



SuperPatchMatch: an Algorithm for Robust Correspondences using Superpixel Patches

Rémi Giraud, Vinh-Thong Ta, Aurélie Bugeau, Pierrick Coupé, Nicolas Papadakis

► To cite this version:

Rémi Giraud, Vinh-Thong Ta, Aurélie Bugeau, Pierrick Coupé, Nicolas Papadakis. SuperPatchMatch: an Algorithm for Robust Correspondences using Superpixel Patches. IEEE Transactions on Image Processing, 2017, 10.1109/TIP.2017.2708504 . hal-01432116

HAL Id: hal-01432116

<https://hal.science/hal-01432116>

Submitted on 9 Jun 2017

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.

SuperPatchMatch: an Algorithm for Robust Correspondences using Superpixel Patches

Rémi Giraud, Vinh-Thong Ta, Aurélie Bugeau, Pierrick Coupé, and Nicolas Papadakis

Abstract—Superpixels have become very popular in many computer vision applications. Nevertheless, they remain underexploited since the superpixel decomposition may produce irregular and non stable segmentation results due to the dependency to the image content. In this paper, we first introduce a novel structure, a superpixel-based patch, called SuperPatch. The proposed structure, based on superpixel neighborhood, leads to a robust descriptor since spatial information is naturally included. The generalization of the PatchMatch method to SuperPatches, named SuperPatchMatch, is introduced. Finally, we propose a framework to perform fast segmentation and labeling from an image database, and demonstrate the potential of our approach since we outperform, in terms of computational cost and accuracy, the results of state-of-the-art methods on both face labeling and medical image segmentation.

Index Terms—Patch-based method, Superpixels, Patches of superpixels

I. INTRODUCTION

Image segmentation is a useful tool to analyze the image content. The goal of segmentation is to decompose the image into meaningful segments, for instance, to separate objects from the background. A segmentation is computed with respect to some priors such as shape, color or texture. To reduce the computational cost, *superpixel* decomposition methods have been developed for grouping pixels into homogeneous regions, while respecting the image contours (for instance see [1] and references therein). Superpixels are able to drastically decrease the number of elements to process while keeping all the geometrical information that is lost with multi-resolution approaches. Small objects disappear at low resolution levels, whereas they can still be represented with one or several superpixels. Nevertheless, superpixels remain underexploited due to their irregular decomposition of the image content.

Many image processing and computer vision methods use reference images. For instance, for labeling applications, these images can be provided with their ground truth segmentation,

labels, or semantic information that are used to process the input image. In this context, matching algorithms can be useful to find associations between the considered elements. In most frameworks, patch-based approximate nearest neighbor (ANN) search methods are used to find correspondences. Numerous methods have been proposed to find ANN [2], [3], [4], [5] within the same image, and between an image and one or several reference ones. Among these methods, the PatchMatch (PM) method [2] was designed to compute correspondences between pixel-based patches.

When applying PM to large images, or when looking for ANN in a database, the search for good ANN may require many iterations. Therefore, multi-resolution PM [6] can be considered to initialize the ANN correspondence map. However, as usually observed with such coarse-to-fine frameworks, details are lost and a poor ANN is estimated for small scale patterns. A regular decomposition of the image could decrease the problem dimension, but it would not respect the object contours, leading to non accurate processing. In this context, the use of superpixels may be interesting to preserve the image geometry and the respect of the image object contours. Local superpixel-based matching models have been proposed for many applications, *e.g.*, video tracking [7], [8]. However, superpixel-based ANN search algorithms have been little investigated in the literature, and recent works such as [9], [10] that compute superpixel correspondences between the decompositions of two images, use complex models that report prohibitive computational times.

Finally, for ANN matching, the neighborhood information greatly helps in finding good correspondences, as demonstrated in the patch-based literature. Therefore, to jointly decrease the number of elements to process, keep the geometrical information, and find accurate matches, it appears necessary to consider superpixels and to describe them using their neighborhoods in a structure that includes spatial information. Nevertheless, the lack of regularity between two superpixel decompositions makes difficult the use of neighborhood for computing relevant correspondences. Some attempts to use superpixel neighborhood information have been proposed [11], [12]. However, these methods are not adapted to the search of ANN, since they perform a regularization on a graph built from superpixel neighbors but do not include the relative spatial information between superpixels in a dedicated structure.

