



HAL
open science

Hierarchical super-resolution-based inpainting

Olivier Le Meur, Mounira Ebdelli, Christine Guillemot

► **To cite this version:**

Olivier Le Meur, Mounira Ebdelli, Christine Guillemot. Hierarchical super-resolution-based inpainting. IEEE Transactions on Image Processing, 2013, 22 (10), pp.3779-3790. 10.1109/TIP.2013.2261308 . hal-00876168

HAL Id: hal-00876168

<https://inria.hal.science/hal-00876168v1>

Submitted on 21 Nov 2014

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.

Hierarchical super-resolution-based inpainting

Olivier Le Meur, Mounira Ebdelli, Christine Guillemot

Abstract—This paper introduces a novel framework for exemplar-based inpainting. It consists in performing first the inpainting on a coarse version of the input image. A hierarchical super-resolution algorithm is then used to recover details on the missing areas. The advantage of this approach is that it is easier to inpaint low-resolution pictures than high resolution ones. The gain is both in terms of computational complexity and visual quality. However, to be less sensitive to the parameter setting of the inpainting method, the low-resolution input picture is inpainted several times with different configurations. Results are efficiently combined with a loopy belief propagation and details are recovered by a single-image super-resolution algorithm. Experimental results in a context of image editing and texture synthesis demonstrate the effectiveness of the proposed method. Results are compared to five state-of-the-art inpainting methods.

Index Terms—Exemplar-based inpainting, single-image super-resolution

I. INTRODUCTION

IMAGE inpainting refers to methods which consist in filling in missing regions (holes) in an image [1]. Existing methods can be classified into two main categories.

The first category concerns diffusion-based approaches which propagate linear structures or level lines (so-called isophotes) via diffusion based on partial differential equations [1], [2] and variational methods [3]. The diffusion-based methods tend to introduce some blur when the hole to be filled in is large. The second family of approaches concerns exemplar-based methods which sample and copy best matching texture patches from the known image neighbourhood [4], [5], [6], [7]. These methods have been inspired from texture synthesis techniques [8] and are known to work well in cases of regular or repeatable textures. The first attempt to use exemplar-based techniques for object removal has been reported in [6]. The authors in [5] improve the search for similar patches by introducing an a priori rough estimate of the inpainted values using a multi-scale approach which then results in an iterative approximation of the missing regions from coarse-to-fine levels. The two types of methods (diffusion- and exemplar-based) can be efficiently combined, e.g. by using structure tensors to compute the priority of the patches to be filled in as in [9]. A recent approach [10] combines an exemplar-based approach with super-resolution. It is a two-steps

algorithm. First a coarse version of the input picture is inpainted. The second step consists in creating an enhanced resolution picture from the coarse inpainted image.

Although tremendous progress has been made in the past years on exemplar-based inpainting, there still exists a number of problems. We believe that the most important one is related to the parameter settings such as the filling order and the patch size.

This problem is here addressed by considering multiple inpainted versions of the input image. To generate this set of inpainted pictures, different settings are used. The inpainted pictures are then combined yielding the final inpainted result. Notice that the inpainting algorithm is preferably applied on a coarse version of the input image; this is particularly interesting when the hole to be filled in is large. This provides the advantage to be less demanding in terms of computational resources and less sensitive to noise and local singularities. In this case the final full resolution inpainted image is recovered by using a super-resolution (SR) method similarly to [10]. SR methods refer to the process of creating one enhanced resolution image from one or multiple input low resolution images. These problems are then referred to as single or multiple images SR, respectively. In both cases, the problem is of estimating high frequency details which are missing in the input image(s). The proposed SR-aided inpainting method falls within the context of single-image SR.

The SR problem is ill-posed since multiple high-resolution images can produce the same low-resolution image. Solving the problem hence requires introducing some prior information. The prior information can be an energy functional defined on a class of images which is then used as a regularization term together with interpolation techniques [11]. This prior information can also take the form of example images or of corresponding LR-HR (Low Resolution - High Resolution) pairs of patches learned from a set of un-related training images [12] or from the input low resolution image itself [13]. This latter family of approaches is known as exemplar-based SR methods [12]. An exemplar-based super-resolution method embedding K nearest neighbours found in an external patch database has also been described in [14]. Instead of constructing the LR-HR pairs of patches from a set of un-related training images, the authors in [13] extract these correspondences by searching for matches across different scales of a multi-resolution pyramid constructed from the input low-resolution image.

The proposed method builds upon the super-resolution-based inpainting method proposed in [10] which is based on exemplar-based inpainting (in particular Criminisi-like approach [4]) and single-image exemplar-based super-

Copyright (c) 2013 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

O. Le Meur is with the University of Rennes 1, France. WebPage: http://people.irisa.fr/Olivier.Le_Meur/

M. Ebdelli is with INRIA, France. E-mail: mounira.ebdelli@inria.fr

C. Guillemot is with INRIA, France. WebPage: <http://people.rennes.inria.fr/Christine.Guillemot/>




Fig. 1. The framework of the proposed method.

resolution [13]. The main novelty of the proposed algorithm is the combination of multiple inpainted versions of the input picture. The rationale behind this approach is to cope with the sensitivity of exemplar-based algorithms to parameters such as the patch size and the filling order. Different combinations have been tested and compared. Besides this major point, different adjustments regarding exemplar-based inpainting and SR methods are described such as the use of the coherence measure to constrain the candidate search [15].

In summary, the proposed method improves on the state-of-the-art exemplar-based inpainting methods by proposing a new framework involving a combination of multiple inpainting versions of the input picture followed by a single-image exemplar-based SR method. Notice that the SR method is used only when the inpainting method is applied on a low resolution of the input picture.

The paper is organized as follows. In Section II, an overview of the proposed exemplar-based inpainting algorithm is presented. In Section III, the details of the inpainting algorithm as well as the combination of the inpainted pictures are given. Section IV presents the super-resolution method. Experiments and comparisons with state-of-the-art algorithms are performed in Section V. Finally we conclude this work in Section VI.

