement, but they give surprisingly good results.

For instance, it has been long since patches have been used for solving the problem of estimating a displacement field between two images. This alignment problem can be classically (and partially) solved by the so-called *Block Matching* algorithm [18, 35], which consists in comparing each neighborhood (patch) of one image with all of the other one, finding the best match for each pixel location. Unfortunately, the regularity of the resulting motion field is not ensured by such a crude method and regularized techniques [5, 24] are sometimes needed. More recently, patch-based algorithms have been proposed to tackle the problem of synthesizing a texture similar to an input model [1, 16, 34]. It has been shown that this quite complicated task can be very well achieved by a simple iterated copy-paste procedure of different patches coming from the model, found to fit the best with the local neighborhoods of the image to synthesize. Some variants of these techniques have been also applied for transferring textures from an image to another one [2, 20], and for textured inpainting [14, 19], *i.e.*the reconstruction of missing textured regions in an image. Despite their relative simplicity, one must admit that these algorithms give outstanding results. It is also worth to cite the set of recent patch-based denoising methods, initiated with the *Non Local Means* scheme [13] and continued with various derivatives [4, 11, 12, 21]. Such methods are mainly based on an iterative weighted average of image patches. Here again, they have rapidly entered the hall of fame of image denoising techniques.

On the other side, variational methods for image processing [3, 8, 29] are known to be mathematically well-posed, flexible and competitive : They have the great ability to put complex a priori constraints (such as regularity constraints) on the obtained solutions. While regularity is desired for many applications (*e.g.*for motion estimation), it may sometimes avoid the reconstruction of textured solutions (*e.g.*for denoising/inpainting), textures being oscillating patterns by nature. Few attempts have been made to incorporate textured features in specific variational denoising/inpainting formulations by putting explicit patches [4, 9, 22, 25] or various image transforms [17] into the functionals to minimize, with few success. Mixing the performance of patches and the flexibility of variational methods is still a very exciting goal and an open challenge.

In this paper, we are looking toward a more general strategy to find patch-based counterparts of image processing techniques expressed as variational problems. Mainly, the idea lies on the construction of an alternate patch space on which the image to process/analyze is mapped. Thus, in this high-dimensional space, each existing image patch is represented by a single point. Variational formulations can be then defined and solved directly on the patch space, projecting back the obtained solution to the original image domain, if needed. So, instead of explicitly incorporating image patches into *ad-hoc* energy functionals, we rather consider the extension of existing classical (non-patch) functionals to a higher-dimensional space of patches. This is quite straightforward : Energy functionals are generally expressed with terms that can be easily extended for an arbitrary number of dimensions (*e.g.*gradients). As a result, we obtain algorithms which are the natural patch-based counterparts of “pointwise” variational techniques.

This paper is organized as follows : First, we define the reversible mapping of a multi-valued image into its patch space (section 2). Then, in this space, we consider the minimization of the *Tikhonov* regularization functional [10, 31] (section 3). We show that the natural patch-based version of the resulting minimizing flow can be interpreted as a variant of the *Non Local Means* filter [13]. In a second attempt, we tackle the problem of image alignment similarly, by minimizing an energy functional inspired by the *Lucas-Kanade* method [5, 24]. We end up with an interesting variational version of the *Block Matching* algorithm, implemented by the evolution of semi-local nonlinear PDE's. Application results and discussions on possible future applications of this general variational framework on patch spaces conclude this paper (sections 5-6).

2. DEFINITION OF THE PATCH SPACE

Let us consider a 2D multi-valued image $\mathbf{I}: \Omega \rightarrow \mathbb{R}^n$ defined on a continuous domain $\Omega \subset \mathbb{R}^2$. In this paper, we will mainly illustrate applications for $n = 3$, *i.e.*color images defined in the (R, G, B) color space. The i^{th} component of $\mathbf{I}$ is a scalar image, denoted by $I_i : \Omega \rightarrow \mathbb{R}$, so that $\forall (x, y) \in \Omega, \quad \mathbf{I}_{(x,y)} = \left(I_{1(x,y)}, I_{2(x,y)}, \ldots, I_{n(x,y)}\right)^T$. More generally, the i^{th} component of a vector $\mathbf{X}$ will be written X_i and the restriction of $\mathbf{X}$ to a set of consecutive components $i \ldots j$, as $\mathbf{X}_{i\ldots j}$.

Patch definition : The patch $\mathcal{P}^{\mathbf{I}}_{(x,y)}$ located at $(x, y) \in \Omega$ on the image $\mathbf{I}$ is defined as the set of all image values belonging to a spatially discretized local $p \times p$ (square) neighborhood of $\mathbf{I}$ centered at (x, y). The spatial discretization step, related to the analysis scale of the considered patches, is assumed to be 1 in order to simplify notations, even though any step is possible. The size p is considered as odd, *i.e.*$p = 2q + 1 \quad (q \in \mathbb{N}^*)$. Actually, a patch $\mathcal{P}^{\mathbf{I}}_{(x,y)}$ can be ordered in a np^2-dimensional vector as :

$$\mathcal{P}^{\mathbf{I}}_{(x,y)} = \big(I_{1(x-q,y-q)}, \ldots, I_{1(x+q,y+q)}, I_{2(x-q,y-q)}, \ldots, I_{n(x+q,y+q)}\big)^T$$

One may see $\mathcal{P}^{\mathbf{I}}_{(x,y)}$ as the concatenation of the patch vectors $\mathcal{P}^{I_i}_{(x,y)}$ for all image channels I_i, with

$$\mathcal{P}^{I_i}_{(x,y)} = \big(I_{i(x-q,y-q)}, \ldots, I_{i(x+q,y+q)}\big)^T$$

An interesting mathematical study of the manifold formed by the set of all these patches have been initiated in [27, 28]. This study is fascinating and share some ideas with the current paper.