A. Contributions

In this paper, we propose a novel structure of superpixel neighborhood called *SuperPatch*. Since the superpixel

This study has been carried out with financial support from the French State, managed by the French National Research Agency (ANR) in the frame of the GOTMI project (ANR-16-CE33-0010-01) and the Investments for the future Program IdEx Bordeaux (ANR-10-IDEX-03-02) with the Cluster of excellence CPU and TRAIL (HR-DTI ANR-10-LABX-57).

R. Giraud is with LaBRI, UMR 5800 and IMB, UMR 5251, F-33400 Talence, France. email: remi.giraud@labri.fr. phone: +33(0) 5 4000 6937.

V.-T. Ta is with Bordeaux INP, LaBRI, UMR 5800, F-33400 Talence, France. email: vinh-thong.ta@labri.fr. phone: +33(0) 5 4000 3538.

A. Bugeau is with University of Bordeaux, LaBRI, UMR 5800, F-33400 Talence, France. email: aurelie.bugeau@labri.fr. phone: +33(0) 5 4000 3528.

P. Coupé is with CNRS, LaBRI, UMR 5800, F-33400 Talence, France. email: pierrick.coupe@labri.fr. phone: +33(0) 5 4000 3538.

N. Papadakis is with CNRS, IMB, UMR 5251, F-33400 Talence, France. email: nicolas.papadakis@math.u-bordeaux.fr. phone: +33(0) 5 4000 2116.




Fig. 1. Superpixels vs superpatches for superpixel matching. (a) and (b): two decompositions using [1] and [13]. (c) and (d): superpixel-based [14] and our superpatch-based matching results. The same experiment is performed between (a) and the sheared image (e), with superpixel [14] (f) and superpatch matching results (g). The displacement is illustrated with optical flow representation (h). The more the colored result is close to white, the lower the displacement is.




Fig. 2. Superpixel matching on textured images. Two different parts of two close textures are combined in (a) and (b). (c) and (d): superpixel-based [14] and our superpatch-based matching results where red superpixels indicate wrong matched texture. (e) and (f): two decompositions of a natural textured image using [1] and [13]. Comparison of superpatch matching with color (g) and combination of color and texture features (h).

neighborhoods of two superpatches are not necessarily the same (in terms of shape or number of elements), a generic framework for comparing superpatches is introduced. A novel method, called *SuperPatchMatch* (SPM), that generalizes the PM algorithm [2], is proposed to perform fast and accurate searches of ANN superpatches within images.

To the best of our knowledge, the specific combination of PM with superpixels has been proposed in [14], [15] that propose to match single superpixels using moves similar to PM. For instance in [14], the superpixel features are pre-computed using a learned distance metric, while the reported labeling results do not reach the ones of state-of-the-art methods. In [16], a more restricted framework is considered for optical flow estimation: PM is used to refine the results within selected superpixel bounding boxes. The purpose of our work is thus completely different since we compare neighborhood of structures defined on irregular image sub-domains.

To emphasize the interest of our method, we propose a framework to perform fast segmentation and labeling from an image database. *SuperPatchMatch* is well adapted to deal with huge and constantly growing databases since no learning phase is required, contrary to most existing approaches based on supervised machine learning [17], [18], or recent neural network methods [19], [20]. We apply *SuperPatchMatch* to the challenging Labeled Faces in the Wild (LFW) database [21], where the goal is to extract hair, face, and background within images decomposed into superpixels, and to the segmentation of tumors on non-registered Magnetic Resonance Images (MRI). Finally, *SuperPatchMatch* outperforms, in terms of computational cost and accuracy, state-of-the-art methods.