II. NOTATIONS AND ALGORITHM OVERVIEW

A. Notations

The following notations are used throughout this paper:

- \mathbf{I} is a color input 2D image:

$$\mathbf{I} : \begin{cases} \Omega \subset \mathcal{R}^n \rightarrow \mathcal{R}^m \\ \mathbf{x} \rightarrow \mathbf{I}(\mathbf{x}) \end{cases} \quad (1)$$

where $n = 2$ and in this case $\mathbf{x} = (x, y)$ represents a vector indicating spatial coordinates of a pixel $p_{\mathbf{x}}$. Ω is the generic definition domain of images. \mathbf{I} is a color image composed of three components ($m = 3$). Here, we consider the (R, G, B) color space.

- $\mathbf{I}_i : \Omega \rightarrow \mathcal{R}$ represents the i^{th} image channel of \mathbf{I} .
- the definition domain Ω is here composed of two parts: $\Omega = S \cup T$, S being the known part of \mathbf{I} (source region) and T the target one or unknown parts of \mathbf{I} .
- $\hat{\mathbf{I}}^{(c)}$ is the inpainted picture obtained by using the c^{th} setting configuration. The setting refers to the patch's

size, the size of the search window, the picture's orientation, to name a few. We consider that c is in the range 1 to M .

- $\hat{\mathbf{I}}^{(*)}$ is the final inpainted picture.
- We denote by $p_{\mathbf{x}}^{(k)}$ the pixel located at \mathbf{x} in the image $\hat{\mathbf{I}}^{(k)}$.
- $\psi_{p_{\mathbf{x}}}$ is the patch centered on the pixel $p_{\mathbf{x}}$.
- $\psi_{p_{\mathbf{x}}}^{uk}$ denotes the unknown pixels of the patch whereas $\psi_{p_{\mathbf{x}}}^k$ denotes its known pixels.
- LR and HR denote Low Resolution and High Resolution respectively. $\psi_{p_{\mathbf{x}}}^{HR}$ is a patch of high resolution.
- \mathcal{D}^{HR} is a set of HR patches (a dictionary) which is used by the super-resolution algorithm.

B. Algorithm overview

Image completion of large missing regions is a challenging task. As presented in the previous section, there are a number of solutions to tackle the inpainting problem. In this paper, we propose a new inpainting framework relying on both the combination of low-resolution inpainting pictures method and a single-image super-resolution algorithm. In the following sections, we briefly present the main ideas of this paper and the reasons why the proposed method is new and innovative.

The proposed method is composed of two main and sequential operations. The first one is a non-parametric patch sampling method used to fill in missing regions. The inpainting algorithm is preferably applied on a coarse version of the input picture. Indeed a low-resolution picture is mainly represented by its dominant and important structures of the scene. We believe that performing the inpainting of such a low-resolution image is much easier than performing it on the full resolution. A low-resolution image is less contaminated by noise and is composed by the main scene structures. In other words, in this kind of picture, local orientation singularities which could affect the filling order computation are strongly reduced. Second, as the picture to inpaint is smaller than the original one, the computational time is significantly reduced compared to the one necessary to inpaint the full resolution image. To give more robustness, we inpaint the low-resolution picture with different settings (patch's size, filling order, etc). By combining these results, a final low-resolution inpainted picture is obtained. Results will show that the robustness

and the visual relevance of inpainting is improved. The second operation is run on the output of the first step. Its goal is to enhance the resolution and the subjective quality of the inpainted areas. Given a low-resolution input image, which is the result of the first inpainting step, we recover its high-resolution using a single-image super-resolution approach. Figure 1 illustrates the main concept underlying the proposed method namely:

- 1) a low-resolution image is first built from the original picture;
- 2) an inpainting algorithm is applied to fill in the holes of the low-resolution picture. Different settings are used and inpainted pictures are combined;
- 3) the quality of the inpainted regions is improved by using a single-image super-resolution method.

III. MULTIPLE EXAMPLAR-BASED INPAINTING

This section aims at presenting the proposed inpainting method and the combination of the different inpainted images.

A. Exemplar-based inpainting

The proposed exemplar-based method follows the two classical steps as described in [4]: the filling order computation and the texture synthesis. These are described in the next sections.

1) *Patch priority*: The filling order computation defines a measure of priority for each patch in order to distinguish the structures from the textures. Classically, a high priority indicates the presence of structure. The priority of a patch centered on p_x is just given by a data and confidence term. The latter is exactly the one defined in [4]. Regarding the data term, a tensor-based [9] and a sparsity-based [16] data terms have been used. These terms are briefly described below.

The *tensor-based priority term* is based on a structure tensor also called Di Zenzo matrix [17]; this is given by:

$$\mathbf{J} = \sum_{i=1}^m \nabla I_i \nabla I_i^T \quad (2)$$

\mathbf{J} is the sum of the scalar structure tensors $\nabla I_i \nabla I_i^T$ of each image channel I_i (R,G,B). The tensor can be smoothed without cancellation effects : $\mathbf{J}_\sigma = \mathbf{J} * G_\sigma$ where $G_\sigma = \frac{1}{2\pi\sigma^2} \exp(-\frac{x^2+y^2}{2\sigma^2})$, with standard deviation σ . One of the main advantages of a structure tensor is that a structure coherence indicator can be deduced from its eigenvalues. Based on the discrepancy of the eigenvalues, the degree of anisotropy of a local region can be evaluated. The local vector geometry is computed from the structure tensor \mathbf{J}_σ . Its eigenvectors $\mathbf{v}_1, \mathbf{v}_2$ ($\mathbf{v}_i \in R^n$) define an oriented orthogonal basis and its eigenvalues $\lambda_{1,2}$ define the amount of structure variation. The vector \mathbf{v}_1 is the orientation with the highest fluctuations (orthogonal to the image contours), whereas \mathbf{v}_2 gives the preferred local orientation. This eigenvector (having the smallest eigenvalue) indicates

the isophote orientation. A data term D is then defined as [18]:

$$D(p_x) = \alpha + (1 - \alpha) \exp\left(-\frac{\eta}{(\lambda_1 - \lambda_2)^2}\right) \quad (3)$$

where η is a positive value and $\alpha \in [0, 1]$ ($\eta = 8$ and $\alpha = 0.01$). On flat regions ($\lambda_1 \approx \lambda_2$), any direction is favoured for the propagation (isotropic filling order). When $\lambda_1 \gg \lambda_2$ indicating the presence of a structure, the data term is important.