Mapping to the patch space : We define the $(np^2 + 2)$-dimensional *patch space* $\Gamma = \Omega \times \mathbb{R}^{np^2}$. Each point $\mathbf{p} \in \Gamma$ is a high-dimensional vector whose coordinates may contain informations of any (x, y) location in Ω, as well as all values of any $p \times p$ patch $\mathcal{P} \in \mathbb{R}^{np^2}$. Obviously, in this patch space Γ, we want to highlight all the points of the form $\mathbf{p} = (x, y, \mathcal{P}^{\mathbf{I}}_{(x,y)})$, *i.e.*the points which precisely correspond to existing patches in $\mathbf{I}$. In this context, we call $\mathbf{p}$ a *located patch*. We define then a function $\tilde{\mathbf{I}}$ in Γ such that $\tilde{\mathbf{I}}_{(\mathbf{p})}$ is non zero only for these particular located patches of $\mathbf{I}$:

$$\tilde{\mathbf{I}} : \Gamma \rightarrow \mathbb{R}^{np^2+1}, \quad \text{s.a.} \quad \forall \mathbf{p} \in \Gamma,$$

$$\tilde{\mathbf{I}}_{(\mathbf{p})} = \begin{cases} ({\mathcal{P}^{\mathbf{I}}_{(x,y)}}^T, 1)^T & \text{if} \quad \mathbf{p} = (x, y, \mathcal{P}^{\mathbf{I}}_{(x,y)})^T \\ \vec{0} & \text{elsewhere} \end{cases} \tag{1}$$

Thus, we call the application $\mathcal{F}$ such that $\tilde{\mathbf{I}} = \mathcal{F}(\mathbf{I})$ a *patch transform*. Note that the value space of $\tilde{\mathbf{I}}$ has an extra component set to 1 for points at existing located image patches of $\mathbf{I}$. This dimension can be compared with the one introduced when dealing with projective spaces : It plays a role of weighting action when inverting the patch transform $\mathcal{F}$, *i.e.*retrieving back the multi-valued image $\mathbf{I}$ from $\tilde{\mathbf{I}}$. Intuitively, it defines how much a located patch in Γ is meaningful, and by default, all existing patches of the original image $\mathbf{I}$ have the same importance.

In this paper, we want to show that it is worth to solve variational problems on Γ, expressed as the minimization of energy functionals $E(\tilde{\mathbf{I}})$ rather than energies $E(\mathbf{I})$ on the original image domain Ω. For this purpose, we have to point out the fact that $\tilde{\mathbf{I}}$ is a highly discontinuous function. To avoid derivation problems of the energies E, we will work on a continuous approximation $\tilde{\mathbf{I}}_\epsilon$ of $\tilde{\mathbf{I}}$, where each original located patch $\mathbf{p}$ is not mapped into a single point, but into a normalized Gaussian function G_ϵ with a variance ϵ close to 0. From a mathematical point of view, this defines the approximation as $\tilde{\mathbf{I}}_\epsilon = \tilde{\mathbf{I}} * G_\epsilon$ which is $\mathcal{C}$-infinite. In the sequels, with a slight abuse of notations and when no confusions are possible, we will denote $\tilde{\mathbf{I}}_\epsilon$ simply by $\tilde{\mathbf{I}}$.

Back-projection on the image domain : Due to the high dimensionality of Γ, there are of course no unique ways to compute the inverse transform of the patch-based representation $\tilde{\mathbf{I}} = \mathcal{F}(\mathbf{I})$. We define a back-projection method based on two steps : First, we retrieve for every location $(x, y) \in \Omega$, the most significant located patch $\mathcal{P}^{\tilde{\mathbf{I}}}_{sig(x,y)}$ expressed in $\tilde{\mathbf{I}}$. It is found to be the one with the maximum weight, *i.e.*$\mathcal{P}^{\tilde{\mathbf{I}}}_{sig(x,y)} = \tilde{I}_{1\ldots np^2}(x, y, \mathcal{P}^{\tilde{\mathbf{I}}}_{max(x,y)})$, with

$$\mathcal{P}^{\tilde{\mathbf{I}}}_{max(x,y)} = \text{argmax}_{\mathbf{q}\in\mathbb{R}^{np^2}} \ \tilde{I}_{np^2+1}(x, y, \mathbf{q}^T) \tag{2}$$

Note that if one perturbs only slightly the patch transform $\tilde{\mathbf{I}}$ of an image $\mathbf{I}$, there are very good chances to find the most significant perturbed patches at the same locations as the original ones, *i.e.* $\mathcal{P}^{\tilde{\mathbf{I}}}_{max(x,y)} = \mathcal{P}^{\mathbf{I}}_{(x,y)}$, even though the pixel values of the modified patches may be different, *i.e.* $\mathcal{P}^{\tilde{\mathbf{I}}}_{sig(x,y)} \neq \mathcal{P}^{\mathbf{I}}_{(x,y)}$.