Fig. 1 and 2 consider several experiments to demonstrate that superpatches enable to find more reliable superpixel ANN than the ones obtained with single superpixel matching [14]. In Fig. 1, two decompositions are computed on the same image using [1] and [13] (Fig. 1(a) and (b)). The aim is to find the best superpixel match between (a) and (b) in terms of superpixel feature (here ℓ_2 -norm on normalized color histograms in RGB space). We display the displacement mag-

nitude of matches with optical flow representation (Fig. 1(h)). When matching only superpixels, as in [14], many outliers are obtained (Fig. 1(c)), while the matching of superpatches provides very accurate ANN (Fig. 1(d)). The same experiment on a sheared image decomposed with [1] (Fig. 1(e)), provides a uniform displacement (Fig. 1(g)) that indicates relevant superpatch matching, and robustness of the proposed structure to geometrical deformations. In Fig. 2, two different parts of two close textures are combined in Fig. 2(a) and (b), and we represent wrong matched texture with red superpixels in Fig. 2(c) and (d). Finally, in Fig. 2, we show that for a natural image containing texture (Fig. 2(e) and (f)), the combination of color and texture features (histogram of oriented gradients [22]) can provide more accurate matching (Fig. 2(h)).

B. Outline

In this paper, we first present related works in Section II. Then, we define the new superpatch structure and a comparison framework between superpatches in Section III. The *SuperPatchMatch* algorithm is next designed to perform superpixel-based ANN search in Section IV. We further emphasize the interest of our method by proposing in Section V a framework to perform labeling from an image database. Finally, we present experiments of face labeling and segmentation of medical images, and *SuperPatchMatch* results outperform the ones of state-of-the-art methods.

II. RELATED WORKS

A. Superpixel Methods

Superpixel decomposition approaches try to group the pixels of an image into meaningful homogeneous regions. They were progressively introduced, for instance, from watershed [23] to Quick shift [24] approaches. In the past years, most decomposition methods start from an initial regular grid and refine the superpixel boundaries by computing a trade-off between color distance and superpixel shape regularity, e.g., [1], [25]. Recently, works such as [26], [27] propose to use gradient

and contour information in the process to further increase the superpixel decomposition accuracy with respect to the image content. Finally, the computational cost is considered since superpixels are mainly used as pre-processing, and recent implementations report real-time performances, *e.g.*, [28].

By considering features at the superpixel scale, the computational complexity of computer vision and image processing tasks can be drastically reduced, while still considering the image geometry and content. Superpixels have therefore become key building blocks of many recent image processing and computer vision pipelines such as multi-class object segmentation [29], [30], [31], [32], body model estimation [33], face and hair labeling [18], data associations across views [12], object localization [34] or contour detection [35]. With these considerations, we propose in this work to use the superpixel representation as the basis of our framework.

B. Including Spatial Information within Image Features

Pixel-based patches enable to describe the pixel neighborhood and to find similar patterns with the same geometric structure. They have progressively proven their efficiency for several applications such as texture synthesis [36] and image denoising [37], and in the design of computer vision descriptors [38], [39] that include spatial information.

Recent works in object retrieval have demonstrated that describing the objects with spatial information enables to reach higher detection accuracy. In [40], [41], Force-Histogram Decomposition descriptors are used to encode the pairwise spatial relations between objects. Deformable part models [42], [43] or adaptive bounding boxes of poselets [44] have also been successfully applied to image retrieval, segmentation or recognition. Finally, the necessity for including spatial information is also studied in [45] that investigates fuzzy approaches to define spatial relationships.

The superpixel itself is not sufficient to provide a robust image descriptor, since the consistency of its neighborhood is not considered. The superpixel neighborhood has been used in [11] for saliency detection based on energy minimization. For each superpixel, the two first adjacent neighbor rings are used in a regularization term. However, the superpixel features are separately included in a data term, leading to a lack of spatial information consistency. The approach is thus dependent on the superpixel decomposition and poorly robust to very irregular decomposition. Consequently, we propose to go further in this work and to take advantage of the superpixel neighborhood to construct a novel representation, namely the superpatch, that naturally includes spatial information.