The *sparsity-based priority* has been proposed recently by Xu et al. [16]. In a search window, a template matching is performed between the current patch ψ_{p_x} and neighbouring patches ψ_{p_j} that belong to the known part of the image. By using a non-local means approach [15], a similarity weight w_{p_x, p_j} (i.e. proportional to the similarity between the two patches centered on p_x and p_j) is computed for each pair of patches. The sparsity term is defined as:

$$D(p_x) = \|\mathbf{w}_{p_x}\|_2 \times \sqrt{\frac{|N_s(p_x)|}{|N(p_x)|}} \quad (4)$$

where N_s and N represent the number of valid patches (having all its pixels known) and the total number of candidates in the search window. When $\|\mathbf{w}_{p_x}\|_2$ is high, it means larger sparseness whereas a small value indicates that the current input patch is highly predictable by many candidates.

2) *Texture synthesis*: The filling process starts with the patch having the highest priority. To fill in the unknown part of the current patch $\psi_{p_x}^{uk}$, the most similar patch located in a local neighbourhood \mathcal{W} centered on the current patch is sought. A similarity metric is used for this purpose. The chosen patch $\psi_{p_x}^*$ maximizes the similarity between the known pixel values of the current patch to be filled in $\psi_{p_x}^k$ and co-located pixel values of patches belonging to \mathcal{W} :

$$\begin{aligned} \psi_{p_x}^* &= \arg \min_{\psi_{p_j} \in \mathcal{W}} d(\psi_{p_x}^k, \psi_{p_j}^k) \\ &s.t \quad Coh(\psi_{p_x}^{uk}) < \lambda_{coh} \end{aligned} \quad (5)$$

where $d(\cdot)$ is the weighted Bhattacharya used in [10]. $Coh(\cdot)$ is the coherence measure initially proposed by Wexler et al. [15]:

$$Coh(\psi_{p_x}^{uk}) = \min_{p_j \in S} \left(d_{SSD}(\psi_{p_x}^{uk}, \psi_{p_j}^{uk}) \right) \quad (6)$$

where d_{SSD} is the sum of square differences. The coherence measure Coh simply indicates the degree of similarity between the synthesized patch $\psi_{p_x}^{uk}$ and original patches. Therefore, the constraint in equation (5) prevents pasting in the unknown regions a texture that would be too different from original textures. If none of the candidates fulfil the constraint (5), the filling process is stopped and the priority of the current patch is decreased. The process restarts by seeking the patch having the highest priority. It is interesting to note that a recent study [19] uses a similar term to predict the quality of the inpainting.

Compared to our previous work [10], there is another substantial difference: we only use the best match to fill in the hole whereas a linear combination of the K most similar patches is generally performed to compute the patch $\psi_{p_x}^*$ in [10], [15], [16], [20]. In these cases, the estimated patch is then given by:

$$\psi_{p_x}^* = \sum_i^K w_{p_x, p_i} \times \psi_{p_i}^k \quad (7)$$

where K is the number of candidates which is often adapted locally so that the similarity of chosen neighbours lies within a range $(1 + \alpha) \times d_{min}$, where d_{min} is the distance between the current patch and its closest neighbours. Different methods can be used to compute the weights. It could be based on either a non-negative matrix factorization (NMF) [21] or a non local means filter [22], [15], to name a few. Combining several candidates increases the algorithm robustness. However, it tends to introduce blur on fine textures as illustrated by figure 2. In our method, only the best candidate is chosen. Its unknown parts are pasted into the missing areas. A Poisson fusion [23] is applied to hide the seams between known and unknown parts. An α -blending is also applied to combine known parts of the candidate and the source patch ($\psi_{p_x}^k \leftarrow \alpha \psi_{p_x}^k + (1 - \alpha) \psi_{p_x}^{k,*}$, with $\alpha = 0.75$). It gives more robustness by locally regularizing the results.

Although the proposed method is able to fill in holes in a visually pleasant fashion (as illustrated by figure 2 and 9), it still suffers from problems of one-pass greedy algorithms. Indeed for most of existing approaches, the setting such as the patch size and the filling order, to name the most important factors, may dramatically impact the quality of results. To overcome this issue, we combine inpainted pictures obtained when different settings are used. In this study, we consider $M = 13$, meaning that the low-resolution picture is inpainted 13 times. Parameters are given in Table I: the patch size is chosen between 5×5 , 7×7 , 9×9 and 11×11 . The filling order is computed by either the sparsity-based or the tensor-based method. The input picture can also be rotated by 180 degrees. This allows changing the filling order.

B. Combining multiple inpainted images

The combination aims at producing a final inpainted picture from the M inpainted pictures. Before delving into this subject in details, Figure 3 illustrates some inpainted results obtained for a given setting. We notice again that the setting plays an important role. To obtain the final inpainted picture, three kinds of combination have been considered. The first two methods are very simple since every pixel value in the final picture is achieved by either the average or the median operator as given below:

$$\widehat{\mathbf{I}}^{(*)}(p_x) = \frac{1}{M} \sum_{i=1}^M \widehat{\mathbf{I}}^{(i)}(p_x) \quad (8)$$

$$\widehat{\mathbf{I}}^{(*)}(p_x) = MED_{i=1}^M \widehat{\mathbf{I}}^{(i)}(p_x) \quad (9)$$

TABLE I
THIRTEEN CONFIGURATIONS USED TO FILL IN THE UNKNOWN PARTS OF THE PICTURES.

Setting	Parameters
1 (default)	Patch's size 5×5 Decimation factor $n = 3$ Search window 80×80 Sparsity-based filling order
2	default + rotation by 180 degrees
3	default + patch's size 7×7
4	default + rotation by 180 degrees + patch's size 7×7
5	default + patch's size 11×11
6	default + rotation by 180 degrees + patch's size 11×11
7	default + patch's size 9×9
8	default + rotation by 180 degrees + patch's size 9×9
9	default + patch's size 9×9 + Tensor-based filling order
10	default + patch's size 7×7 + Tensor-based filling order
11	default + patch's size 5×5 + Tensor-based filling order
12	default + patch's size 11×11 + Tensor-based filling order
13	default + rotation by 180 degrees + patch's size 9×9 + Tensor-based filling order

The advantage of these operators is their simplicity. However they suffer from at least two main drawbacks. The average operator as well as the median one do not consider the neighbours of the current pixel to take a decision. Results might be more spatially coherent by considering the local neighbourhood. In addition, the average operator inevitably introduces blur as illustrated by Figure 5.