In a second step, the back-projected image $\mathbf{I}$ is reconstructed by combining these most significant patches together. Several strategies are possible : Here, we will use the simplest one, which simply consists in copying the center pixel of each significant patch $\mathcal{P}^{\tilde{\mathbf{I}}}_{sig(x,y)}$ at its corresponding known location (x, y), normalizing it by its weight : $\forall i \in [1, n], \;\; \forall (x, y) \in \Omega$,

$$I_{i(x,y)} = \frac{\tilde{I}_{ip^2+\frac{p^2+1}{2}}(x, y, \mathcal{P}^{\tilde{\mathbf{I}}}_{max(x,y)})}{\tilde{I}_{np^2+1}(x, y, \mathcal{P}^{\tilde{\mathbf{I}}}_{max(x,y)})} \quad (3)$$

More generally, we could have copied an entire sub-patch of $\mathcal{P}^{\tilde{\mathbf{I}}}_{sig}$ at (x, y), while blending overlapped neighborhood patches according to their relative weights. This consideration on how to copy image patches appears frequently in the literature related to patch-based methods, and no "best" solutions have appeared yet. It is interesting to see that in our case, this choice is clearly delimited as a part of an inverse transform only.

Now that the direct and inverse patch transforms are clearly defined, we study the application of some variational methods on the patch space Γ rather than on the original image domain Ω, for image denoising (section 3) and registration (section 4).

3. IMAGE DENOISING BY PATCH-BASED TIKHONOV REGULARIZATION

Suppose we have a multi-valued image $\mathbf{I}^{noisy} : \Omega \rightarrow \mathbb{R}^n$ corrupted by some kind of noise. So will be its patch transform $\tilde{\mathbf{I}}^{noisy}$. We are looking for a patch-based minimizing flow able to regularize $\tilde{\mathbf{I}}^{noisy}$ rather than process directly $\mathbf{I}^{noisy}$. For this purpose, we minimize the following energy E_1, classically denoted as the *Tikhonov* regularization functional, which have been simply extended to the high-dimensional space Γ :

$$E_1(\tilde{\mathbf{I}}) = \int_{\Gamma} \|\nabla \tilde{\mathbf{I}}_{(\mathbf{p})}\|^2 \; d\mathbf{p} \quad (4)$$

where $\|\nabla \tilde{\mathbf{I}}_{(\mathbf{p})}\| = \sqrt{\sum_{i=1}^{np^2+1} \|\nabla \tilde{I}_{i(\mathbf{p})}\|^2}$ is the habitual extension of the gradient norm for multi-valued datasets [15]. Note that this multi-valued gradient includes also the gradient $\|\nabla \tilde{I}_{np^2+1}\|$ of the patch weights.

Minimizing flow : The PDE flow that minimizes (4) is found by the derivation of $E_1(\tilde{\mathbf{I}})$ using the Euler-Lagrange equations and by the expression of the corresponding gradient descent algorithm. It leads to the well known *heat flow* equation, which is in our case performed on the high-dimensional patch space Γ :

$$\begin{cases} \tilde{\mathbf{I}}_{[t=0]} = \tilde{\mathbf{I}}^{noisy} \\ \frac{\partial \tilde{I}_i}{\partial t} = \Delta \tilde{I}_i \end{cases} \quad (5)$$

where Δ is the Laplacian operator on Γ.

Similarly to denoising techniques based on classical diffusion PDE's [3, 6, 7, 29, 30, 33], we are not particularly interested by the steady-state solution of (5), since it would roughly consist of a constant solution. We are rather looking for a solution of this multi-dimensional heat flow at a particular *finite time* t_1. It has been already proven [23] that this solution is the convolution of the initial estimate $\tilde{\mathbf{I}}^{noisy}$ with a normalized Gaussian kernel G_σ with a standard deviation $\sigma = \sqrt{2\,t_1}$. Here, this convolution has to be done on the high-dimensional patch space Γ :

$$\tilde{\mathbf{I}} = \tilde{\mathbf{I}}^{noisy} * G_\sigma, \;\; \text{with}$$

$$\forall \mathbf{p} \in \Gamma, \;\; G_{\sigma(\mathbf{p})} = \frac{1}{(2\pi\sigma^2)^{\frac{np^2+2}{2}}} \; e^{-\frac{\|\mathbf{p}\|^2}{2\sigma^2}} \quad (6)$$

This is a quite straightforward way of defining the patch-based counterpart of the *Tikhonov* regularization process.

Interpretation in the image domain : Interesting things arise when one tries to interpret the result of this patch-based heat flow (5) in the original image domain Ω. In fact, the patch transform (1) tells us that $\tilde{\mathbf{I}}^{noisy}$ vanishes almost everywhere in Γ, excepted on the points $\mathbf{p} = (x, y, \mathcal{P}^{\mathbf{I}^{noisy}}_{(x,y)})^T$ corresponding to the locations of the original image patches in $\mathbf{I}^{noisy}$. So, the convolution (6) can be simplified as : $\forall (x, y, \mathcal{P}) \in \Gamma$,

$$\tilde{\mathbf{I}}_{(x,y,\mathcal{P})} = \int_{\Omega} \tilde{\mathbf{I}}^{noisy}_{(p,q,\mathcal{P}^{\mathbf{I}^{noisy}}_{(p,q)})} \; G_{\sigma(p-x,q-y,\mathcal{P}^{\mathbf{I}^{noisy}}_{(p,q)}-\mathcal{P})} \; dp\; dq$$