C. Patch Matching Methods

Patch-based methods have demonstrated state-of-the-art results over various computer vision and image processing applications such as: texture synthesis [36], denoising [37] or super-resolution [46]. These approaches rely on the search of ANN, *i.e.*, similar patches. Many methods were proposed to find ANN within the image itself, between two images or in an entire database [2], [3], [4], [5]. When facing huge databases, dimension reduction methods are usually considered to have




Fig. 3. Superpatch illustration. In blue: circle search of radius R centered on c_i , barycenter of A_i (yellow). The superpatch $\mathbf{A}_i$ is composed by all superpixels having their barycenter within the circle.

fast computation of ANN, but they depend on the size of the data. In this context, the PatchMatch (PM) algorithm [2] is an efficient tool to compute ANN. Within an image itself, the found ANN enable to perform several processings such as image retargeting or completion [2]. Nevertheless, PM can also find matches between several images, and easily handles large databases, since its complexity only depends on the size of the image to process, as shown in [47], [48] where the ANN are used for exemplar-based segmentation of 3D medical images.

In this work, we introduce the SuperPatchMatch method (SPM), that combines both the advantages of the PM algorithm, and the superpixel decomposition of an image, to compute robust correspondences of superpixels using superpatches. The proposed superpatch structure enables to match similar patterns at the superpixel level since it considers the geometrical information between the contained superpixels, which are described by image features such as color or texture.

III. SUPERPATCH

A. Superpatch Definition

Similarly to a patch of pixels, a *superpatch* is a patch (a set) of neighboring superpixels. Let A be an image, decomposed by any superpixel decomposition method, into $|A|$ superpixels such that $A = \{A_i\}_{i \in \{1, \dots, |A|\}}$, where $|\cdot|$ denotes the cardinality, and for two superpixels $A_i, A_j \in A$, $A_i \cap A_j = \emptyset$. A superpatch $\mathbf{A}_i$ is centered on a superpixel A_i and is composed of its neighboring superpixels $A_{i'}$ such that: $\mathbf{A}_i = \{A_{i'}, \text{ with } \|c_i - c_{i'}\| \leq R\}$, with c_i the spatial barycenter of the pixels contained in A_i . In other words, the superpatch centered on a superpixel is defined by considering all superpixels within a fixed radius $R \geq 0$. Note that each superpatch $\mathbf{A}_i$ contains at least the superpixel A_i . Fig. 3 illustrates the superpatch definition.

For the sake of clarity, we denote $\mathcal{I}_i^A = \{i', \text{ with } A_{i'} \in \mathbf{A}_i\}$, the index set of superpixels $A_{i'} \in \mathbf{A}_i$. Each superpixel A_i is described by a set of features F_i^A . These features can be, for instance, the coordinates of c_i , the mean color, or any superpixel descriptors that can be found in the literature.

B. Superpatch Comparison Framework

The comparison between two regular square patches is commonly performed using the sum of squared differences (SSD), computed in a scan order. When considering two superpatches, their number of elements and geometry are generally different, which makes difficult their comparison. In the following, we consider two superpatches $\mathbf{A}_i$ and $\mathbf{B}_j$, in different images A and B . We propose to first register the relative positions of all superpixels within the superpatches. To overlap two superpatches, all positions $c_{j'}$ of superpixels $B_{j'} \in \mathbf{B}_j$ are registered with the vector $v_{ij} = c_i - c_j$, where c_i and c_j are the spatial barycenters of A_i and B_j , respectively. Contrarily to the classical pixel setting, the number of elements and geometry of two superpatches are likely to differ since their construction depends on the initial superpixel decomposition. Therefore, a registered superpixel $B_{j'}$ can overlap with several superpixels $A_{i'}$, and this information has to be considered.