To cope with these problems, namely blur and spatial consistency of the final result, the combination is achieved by minimizing an objective function. Rather than using a global minimization that would solve exactly the problem, we use a Loopy Belief Propagation which in practice provides a good approximation of the problem to be solved. This approach is described in the next section.

1) *Loopy Belief Propagation*: As in [24], the problem is to assign a label to each pixel p_x of the unknown regions T of the picture $\widehat{\mathbf{I}}^{(*)}$. The major drawback of the belief propagation is that it is slow especially when the number of labels is high. Komodakis and Tziritas [24] have designed a priority Belief Propagation in order to deal with this complexity bottleneck. Indeed, the number of labels in [24] is equal to the number of patches in the source region. Here the approach is simpler since the number of labels is rather small; a label is simply the index of the inpainted picture from which the patch is extracted. A finite set of labels \mathcal{L} is then composed of M values ($M = 13$ here), going from 1 to M . This problem can be formalized with a Markov Random Field (MRF) $G = (\nu, \epsilon)$ defined over the target region T . The MRF nodes ν are the lattice composed of pixels inside T . Edges ϵ are the four-connected image grid graph centered around each node. We denote \mathcal{N}_4 this neighbourhood system. The labelling assigns a label l ($l \in \mathcal{L}$) to each node/pixel p_x ($p_x \in T$)




Fig. 2. (a) original picture with a black hole to be filled in; (b) result obtained with a simple template matching (TM), with the linear combination of the K most similar neighbours in which weights are computed by using a non local means method (c) and by a non negative matrix factorization (d). For this example, the maximum number of neighbours is 10. (e) to (g) zoom in the inpainted part for TM, NLM and NMF, respectively.




Fig. 3. The low-resolution picture is inpainted with different settings S_x (see table I for the setting details). (a) original picture with the hole to be filled in. Pictures S_1 up to S_8 (from (b) to (i)) are the inpainting results (note the setting sensibility of the inpainting algorithm).

so that the total energy E of the MRF is minimized (we denote by l_p the label of pixel p_x) [25], [26]:

$$E(l) = \sum_{p \in \nu} V_d(l_p) + \sum_{(n,m) \in \mathcal{N}_4} V_s(l_n, l_m) \quad (10)$$

where,

- $V_d(l_p)$ is called the label cost (or the data cost) [24]. This represents the cost of assigning a label l_p to a pixel p_x . This is given by:

$$V_d(l_p) = \sum_{n \in \mathcal{L}} \sum_{\mathbf{u} \in \nu} \left\{ \widehat{\mathbf{I}}^{(l)}(\mathbf{x} + \mathbf{u}) - \widehat{\mathbf{I}}^{(n)}(\mathbf{x} + \mathbf{u}) \right\}^2 \quad (11)$$

where ν is a square neighbourhood (3×3) centered on the current pixel. The cost increases when the dissimilarity between the current patch and collocated patches in other inpainted pictures is high.

- The pairwise potential or the discontinuity cost $V_s(l_n, l_m)$ is a quadratic cost function given by:

$$V_s(l_n, l_m) = \lambda \times (l_n - l_m)^2 \quad (12)$$

where λ is a weighting factor and is set to 100 (Figure 4 illustrates the influence of this factor). The discontinuity cost is here based on a difference between labels rather than the difference between pixel values.

The minimization of the energy E over the target region T can be achieved using loopy belief propagation (LBP) [27] and corresponds to the maximum a posteriori

(MAP) estimation problem for an appropriately defined MRF [28]. Figure 4 illustrates the minimization of the total energy as well as the influence of the smoothness term. When $\lambda = 0$, there is no smoothness term. Some artefacts indicated by arrows are visible. When $\lambda = 500$, artefacts are visible. A good trade-off is obtained by setting the value λ to 100. On the same figure 4, the labelling convergence is given for four iterations (pictures on top of this figure represent label values, not pixel values). This illustrates the fact that the label choice is not greedy.

2) *Comparison of the combination methods:* Figure 5 illustrates the performance of the combination methods. As expected, when the different inpainted pictures are averaged, the reconstructed areas are blurred. The blur is less striking when the median operator is used to combine pictures. The LBP method provides the best result. The texture is well retrieved and thanks to the global energy minimization results are spatially consistent. In the following, we use the LBP method to combine low-resolution inpainted pictures.

3) *Possible extensions:* Two extensions of the proposed method could be also envisioned. First, rather than combining the low-resolution inpainted pictures, the best inpainting picture could be selected either by a user (semi-supervised inpainting) or by automatically evaluating the inpainting quality using a coherence measure as proposed in [19].

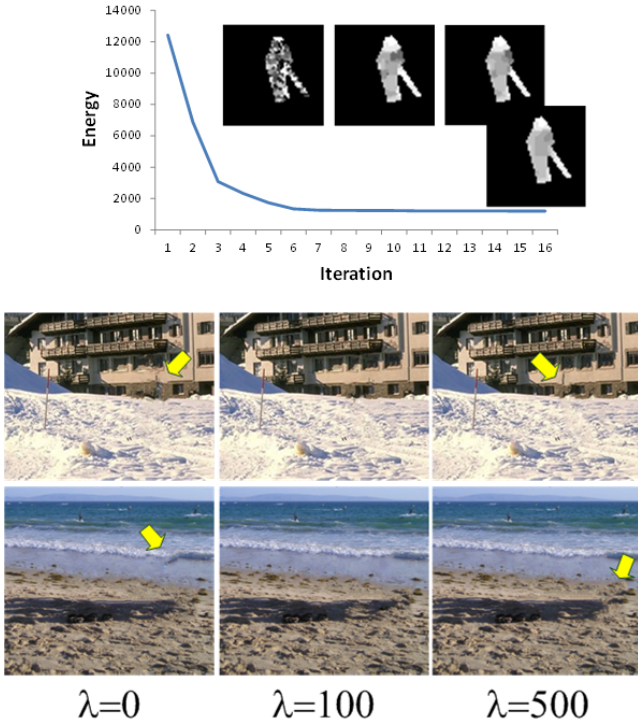


Fig. 4. Illustration of the energy minimization. On the top, the convergence process of the LBP is illustrated. The labelling obtained for four different iterations are also shown. The bottom images show the influence of the smoothness term for three values of λ (see equation (12)).