We also notice that the locations of the most significant patches in $\tilde{\mathbf{I}}$, as defined in (2), will be the same as the ones in $\tilde{\mathbf{I}}^{noisy}$, since the convolution (6) by a Gaussian kernel leaves the maxima of the patch weights $\tilde{I}_{np^2+1}$ unchanged [23]. So, the inverse patch transform (3) of $\tilde{\mathbf{I}}$ has a simple closed-form expression in Ω, and can be written directly from $\mathbf{I}^{noisy}$, considering that we have

$$\forall (x, y) \in \Omega, \quad \mathcal{P}^{\tilde{\mathbf{I}}}_{max(x,y)} = \mathcal{P}^{\tilde{\mathbf{I}}^{noisy}}_{max(x,y)} = \mathcal{P}^{\mathbf{I}_{noisy}}_{(x,y)}$$

as well as $\tilde{I}^{noisy}_{ip^2+\frac{p^2+1}{2}}(x, y, \mathcal{P}^{I_{noisy}}_{(x,y)}) = I^{noisy}_{i(x,y)}$,

and $\tilde{I}^{noisy}_{np^2+1(x,y,\mathcal{P}^{I_{noisy}}_{(x,y)})} = 1$.

Using these relations together with (3), and after few calculus, we end up with :

$$\forall (x, y) \in \Omega, \quad \mathbf{I}_{(x,y)} = \frac{\int_{\Omega} \mathbf{I}^{noisy}_{(p,q)} \; w_{(x,y,p,q)} dp\; dq}{\int_{\Omega} w_{(x,y,p,q)} \; dp\; dq} \quad (7)$$

with $w_{(x,y,p,q)} = w^s_{(x,y,p,q)} \; w^p_{(x,y,p,q)}$ and

$$w^s_{(x,y,p,q)} = \frac{1}{2\pi\sigma^2} \; e^{-\frac{(x-p)^2+(y-q)^2}{2\sigma^2}}$$

$$w^p_{(x,y,p,q)} = \frac{1}{(2\pi\sigma^2)^{\frac{np^2}{2}}} \; e^{-\frac{\|\mathcal{P}^{\mathbf{I}^{noisy}}_{(x,y)} - \mathcal{P}^{\mathbf{I}^{noisy}}_{(p,q)}\|^2}{2\sigma^2}} \quad (8)$$

Thus, each pixel value $\mathbf{I}_{(x,y)}$ of the regularized image is the result of a weighted averaging of all noisy pixels $\mathbf{I}^{noisy}_{(p,q)}$, the weight depending both on the spatial distance between points (x, y) and (p, q) (first term in (8)), as well as the similarity between corresponding patches centered on (x, y) and (p, q) (second term in (8)). Of course, this interpretation as a filtering process in Ω greatly simplifies the implementation step since it naturally suggests an algorithm in a lower-dimensional space Ω. But we have to keep in mind that it is actually the outcome of a classical *Tikhonov* minimizing flow on the patch space Γ. Note that the Gaussian functions in the spatial and patch-based terms of the weighing function $w_{(x,y,p,q)}$ (8) have the same standard deviation σ. Having different σ can be easily simulated by pre-multiplying the image $\mathbf{I}^{noisy}$ by a ratio factor λ before processing, being then equivalent than having two different standard deviations $\sigma_{patch} = \sigma_{spatial}/\lambda$.

Link with other filtering methods : The patch-based *Tikhonov* regularization method (7) is actually quite similar to the *Non Local Means* technique, as defined in [13]. Differences are twofold : First, it is naturally defined as the solution of a *minimizing flow* acting on multi-valued images, while the original formulation has been expressed as an explicit nonlinear filter acting on scalar images. Second, the averaging weights $w_{i,j}$ of the *Non Local Means* algorithm are only based on the similarity between patches. Here, the weighting function (8) also considers the spatial distances between patches to average. Also, in the extreme case when the patch size p is reduced to 1 (so, a patch is just one point), the regularization method (7) becomes the natural multi-valued version of another well know nonlinear image filter, namely the *Bilateral Filtering* method [6, 7, 26, 32]. These filters are known to be anisotropic in the image domain Ω. The benefit of minimizing the *Tikhonov* functional on the patch space Γ rather than on the image domain Ω is obvious in terms of regularization quality : It avoids the typical isotropic smoothing behavior of the *Tikhonov* regularization that usually over-smoothes the important image structures, such as the edges, corners and textures. A comparative figure of all these regularization algorithms is shown and commented in section 5.