To compute a distance between irregular structures, such as superpixels, [12] proposes to use the editing distance. However, such distance computes one-to-one matching between the structure elements and cannot accurately deal with the overlap of superpixels that requires a one-to-many mapping. Another limitation is that it mixes two different information: superpixel similarities and the cost of removing or adding superpixels. Therefore, this distance should be carefully tuned with respect to the considered application. Consequently, to define a relevant metric between superpatches, it is necessary to consider the geometry of the superpatches within the distance. We propose to define the symmetric distance D between two superpatches $\mathbf{A}_i$ and $\mathbf{B}_j$ as:

$$D(\mathbf{A}_i, \mathbf{B}_j) = \frac{\sum_{i' \in \mathcal{I}_i^A} \sum_{j' \in \mathcal{I}_j^B} w(A_{i'}, B_{j'}) d(F_{i'}^A, F_{j'}^B)}{\sum_{i' \in \mathcal{I}_i^A} \sum_{j' \in \mathcal{I}_j^B} w(A_{i'}, B_{j'})}, \quad (1)$$

where d is the Euclidean distance between the superpixel features $F_{i'}^A$ and $F_{j'}^B$, for instance, average superpixel color or normalized cumulative color histogram, and w is a weight depending on the relative position of the superpixels $A_{i'}$ and $B_{j'}$. Note that we consider an Euclidean distance, but any distance on superpixel features can be computed with d .

1) *Fast distance between superpixels:* To compute the weight between two superpixels $A_{i'}$ and $B_{j'}$, we would ideally like to measure their relative overlapping area, i.e., setting $w(A_{i'}, B_{j'}) = |A_{i'} \cap B_{j'}| / |A_{i'} \cup B_{j'}|$. Nevertheless, this computation requires the expensive count of overlapping pixels that cancels the computational advantage of the superpixel representation. A fast method would be to compare a superpixel $A_{i'}$ to the spatially closest $B_{j'}$, but we propose a more robust framework that considers a spatial distance between the superpixel barycenters. We define the symmetric spatial weight between two superpixels $A_{i'}$ and $B_{j'}$ as:

$$w(A_{i'}, B_{j'}) = \exp(-x_{ij'}^T x_{ij'} / \sigma_1^2) w_s(A_{i'}) w_s(B_{j'}), \quad (2)$$

where $x_{ij'} = c_{j'} - c_{i'} + v_{ij}$ is the relative distance between $c_{i'}$ and $c_{j'}$, $w_s(A_{i'})$ weights the influence of $A_{i'}$ according to its spatial distance to A_i such that $w_s(A_{i'}) = \exp(-\|c_i - c_{i'}\|_2^2 / \sigma_2^2)$, and σ_1 and σ_2 are two scaling parameters. The setting of σ_1 depends of the superpixel decomposition scale.

Since the superpatches have been registered, and the aim is to compare a superpixel $A_i \in \mathbf{A}_i$ to the closest ones $B_j \in \mathbf{B}_j$, σ_1 can be set to half the average superpixel size, i.e., half the average distance between superpixel barycenters, such that $\sigma_1 = \frac{1}{2} \sqrt{h \times w / K}$, for an image of size $h \times w$ pixels decomposed into K superpixels. Finally, σ_2 depends on the superpatch size and can be set to $\sqrt{2}R$ to weight the contribution of closest superpixels.

2) *Generalization of pixel-based patches:* In the limit case where each superpixel only contains one pixel, i.e., $A_{i'} = c_{i'}$, $B_{j'} = c_{j'}$, $\mathbf{A}_i$ and $\mathbf{B}_j$ have the same regular structure and the same number of elements. With $\sigma_1 \rightarrow 0$ and $\sigma_2 \rightarrow \infty$ in (2), $w(A_{i'}, B_{j'}) = w(c_{i'}, c_{j'}) = 1$ if $c_{i'} - c_i = c_{j'} - c_j$ and 0 otherwise, and the proposed distance D (1) is a generalization of the distance between patches, since it reduces to a normalized standard SSD between two pixel-based patches:

$$\begin{aligned} D(\mathbf{A}_i, \mathbf{B}_j) &= \frac{1}{|\mathbf{A}_i|} \sum_{i' \in \mathcal{I}_i^A, j' \in \mathcal{I}_j^B} d(F_{i'}^A, F_{j'}^B) \delta_{c_{i'} - c_i = c_{j'} - c_j}, \\ &= \text{SSD}(\mathbf{A}_i, \mathbf{B}_j), \end{aligned} \quad (3)$$

where $\delta_a = 1$ when a is true and 0 otherwise.