Fig. 5. Comparison of combination methods. (a) Input picture; (b) results obtained by averaging all inpainted pictures, (c) by taking the median pixel values and (d) by using a Loopy Belief Propagation.

IV. SUPER-RESOLUTION ALGORITHM

Once the combination of the low-resolution inpainted pictures is completed, a hierarchical single-image super-resolution approach is used to reconstruct the high resolution details of the image. We stress the point that the single-image SR method is applied only when the input picture has been downsampled for the inpainting purpose. Otherwise the SR method is not required. As in [12], [10], the problem is to find a patch of higher-resolution from a database of examples. The main steps, illustrated in figure 6 are described below:

- 1) Dictionary building: it consists of the correspondences between low and high resolution image patches. The unique constraint is that the high-

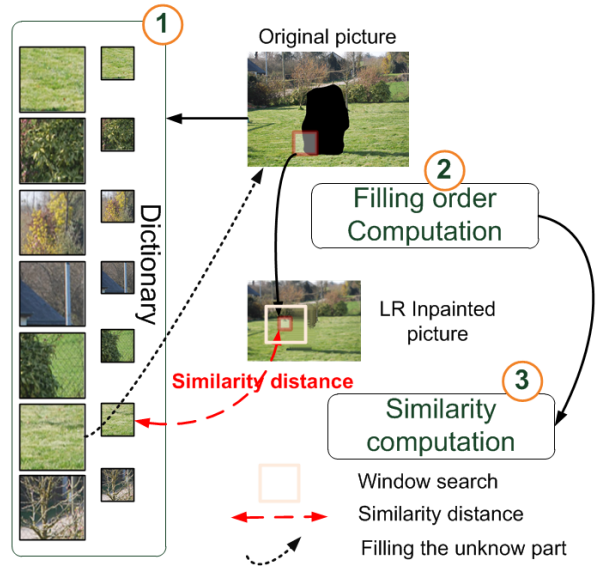


Fig. 6. Flowchart of the super-resolution algorithm. The missing parts of the red block are filled in by the best candidate stemming either from the dictionary or from the local neighbourhood. The top image represents the original image with the missing areas whereas the bottom one is the result of the low-resolution inpainting.

resolution patches have to be valid, i.e. entirely composed of known pixels. In the proposed approach, high-resolution and valid patches are evenly extracted from the known parts of the image. The size of the dictionary is a user-parameter which might influence the overall speed/quality trade-off. An array is used to store the spatial coordinates of HR patches (\mathcal{D}^{HR}). Those of LR patches are simply deduced by using the decimation factor equal to 2;

- 2) Filling order of the HR picture: The computation of the filling order is similar to the one described in Section III. It is computed on the HR picture with the sparsity-based method. The filling process starts with the patch $\psi_{p \times}^{HR}$ having the highest priority and which is composed of known and unknown parts. Compared to a raster-scan filling order, it allows us to start with the structures and then to preserve them;
- 3) For the LR patch corresponding to the HR patch having the highest priority, its best neighbour in the inpainted images of lower resolution is sought. This search is performed in the dictionary and within a local neighbourhood. Only the best candidate is kept. From this LR candidate, a HR patch is simply deduced. Its pixel values are then copied into the unknown parts of the current HR patch $\psi_{p \times}^{HR}$.

After the filling of the current patch, the priority value is propagated and the aforementioned steps are iterated while there exist unknown areas. A Poisson and alpha-blending are again used to hide seams between known and unknown parts and to improve robustness.

Compared to [10], the SR method is applied in a hierarchical manner. For instance, if the input picture of resolution (X, Y) has been down-sampled by four in both



Fig. 7. Results of the proposed method. (a) input images with unknown regions; (b) inpainted low-resolution images; (c) final inpainted images (for the sake of visibility, we do not respect the down sampling factor of 4 between pictures (b) and (c)).

directions, the SR algorithm is applied twice: a first time to recover the resolution $(\frac{X}{2}, \frac{Y}{2})$ and a second time to recover the native resolution.

V. EXPERIMENTAL RESULTS

In this section, the proposed approach is tested on a variety of natural images and compared to five state-of-the-art inpainting methods.

A. Intrinsic performance of the proposed method

Figures 7 and 8 illustrate the results of the proposed approach. A down sampling factor of 4 in both directions is used. Pictures have a resolution varying in the range 420×380 to 720×512 . Figure 7 (b) gives the inpainted low-resolution picture whereas the column (c) illustrates the final result. The proposed approach faithfully recovers the texture such as the grass, the sand and the snow. Structures are also well recovered. Figure 8 presents more results which are visually plausible and pleasing in most of the cases. The less-favourable results are obtained when the hole to be filled in is rather small. In this case it might be better to reduce the down sampling factor or even to perform the inpainting at the full resolution.

B. Parameters analysis

When the proposed method, called the baseline method in the following, has been defined, we have fixed several

parameters such as the number of K-NN for the inpainting ($K = 1$), the similarity metric (weighted Bhattacharya) and the color space (RGB). We evaluate subjectively the relevance of these choices on two pictures. Results are illustrated in Figure 9.

1) *K-NN patches for inpainting*: In Section III-A2, we set the number of K-NN for the exemplar-based inpainting to 1 preventing blur apparition as illustrated by figure 2. However, as the inpainting is applied on a low-resolution picture, it might make sense to use more than one candidate. We compare results obtained by the baseline approach to those obtained for three different K values: 1 (baseline algorithm), 4 and 8. Results are illustrated in figure 9. Results are similar for $K = 4$. For $K = 8$, the inpainted quality is not as visually pleasing as the one obtained by the two previous settings.