4. IMAGE REGISTRATION BY PATCH-BASED VARIATIONAL METHOD

Let us now consider the problem of estimating a displacement field $\mathbf{u} : \Omega \to \mathbb{R}^2$ between two multi-valued images $\mathbf{I}^{t_1} : \Omega \to \mathbb{R}^n$ (the *reference* image) and $\mathbf{I}^{t_2}$ (the *target* image). This image alignment problem can be typically solved by a variational method, which aims at finding the vector field $\mathbf{u}$ that minimizes the following energy :

$$E_2(\mathbf{u}) = \int_\Omega \alpha\, \|\nabla \mathbf{u}_{(\mathbf{p})}\|^2 + \|\mathcal{D}_{(\mathbf{p},\mathbf{p}+\mathbf{u})}\|^2 \, d\mathbf{p} \quad (9)$$

The user-defined parameter $\alpha \in \mathbb{R}^+$ imposes a regularity constraint on the estimated vector field $\mathbf{u}$, if it chosen to be non zero. $\mathcal{D}_{(\mathbf{p},\mathbf{q})}$ is a measure of the *dissimilarity* between image pixels $\mathbf{I}^{t_1}_{(\mathbf{p})}$ and $\mathbf{I}^{t_2}_{(\mathbf{q})}$. Lot of different expressions for $\mathcal{D}$ have been already proposed in the literature [3, 5, 24]. One of the most common choice makes the assumption that $\mathbf{I}^{t_1}$ and $\mathbf{I}^{t_2}$ are acquired under the same global illumination conditions (for instance, they are successive frames of a video sequence) and then, that the constraint of the *Brightness Consistency* holds.
In this case, $\mathcal{D}$ can be reasonably chosen to be :

$$\mathcal{D}_{1(\mathbf{p},\mathbf{q})} = \mathbf{I}^{t_1}_{\sigma(\mathbf{p})} - \mathbf{I}^{t_2}_{\sigma(\mathbf{q})} \quad (10)$$

where $\mathbf{I}^{t_k}_\sigma = \mathbf{I}^{t_k} * G_\sigma$ are filtered versions of the images $\mathbf{I}^{t_k}$, convolved by a normalized Gaussian kernel G_σ. It allows the consideration of regularized reference and target images instead of possibly noisy ones. The choice $\mathcal{D} = \mathcal{D}_1$ in (9) leads to a variational extension of the *Lucas-Kanade* method [5, 24] for multi-valued images.

Dissimilarity measure on the patch space : The dissimilarity measure $\mathcal{D}_1$ can be naturally extended to deal with patches, by expressing it with the patch transforms $\tilde{\mathbf{I}}^{t_1}$ and $\tilde{\mathbf{I}}^{t_2}$ instead of the original images $\mathbf{I}^{t_1}$ and $\mathbf{I}^{t_2}$. This is quite straightforward : We just replace the pointwise intensities by the most significant patches (2) present in the regularized versions of the patch transforms $\tilde{\mathbf{I}}^{t_k}_\sigma = \tilde{\mathbf{I}}^{t_k} * G_\sigma$:

$$\mathcal{D}_{2(\mathbf{p},\mathbf{q})} = \tilde{\mathbf{I}}^{t_1}_{\sigma(\mathbf{p},\mathcal{P}^{\tilde{\mathbf{I}}^{t_1}}_{max(\mathbf{p})})} - \tilde{\mathbf{I}}^{t_2}_{\sigma(\mathbf{q},\mathcal{P}^{\tilde{\mathbf{I}}^{t_2}}_{max(\mathbf{q})})} \quad (11)$$

Here, the smoothed patch transforms $\tilde{\mathbf{I}}^{t_k}_\sigma$ represent edge-preserving filtered versions of the reference and target images $\mathbf{I}^{t_k}$, as mentioned in section 3. It means also that the localization of the most significant patches is known to be $\mathcal{P}^{\tilde{\mathbf{I}}^{t_k}}_{\sigma max} = \mathcal{P}^{\mathbf{I}^{t_k}}$.
In the particular case where $\sigma = 0$, the dissimilarity function $\mathcal{D}_2$ is the same as the one used for the classical *Block Matching* method. Moreover, when $\alpha = 0$ (no regularization constraints are considered on $\mathbf{u}$), the *Block Matching* is a global minimizer of (9).

Minimizing flow : In a more general setting, the Euler-Lagrange derivation of (9) gives the set of coupled PDE's which locally minimizes the energy functional E_2 :
$\forall j \in [1, 2], \forall \mathbf{x} \in \Omega$,

$$\begin{cases} \mathbf{u}_{[t=0]} = \vec{0} \\ \dfrac{\partial u_{j(\mathbf{x})}}{\partial t} = \alpha\, \Delta u_j + \\ \displaystyle\sum_{i=1}^{np^2+1} \left(\tilde{I}^{t_1}_{\sigma i(\mathbf{x},\mathcal{P}^{\mathbf{I}^{t_1}}_{(\mathbf{x})})} - \tilde{I}^{t_2}_{\sigma i(\mathbf{x}+\mathbf{u},\mathcal{P}^{\mathbf{I}^{t_2}}_{(\mathbf{x}+\mathbf{u})})} \right) [\nabla \mathcal{G}_i]_{j(\mathbf{x}+\mathbf{u})} \end{cases} \quad (12)$$

where $\mathcal{G}_{i(\mathbf{x})} = \tilde{I}^{t_2}_{\sigma i(\mathbf{x},\mathcal{P}^{\mathbf{I}^{t_2}}_{(\mathbf{x})})}$. Here again, this minimizing flow (12) can be implemented directly in the image domain Ω, without having to explicitly compute and store the patch transforms $\tilde{\mathbf{I}}^{t_k}_\sigma$. Indeed, it requires only the retrieval of the most significant patches in $\tilde{\mathbf{I}}^{t_k}_\sigma$, which can be proven to be $\mathcal{P}^{\tilde{\mathbf{I}}^{t_k}}_{\sigma sig} = \mathcal{P}^{\mathbf{I}^{t_k}_{regul}}$, where $\mathbf{I}^{t_k}_{regul}$ is the nonlinear filtered version of $\mathbf{I}^{t_k}$ by the *Tikhonov* regularization flow on the patch space (7).