IV. SUPERPATCHMATCH

A. The SuperPatchMatch Algorithm

We propose the SuperPatchMatch method (SPM), an extension of the PatchMatch (PM) algorithm [2] dedicated to our superpatch framework, for fast matching of irregular structures from superpixel decompositions. In this section, only direct adjacent neighborhood relationship needs to be considered to design our algorithm. Nevertheless, as for pixels described by regular patches, we demonstrate in Section V that the proposed framework is significantly more efficient using superpatches. In the following, as in Section III-B, we illustrate the proposed method by considering two superpatches $\mathbf{A}_i$ and $\mathbf{B}_j$, in different images A and B , but our approach can be applied to an entire database, as demonstrated in Section V.

PM is a method that computes pixel-based patch correspondences between two images. The key point of this method is that good correspondences can be propagated to the adjacent patches within an image. The algorithm has three steps: initialization, propagation and random search. The initialization consists in randomly associating each patch of the image A with a patch of the image B , leading to an initial ANN field. The following steps are then iteratively performed to improve the correspondences. The propagation uses the assumption that when a patch in A corresponds to a patch in B , then the adjacent patches in A should also match the adjacent patches in B . The random search consists in a sampling around the current ANN to escape from possible local minima.

The lack of regular geometry between superpatches is the main issue for adapting the PM algorithm. The notion of adjacent patches has to be defined for an irregular superpixel decomposition. For the sake of clarity, the current best ANN of a superpixel $A_i \in A$, is denoted as $B(i) = B_{\mathcal{M}(i)}$, with $\mathcal{M}$ the ANN map which stores, for superpixels in A , the index in B of their corresponding ANN.

1) *SPM initialization step*: For each superpixel $A_i \in A$, we assign a random superpixel $B(i) \in B$. Fig. 4 shows initialization examples. After this step, propagation and random search are iteratively performed to improve the initial matches.




Fig. 4. SPM initialization step. Each superpixel $A_i \in A$ is randomly assigned to a superpixel $B(i) \in B$.

2) *SPM propagation step*: In [2], the propagation step tries to improve the current ANN by using the ones of the directly adjacent neighbors in A . Pixels are processed according to a scan order, e.g., from top-left to bottom-right. The propagation only considers previously processed and directly adjacent neighbors, e.g., top and left. Their ANN are shifted to respect the relative position of pixels in A , providing ANN candidates to the currently processed pixel. With superpixels, the selection of top, left, bottom and right adjacent superpixels is not direct, since there is no regular geometry between them. We propose to define the superpixel scan order from the raw pixel order on A (left to right, top to bottom), and to consider in the propagation step all adjacent superpixels that were processed during the current iteration. The selection of candidates from adjacent neighbors is illustrated in Fig. 5.




Fig. 5. SPM propagation step. According to the scan order, only top-left superpixels of A_1 are considered on even iterations (A_2 , A_3 and A_4). The remaining neighbors (in gray) are tested on odd iterations. Current matches are denoted by $B(i)$, while $C(i)$ represent the new candidates to test as ANN.

When an adjacent superpixel $A_{i'}$ and its ANN $B(i')$ are considered, a neighbor of $B(i')$ is selected as a candidate to improve the correspondence of A_i . However, the ANN $B(i')$ cannot be shifted as done for regular patches, since the superpixels are defined on irregular domains. Therefore, to improve the ANN of A_i , one particular neighbor of $B(i')$, denoted as $C(i') \in B$ is tested. It is given by the superpixel whose relative position to $B(i')$ is the most similar to $\theta_{ii'}$, the angle between A_i and $A_{i'}$. Hence, the ANN candidate $C(i')$ to test is obtained as:

$$C(i') = \underset{k \in \mathcal{N}_{\mathcal{M}(i')}^B}{\operatorname{argmin}} |(\theta_{ii'} + \pi) - \theta_{i'k}|, \quad (4)$$

with $\theta_{i'k}$ the angle between $B(i')$ and its set $\mathcal{N}_{\mathcal{M}(i')}^B$ of adjacent superpixels B_k . Note that all angles are computed from superpixel barycenters. The selection of the candidate for A_3 , which is on top-left position of A_1 , is illustrated in Fig.