2) *RGB space versus Y component*: The color space, used to perform the candidate search for instance, might play an important role. By using the luminance component only, results are visually less good than the baseline one, especially for the first picture.

3) *SSD versus weighted Bhattacharya*: Qualitatively speaking, inpainted results obtained with the weighted Bhattacharya metric are slightly better than those obtained with the SSD metric.

4) *Downsampling factor equal to 2*: By performing the inpainting at a higher resolution, results are not as good as



Fig. 8. Results of proposed approach for a variety of pictures of natural content. The unknown parts in black visible on pictures of odd rows are more or less large and complex to fill in. Pictures of even rows are our results.




Fig. 9. Parameters analysis. (a) baseline method; (b) baseline + $K=4$; (c) baseline + $K=8$; (d) Y space; (e) SSD; (f) with a decimation factor of 2 in both directions (instead of 4).

those obtained by the baseline approach, especially on the first picture (first row). This could be explained by the size of the search window which should be increased for this condition.

C. Comparison with state-of-the-art methods

State-of-the-art inpainting methods are used for comparison purposes. The first one is the well known method proposed by Criminisi et al. [4]. We have used a third part implementation available on <http://www.cc.gatech.edu/~sooraj/inpainting/> (we use a patch's size of 9×9 and a local windows search of 80×80). The second one is the PatchMatch method [7] which is a fast algorithm for computing dense approximate nearest neighbour correspondences between patches of two image regions. This algorithm is available in Adobe Photoshop CS5. The shift-map method [29] is the third algorithm considered in

the test. This approach is treating inpainting as a global optimization in a multi-scale scheme. The goal is to shift pixel values from the known part of the picture to the missing part. A graph cut optimization allows finding the best mapping. Inspired by [29], He and Sun [30] recently proposed a new method. The idea is to find the most important 2D offsets (displacement) in the picture. These offsets are then used to shift the input picture. The shifted images are finally combined. Method of Komodakis and Tziritis [24] relying on priority belief propagation is also tested.

Figure 10 illustrates results of the proposed approach, Criminisi et al. [4] and PatchMatch method [7]. Our method outperforms Criminisi et al. [4]. Compared to PatchMatch method [7], results are similar for the two first pictures and better for the two last pictures. Comparison between the proposed approach and He's method [30], shift-map [29]

and priority belief propagation [24] is given on figure 11. The tested pictures are extracted from [30]. Concerning the comparison between our method and He’s method, results are quite similar except for the first and last pictures for which He’s approach gives more visually pleasing results, although there is no obvious artefacts in our results. Regarding shift-map and BP methods (the three last rows of figure 11), artefacts are visible such as on the stone wall (fourth row) and on the waterfall (BP method on the fifth row). On these pictures, our method is more robust than shift-map and Belief Propagation methods.

D. Extension

1) *Application to texture synthesis*: Figure 12 presents additional results for texture synthesis. A small chunk of texture was placed into the upper left corner of an empty image. The goal is to fill in the missing parts composed here of 256×256 pixels in a visually plausible way. This application is interesting in itself, but it also allows showing a limitation of the proposed approach. When the texture is stochastic as illustrated by figure 12, there exist many ways to fill in the holes in a visually plausible fashion. Performing M different inpaintings and combining them would lead to a poor quality. To overcome this limitation, the input texture is inpainted just once at a low resolution. The SR method is then applied. For texture synthesis, the patch size is arbitrarily set to 9×9 and we down-sample the input chunk by 4 in both directions. In addition, two modifications are also brought to the algorithm. First we limit the search for the closest candidate to the input picture (here it is the original texture chunk). Second the filling order is the raster scan order.

The two first texture chunks illustrated by figure 12 are complex to synthesize. Some artefacts are visible on the first two results whereas for the two last results the synthesised texture has a very good quality.

2) *Using similar images*: Another extension would be to use similar images in a similar fashion to Hays and Efros [31] and Whyte et al. [32]. We have adapted the proposed method to take as input the image to inpaint, similar images to the image source and an image which is used as a guide for the LBP. For that the energy of the LBP has been modified as follows:

$$E(l) = \sum_{\mathbf{p} \in \nu} V_d(l_p) + \sum_{(\mathbf{p}, \mathbf{q}) \in \nu} V_s(l_p, l_q) + \sum_{\mathbf{p} \in \nu} V_g(l_p) \quad (13)$$

where, the first term is the same as previously defined (see equation (11)). The second term is here different since it involves pixel values rather than label values as in [30]. This term is given by:

$$V_s(m, n) = \|\hat{\mathbf{I}}^{(m)}(\mathbf{p}) - \hat{\mathbf{I}}^{(m)}(\mathbf{q})\|^2 + \|\hat{\mathbf{I}}^{(n)}(\mathbf{p}) - \hat{\mathbf{I}}^{(n)}(\mathbf{q})\|^2 \quad (14)$$

with (\mathbf{p}, \mathbf{q}) are neighbouring pixels and m, n represent their labels. The third term V_g is used to guide the labelling. It is given by:

$$V_g(n) = \|\hat{\mathbf{I}}^{(n)}(\mathbf{p}) - \hat{\mathbf{I}}_g(\mathbf{p})\|^2 \quad (15)$$




Fig. 12. Examples of texture synthesis. (a) chunk of texture (256×256) with the unknown parts in black (256×256); (b) Results obtained by the proposed approach.

where $\hat{\mathbf{I}}_g$ is a prediction image of the unknown pixels. The prediction can be obtained by a classical inpainting such as our baseline method presented in section III-A or by an diffusion-based method as illustrated in figure 13. This term is useful to select the most appropriate pixel values in the set of similar pictures: for a given pixel \mathbf{p} , if the candidate pixel value at that location is close to the predicted one $\hat{\mathbf{I}}_g(\mathbf{p})$, the associate cost is low. Figure 13 illustrates this method. From an image to inpaint (top-middle), 21 similar images were found by using *Google Image search*. To guide the LBP, we use the inpainted image from either our baseline inpainting method or the Navier-Stoke diffusion method [33]. These two guides coupled with the similar images already provide quite good results. Result of the method involving multiple inpainting (bottom-right of figure 13) is also shown and further improves the quality.