This patch-based registration PDE (12) is a local minimizer of a *Block Matching*-like objective function. Differences are twofold : First, it *implicitly* considers edge-preserving filtered versions of the reference and target images, instead of isotropically smoothed ones. But most of all, it is able to put an important *a priori* smoothness constraint on the estimated field $\mathbf{u}$. This method combines then the significance of the patch description for the matching of image structures, while keeping the aptitude of the variational methods to impose useful constraints. These interesting properties are discussed and illustrated with a comparative figure in section 5.

5. APPLICATION RESULTS

We applied the different variational techniques on the patch space presented in this paper, on different color images considered in their original (R, G, B) color space.

Color image denoising : The *Tikhonov* regularization flow (7) on the patch space Γ can be used to enhance degraded color images (or other multi-valued datasets). Fig.1 compares it with the most connected algorithms, namely the *Non Local Means*, the *Bilateral Filtering* and the classical *Tikhonov* regularization performed on the image domain Ω. Synthetic white Gaussian noise ($\sigma_{noise} = 20$) has been added to the original color image *Barbara*. For the honesty of the comparison, we have not applied the scalar versions of these related filters as defined in the papers [13, 32], but their multi-valued extensions. It gives indeed better denoising results than applying them channel by channel. The PSNR between the noise-free and restored images, as well as the parameters used for the experiments are displayed. For each method, the parameters have been manually tuned to optimize the obtained PSNR. As our proposed flow (7) is actually very close in its final expression to the *Non Local Means* and *Bilateral* filters, the denoising results are of course very comparable in terms of quality. It seems to perform a little bit better on certain regions (some edges are sharper in the zoomed part). We believe it is due to the fact that our weighting function (8) considers both the similarity as well as the spatial distances between patches. The visual improvement is very subtle anyway. Additional results of our patch-based regularization technique are illustrated on Fig.2. It is interesting to notice that in spite of its intrinsic isotropic nature, the *Tikhonov* flow on the patch space (5) clearly competes against sophisticated anisotropic diffusion PDE's for the smoothing of multi-valued images, such as the method proposed in [33].

Color image registration : The patch-based registration PDE's (12) allow to estimate a smooth displacement field between two color images while considering a patch-based dissimilarity measure. Such a displacement field can be used then to generate a morphing video sequence, where new synthetic frames are added between the reference and target images (*i.e.*a *re-timing* process, where the reference frame corresponds to $t = 0$ and the target to $t = 1$). Fig.3 shows a comparison of the re-timed results obtained by our proposed registration method (12) and by the most related ones, *i.e.*the *Block Matching* algorithm and the pointwise version (10) of the variational registration technique (9). For each method, the estimated displacement field $\mathbf{u}$ is displayed as a colored image (first column). The mid-interpolated frame (*i.e.*reconstructed at $t = 0.5$) and a zoomed part of it are shown on the two last columns. This is indeed the hardest frame to reconstruct since it classically exhibits the largest ghosting artefacts due to motion estimation errors. To be able to handle large displacements, we considered a classical multiscale evolution scheme for the PDE's that minimize the pointwise and patch-based versions of the energy functional (9).
Without any surprise, the original *Block Matching* algorithm fails in reconstructing a smooth motion field, resulting in a unpleasant mid-interpolated frame reconstruction. The pointwise version of the variational registration technique (9) is quite competitive when the smoothness constraint is high enough ($\alpha = 0.1$), but fails otherwise. Surprisingly, our proposed patch-based variational formulation (9) with the dissimilarity measure (11) outputs a quite good reconstructed frame, even when $\alpha = 0$. Considering a patch-based dissimilarity measure seems to have a natural tendency to intrinsically regularize the registration problem. Anyway, best visual results are obtained with our variational patch-based registration method while considering regularization constraints on the displacement field $\mathbf{u}$ (last row).

6. CONCLUSIONS & PERSPECTIVES

In this paper, a reversible patch-based transform of multi-valued images has been presented. More important, we illustrated the fact that considering variational methods in the corresponding patch space is a very nice way of designing the natural patch-based counterparts of classical image processing techniques. The obtained experimental results seems to confirm this impression. We strongly believe that this effort is a first step in reconciling both patch and variational worlds and that our proposed framework can be used to translate many other "pointwise" variational formulations into their natural patch-based counterparts.

7. REFERENCES

[1] Ashikhmin, M.: Synthesizing Natural Textures. Symposium on Interactive 3D Graphics. p.217–226, 2001.

[2] Ashikhmin, M.: Fast Texture Transfer. IEEE Computer Graphics and Applications, Vol.23 (4), p.38–43, 2003.

[3] Aubert, G., Kornprobst, P.: Mathematical Problems in Image Processing: PDE's and the Calculus of Variations. Applied Mathematical Sciences, Vol. 147, Springer-Verlag, 2006.

[4] Azzabou, N., Paragios, N., Guichard, F., Cao, F.: Variable bandwidth image denoising using image-based noise models. IEEE CVPR, 2007.