Note that rather than using similar images, it would be possible to use state-of-the-art inpainted pictures and to combine them. This kind of solution might improve performance because of the better complementary of state-of-the-art inpainting methods. It would be possible to mix results from exemplar-based, diffusion-based and shift-map inpainting methods. Finally, a semi-supervised approach for which the user is able to choose the best low-resolution inpainted picture could be easily designed. In this case, it is no longer necessary to combine them.

E. Running time

Table II gives the running time of the proposed approach. Simulation has been performed on a laptop with an Intel




Fig. 10. Comparison with state-of-the-art methods: (a) original pictures with unknown regions; (b) Criminisi’s algorithm; (c) PatchMatch method; (d) proposed method.

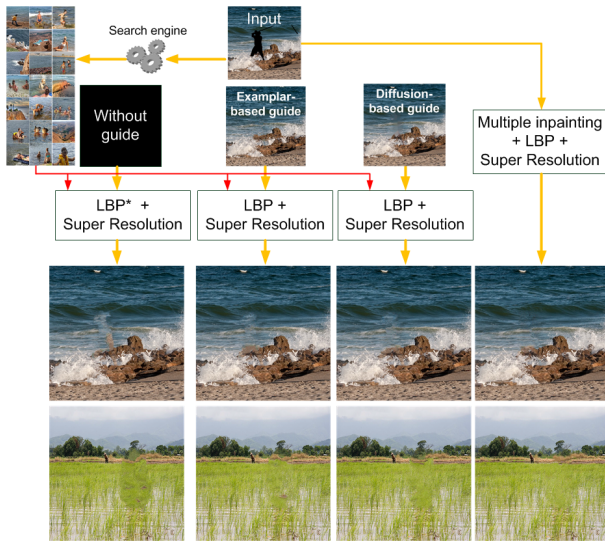


Fig. 13. Extension of the proposed approach. From a set of similar images, an inpainting is performed on a low-resolution version of the input picture. The LBP uses a set of similar images (21 images here (top-left)) and a guide (if any). The SR is applied to recover the details of the missing pixels. Different guides (exemplar-based or diffusion-based inpainting) can be used to drive the LBP. (*) (in the first box on the left) means that the term V_g of equation (13) is null. Bottom-up line gives other results.

Core i7 2.40GHz and 4Go RAM. As the proposed approach is not multithreaded, it just uses one core. In addition, no optimization was made. Table II indicates that the super-resolution process is the most time-consuming step (in

average it represents 80% of the total time). It is due to the template matching which is not parallelized. To improve the performance, the template matching could be replaced by an approximate nearest neighbour search [7].

TABLE II
RUNNING TIME IN FUNCTION OF THE NUMBER OF MISSING PIXELS.
THE RUNNING TIME IS GIVEN FOR THE LOW-RESOLUTION
INPAINTINGS AND THE SUPER-RESOLUTION.

Picture	Resolution	Missing areas	Inpainting and LBP	SR	Total
Elephant	480 × 320	17%	32s	1m58s	2m30s
SnowBall	480 × 320	13%	22s	1m36s	1m58s
Tiger	480 × 320	28%	59s	2m36s	3m35s
Soldier	320 × 480	30%	51s	2m39s	3m30s
Cow	600 × 400	13%	40s	2m11s	2m51s

VI. CONCLUSION

A novel inpainting approach has been presented in this paper. The input picture is first down sampled and several inpaintings are performed. The low-resolution inpainted pictures are combined by globally minimizing an energy term. Once the combination is completed, a hierarchical single image super resolution method is applied to recover details at the native resolution. Experimental results on a wide variety of images have demonstrated the effectiveness of the proposed method.

One interesting avenue of future work would be to extend this approach to the temporal dimension. Also, we plan to test other SR methods to bring more robustness to the



Fig. 11. Comparison with state-of-the-art results. On the three first rows: (a) original image; (b) He’s results (extracted from [30]), (c) proposed approach. On the three last rows: (a) original image; (b) He’s results; (c) Shift-map results [29]; (d) BP results [24] and (d) proposed approach.

method. But the main important improvement is likely the use of geometric constraint and higher-level information such as scene semantics in order to improve the visual relevance.

Windows executable is available on authors webpage: http://people.irisa.fr/Olivier.Le_Meur/.

REFERENCES

- [1] Bertalmio, M., Sapiro, G., Caselles, V., Ballester, C.: Image inpainting. In: SIGGRPAH 2000. (2000)
- [2] Tschumperlé, D., Deriche, R.: Vector-valued image regularization with pdes: a common framework for different applications. *IEEE Trans. on PAMI* **27** (2005) 506–517
- [3] Chan, T., Shen, J.: Variational restoration of non-flat image features: models and algorithms. *SIAM J. Appl. Math.* **61** (2001) 1338–1361
- [4] Criminisi, A., Pérez, P., Toyama, K.: Region filling and object removal by exemplar-based image inpainting. *IEEE Trans. On Image Processing* **13** (2004) 1200–1212
- [5] Drori, I., Cohen-Or, D., Yeshurun, H.: Fragment-based image completion. *ACM Trans. Graph.* **22** (2003) 303–312
- [6] Harrison, P.: A non-hierarchical procedure for re-synthesis of complex texture. In: Proc. Int. Conf. Central Europe Comp. Graphics, Visua. and Comp. Vision. (2001)
- [7] Barnes, C., Shechtman, E., Finkelstein, A., Goldman, D.B.: Patch-Match: A randomized correspondence algorithm for structural image editing. *ACM Transactions on Graphics (Proc. SIGGRAPH)* **28** (2009)
- [8] Efros, A.A., Leung, T.K.: Texture synthesis by non-parametric sampling. In: *IEEE Computer Vision and Pattern Recognition (CVPR)*. (1999) 1033–1038
- [9] Le Meur, O., Gautier, J., Guillemot, C.: Exemplar-based inpainting based on local geometry. In: *ICIP*. (2011)
- [10] Le Meur, O., Guillemot, C.: Super-resolution-based inpainting. In: *ECCV*. (2012)
- [11] Dai, S., Han, M., Xu, W., Wu, Y., Gong, Y., Katsaggelos, A.: Softcuts: a soft edge smoothness prior for color image super-resolution. *IEEE Trans. On Image Processing* **18** (2009) 969–981
- [12] Freeman, W.T., Jones, T.R., Pasztor, E.C.: Example-based super-resolution. *IEEE Computer Graphics and Applications* **22** (2002) 56–65
- [13] Glasner, D., Bagon, S., Irani, M.: Super-resolution from a single image. In: *IEEE International Conference on Computer Vision (ICCV)*. Volume 10. (2009) 349356
- [14] Chang, H., Yeung, D.Y., Xiong, Y.: Super-resolution through neighbor embedding. In: *IEEE Computer Vision and Pattern Recognition (CVPR)*. Volume I. (2004) 275–282
- [15] Wexler, Y., Shechtman, E., Irani, M.: Space-time video completion. In: *IEEE Computer Vision and Pattern Recognition (CVPR)*. (2004)
- [16] Xu, Z., Sun, J.: Image inpainting by patch propagation using patch sparsity. *IEEE Trans. on Image Processing* **19** (2010) 1153–1165
- [17] Di Zeno, S.: A note on the gradient of a multi-image. *Computer Vision, Graphics, and Image Processing* **33** (1986) 116–125
- [18] Weickert, J.: Coherence-enhancing diffusion filtering. *International Journal of Computer Vision* **32** (1999) 111–127
- [19] Kopf, J., Kienzle, W., Drucker, S., Kang, S.B.: Quality prediction for image completion. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia 2012)* **31** (2012)
- [20] Bugeau, A., Bertalmio, M., Caselles, V., Sapiro, G.: A comprehensive framework for image inpainting. *IEEE Trans. on Image Processing* **19** (2010) 2634–2644
- [21] Lee, D.D., Seung, H.S.: Algorithms for non-negative matrix factorization. In: *Advances in Neural Information Process. Syst. (NIPS)*. (2000)
- [22] Buades, A., Coll, B., Morel, J.: A non local algorithm for image denoising. In: *IEEE Computer Vision and Pattern Recognition (CVPR)*. Volume 2. (2005) 60–65
- [23] Pérez, P., Gangnet, M., Blake, A.: Poisson image editing. In: *SIGGRAPH*. (2003)
- [24] Komodakis, N., Tziritis, G.: Image completion using efficient belief propagation via priority scheduling and dynamic pruning. *IEEE Trans. On Image Processing* **16** (2007) 2649 – 2661
- [25] Boykov, Y., Veksler, O., R.Zabih: Efficient approximate energy minimization via graph cuts. *IEEE Trans. On PAMI* **20** (2001) 1222–1239
- [26] Boykov, Y., Kolmogorov, V.: An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Trans. On PAMI* **26** (2004) 1124–1137
- [27] Yedidia, J., Freeman, W., Weiss, Y.: Constructing free energy approximations and generalized belief propagation algorithms. *J.S. Yedidia, W.T. Freeman, and Y. Weiss* **51** (2005) 2282–2312
- [28] Blake, A., Zisserman, A.: *Visual reconstruction*. MIT Press (1987)
- [29] Pritch, Y., Kav-Venaki, E., Peleg, S.: Shift-map image editing. In: *ICCV'09, Kyoto (2009)* 151–158
- [30] He, K., Sun, J.: Statistics of patch offsets for image completion. In: *ECCV*. (2012)
- [31] Hays, J., Efros, A.A.: Scene completion using millions of photographs. *ACM Transactions on Graphics (SIGGRAPH 2007)* **26** (2007)
- [32] Whyte, O., Sivic, J., Zisserman, A.: Get out of my picture! internet-based inpainting. In: *Proceedings of the 20th British Machine Vision Conference, London*. (2009)
- [33] Bertalmio, M., Bertozzi, A.L., Sapiro, G.: Navier-stokes, fluid dynamics, and image and video inpainting. In: *Proc. IEEE Computer Vision and Pattern Recognition (CVPR)*. (2001) 355–362



Olivier Le Meur obtained his PhD degree from the University of Nantes in 2005. From 1999 to 2009, he has worked in the media and broadcasting industry. In 2003 he joined the research center of Thomson-Technicolor at Rennes where he supervised a research project concerning the modelling of the human visual attention. Since 2009 he has been an associate professor for image processing at the University of Rennes 1. In the IRISA/SIROCCO team his research interests are dealing with the understanding of the human visual attention. It includes computational modelling of the visual attention and saliency-based applications (video compression, objective assessment of video quality, retargeting). Home page: http://people.irisa.fr/Olivier.Le_Meur/.



Mounira Ebdelli received the engineering degree from the national school of computer science, Tunisia in 2005 and the M.Sc. degree in perception and digital communication from the national school of engineering, Tunisia and Paris-Descartes, France, in 2010. She is currently working toward the Ph.D. degree with the SIROCCO Research Group, INRIA, Rennes, France. From 2005 to October 2008, she has worked as a developer in STMicroelectronics, Tunisia, and she was a Teaching Assistant with the High Institute of Technological Studies in Zaghouan, Tunisia from 2008 to November 2010. Her research interests include signal processing with an emphasis on image and video processing.



Christine Guillemot is Director of Research at INRIA, head of a research team dealing with image and video modelling, processing, coding and communication. She holds a PhD degree from ENST (Ecole Nationale Supérieure des Télécommunications) Paris, and an 'Habilitation for Research Direction' from the University of Rennes. From 1985 to Oct. 1997, she has been with FRANCE TELECOM, where she has been involved in various projects in the area of image and video coding for TV, HDTV and multimedia.

From Jan. 1990 to mid.1991, she has worked at Bellcore, NJ, USA, as a visiting scientist. She has (co)-authored 15 patents, 8 book chapters, 50 journal papers and 140 conference papers. She has served as associated editor (AE) for the IEEE Trans. on Image processing (2000-2003), for IEEE Trans. on Circuits and Systems for Video Technology (2004-2006) and for IEEE Trans. on Signal Processing (2007-2009). She is currently AE for the Eurasip journal on image communication. She has been a member of the IEEE IMDSP and MMSP technical committees. Home page <http://people.rennes.inria.fr/Christine.Guillemot/>.