[5] Baker, S., Matthews, I.: Lucas-Kanade 20 Years On: A Unifying Framework. International Journal of Computer Vision, Vol.56 (3), p.221–255, 2004.

[6] Barash, D.: A fundamental relationship between bilateral filtering, adaptive smoothing, and the nonlinear diffusion equation. IEEE PAMI, Vol.24 (6), p.844–847, 2002.

[7] Barash, D., Comaniciu, D.: A common framework for Nonlinear Diffusion, Adaptive Smoothing, Bilateral Filtering and Mean Shift. Image Vision Comp., Vol.22 (1), 2004.

[8] Bertalmío, M., Sapiro, G., Caselles, V., Ballester, C.: Image inpainting. SIGGRAPH, Computer Graphics Proceedings, p.417–424, 2000.

[9] Bertalmío, M., Vese, L.A., Sapiro, G., Osher, S.: Simultaneous Structure and Texture Image Inpainting. IEEE CVPR, p.707–712, 2003.

[10] Bouhamidi, A., Jbilou, K.: Sylvester Tikhonov-Regularization Methods in Image Restoration. J. Comput. Appl. Math. Vol.206 (1), p.86–98, 2007.

[11] Boulanger, J., Kervrann, C., Bouthemy, P.: Space-time Adaptation for Patch-based Image Sequence Restoration. IEEE PAMI. Vol.29 (6), p.1096–1102, 2007.

[12] Brox, T., Cremers, D.: Iterated Nonlocal Means for Texture Restoration. SSVM. Vol.4485 of Lecture Notes in Computer Science, Springer p.13–24, 2007.

[13] Buades, A., Coll, B., Morel, J.M.: A Non-Local Algorithm for Image Denoising. IEEE CVPR, p.60–65, 2005.

[14] Criminisi, A., Pérez, P., Toyama, K.: Region Filling and Object Removal by Exemplar-based Image Inpainting. IEEE Trans. on Image Proc., Vol.13 (9).

[15] Zenzo, S.D.: A Note on the Gradient of a Multi-image. Computer Vision, Graphics, and Image Processing, Vol.33 (1), p.116–125, 1986.

[16] Efros, A.A., Leung, T.K.: Texture Synthesis by Non-Parametric Sampling. IEEE ICCV, Vol.2, p.1033–1038, 1999.

[17] Fadili, M., Starck, J.L., Boubchir, L.: Morphological Diversity and Sparse Image Denoising. IEEE ICASSP, Vol.1, p.589–592, 2007.

[18] Haworth, C., Peacock, A.M., Renshaw, D.S.: Performance of Reference Block Updating Techniques when Tracking with the Block Matching Algorithm. IEEE ICIP, Vol.1, p.365–368, 2001.

[19] Hays, J., Efros, A.A.: Scene Completion using Millions of Photographs. ACM Trans. Graph., Vol.26 (3), 2007.

[20] Hertzmann, A., Jacobs, C.E., Oliver, N., Curless, B., Salesin, D.: Image Analogies. SIGGRAPH, Computer Graphics Proceedings, p.327–340, 2001.

[21] Kervrann, C., Boulanger, J.: Optimal Spatial Adaptation for Patch-based Image Denoising. IEEE Transactions on Image Processing, Vol.15 (10), p.2866–2878, 2006.

[22] Kindermann, S., Osher, S., Jones, P.: Deblurring and Denoising of Images by Nonlocal Functionnals. SIAM Multiscale Modeling and Simulation, Vol.4, p.1091–1115, 2005.

[23] Koenderink, J.: The Structures of Images. Biological Cybernetics, Vol.50, p.363–370, 1984.

[24] Lucas, B.D., Kanade, T.: An Iterative Image Registration Technique with an Application to Stereo Vision. Hayes, P.J., Ed.: IJCAI, William Kaufmann, p.674–679, 1981.

[25] Osher, S., Sole, A., Vese, L.: Image Decomposition and Restoration using Total Variation Minimization and the H-1 Norm. UCLA CAM Report, Vol.57, 2002.

[26] Paris, S., Durand, F.: A Fast Approximation of the Bilateral Filter using a Signal Processing Approach. ECCV Vol.3954 (4), LNCS, Springer, p.568–580, 2006.

[27] Peyré, G.: Image Processing with Non-local Spectral Bases. Research Report, Ceremade, 2007.

[28] Peyré, G.: Manifold Models for Signals and Images. Research Report, Ceremade, No.2007-14, 2007.

[29] Sapiro, G.: Geometric Partial Differential Equations and Image Analysis. Cambridge University Press, 2001.

[30] Sochen, N.A., Zeevi, Y.Y.: Representation of Colored Images by Manifolds Embedded in Higher Dimensional Non-Euclidean Space. ICIP, Vol.1, p.166–170, 1998.

[31] Tikhonov, A.: Regularization of Incorrectly Posed Problems. Soviet Math. Vol.4, p.1624–1627, 1963.

[32] Tomasi, C., Manduchi, R.: Bilateral Filtering for Gray and Color images. ICCV. p.839–846, 1998.

[33] Tschumperlé, D., Deriche, R.: Vector-valued Image Regularization with PDE's: A Common Framework for Different Applications. IEEE PAMI, Vol.27 (4), p.506–517, 2005.

[34] Wei, L.Y., Levoy, M.: Fast Texture Synthesis using Tree-structured Vector Quantization. SIGGRAPH, Computer Graphics Proceedings, p.479–488, 2000.

[35] Wiegand, T., Sullivan, G., Bjntegaard, G., Luthra, A.: Overview of the H264 Video Coding Standard. IEEE Trans. on Circuits and Systems for Video Tech., 13(7), p.560–576, 2003.

(a) Noisy image (std=20) (b) *Tikhonov* regularization on Ω (c) *Bilateral Filtering* (d) *Non Local Means*

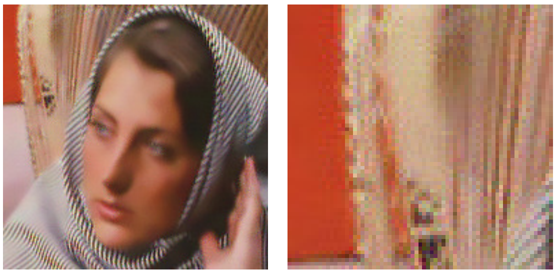

(e) *Tikhonov* regularization on the patch space Γ

Result	PSNR	Parameters	Reference
(b)	25.33 dB	$\sigma = 0.6$	[10, 31]
(c)	26.97 dB	$\sigma_s = 10, \sigma_r = 25$	[6, 26, 32]
(d)	26.98 dB	$\sigma = 21$, patch 7×7	[13]
(e)	27.03 dB	$\sigma = 21$, patch 7×7	.

(f) PSNR(Original/Denoised) and parameters used for experiments

Figure 1. Illustration of the *Tikhonov* regularization on the patch space (7), and comparisons with most related filtering techniques.

(a) Original degraded image (b) *Tikhonov* regularization on the patch space Γ

(a) Original degraded image (b) Using anisotropic diffusion PDE's [33] (c) Using *Tikhonov* on the patch space Γ

Figure 2. Additional results obtained by the *Tikhonov* regularization flow on the patch space (7).

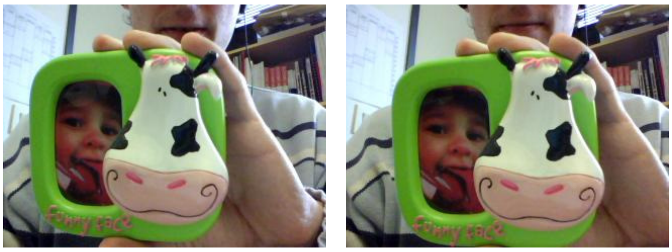

(a) Reference and target color images

(b) 9×9 *Block Matching* registration

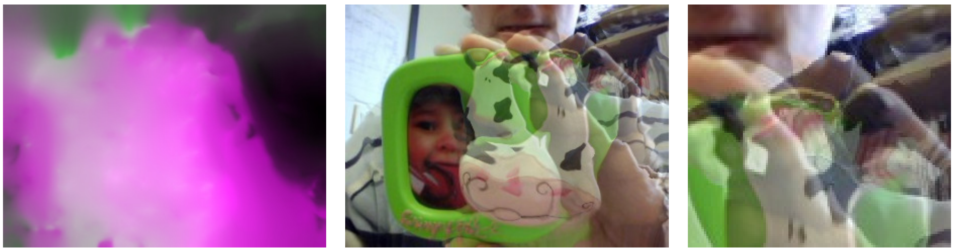

(c) Variational pointwise registration with smoothness constraints ($\alpha = 0.01$)

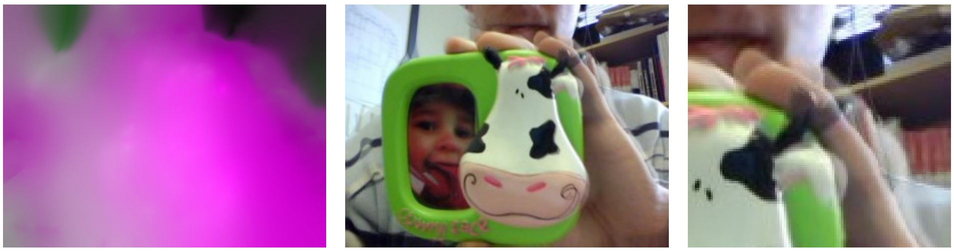

(d) Variational pointwise registration with smoothness constraints ($\alpha = 0.1$)

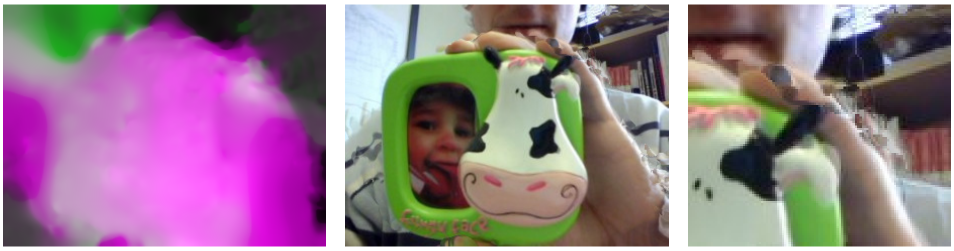

(e) Variational patch-based registration without smoothness constraints ($\alpha = 0$)

(f) Variational patch-based registration with smoothness constraints ($\alpha = 0.01$)

Figure 3. Illustration of the variational image registration method on the patch space (12), and comparisons with *Block Matching* [18, 35] and pointwise variational registration (9) methods.