Fig. 6. To improve the ANN of the superpixel A_1 , the neighbor A_3 is considered. Its current ANN is $B(3)$. The selected superpixel $C(3)$ to test is the adjacent superpixel of $B(3)$ with the most similar symmetric orientation, i.e., the superpixel B_k which angle to $B(3)$, θ_{3k} , is the closest to $\pi + \theta_{13}$.

6. To keep the same relative position, the new superpixel to test as ANN for A_1 is the neighbor of $B(3)$ that is the closest to its bottom-right position, according to the angle $\theta_{i'k}$.

3) *SPM random search step*: The random search step consists in a sampling around the current ANN to escape from possible local minima [2]. Candidates are selected at an exponentially decreasing distance from the barycenter of the best current match. A random pixel position is computed within decaying boxes, and the superpixel containing this pixel is the candidate to test. Fig. 7 illustrates the random search step, where the boxes are depicted in dotted lines.




Fig. 7. SPM random search step. The sampling is performed at a decreasing distance around the barycenter of the ANN $B(1)$ of A_1 . The superpixels containing the selected positions (crosses) are the candidates to test.

B. Library of Training Images

One advantage of the SPM algorithm is that its complexity only depends on the size of the image to process and not on the size of the compared image database. This important fact enables SPM to perform fast ANN searches within a large database with no increase on the computational time.

All example images within the database are grouped into a single library T . In the case of a training database, SPM steps are adapted so the ANN can be found within all images. The initialization is extended: the ANN is randomly selected within T . The propagation step still tests the shifted ANN of the neighbors, that are not necessarily in the same training image. Finally, the random search is performed within the current best image, and within a random image in T , as in [47].

C. Multiple SPM

Contrary to PM, that only estimates one ANN, SPM computes k -ANN matches in the library T , since the diversity of information from multiple ANN may help to perform more

accurate processing. In the literature, an extension of the original PM algorithm to the k -ANN case has been proposed in [6]. The suggested strategy is to build a constantly updated data structure of the best visited correspondences. However, to parallelize such an approach, the current test image must be split into several parts which leads to boundary issues. Therefore, we chose to implement the k -ANN search by k fully independent SPM, leading to a simpler scheme.

V. APPLICATION TO IMAGE LABELING

To demonstrate the interest of the superpatch structure and the SPM algorithm, we adapt our approach to exemplar-based labeling. We consider two experiments: face labeling on the LFW dataset [21], and segmentation and labeling on non-registered medical images from the BRATS dataset [49].

A. Label Fusion Method

The proposed algorithm is particularly interesting for labeling applications. The superpixel decomposition segments the image into homogeneous regions that try to respect existing contours, and SPM finds superpixel-based correspondences whose labels can be transferred. In this application, a library T of training images with their label ground truths is considered, and SPM provides k -ANN matches. We denote as $l(T_j) \in \{1, \dots, M\}$ the label of the training superpixel T_j contained in T . The labels of the selected ANN within T are merged by a patch-based label fusion [50], inspired from [37].

At the end of the ANN search, k -ANN are estimated for all superpixels in the test image A . To obtain the final labeling, for a superpixel A_i and $\mathcal{K}_i^m = \{T_j\}$ the set of its k -ANN matches with label $l(T_j) = m$, its label fusion map $L_m(A_i)$ is defined by:

$$L_m(A_i) = \frac{\sum_{T_j \in \mathcal{K}_i^m} \omega(A_i, T_j)}{\sum_{m=1}^M \sum_{T_j \in \mathcal{K}_i^m} \omega(A_i, T_j)}, \quad (5)$$

where $\omega(A_i, T_j)$ is the weight contributing to label m , and depends on the similarity between the superpatch $\mathbf{A}_i \in A$, and the ANN superpatch $\mathbf{T}_j \in T$. This label map L_m gives the probability of assigning the label m to the superpixel A_i .

Some applications can also deal with registered images, where structures of interest between A and images of T are spatially close. Therefore, good superpatch matches should not be spatially too far in the image domain. In this case, to enforce the spatial coherency of the selected k -ANN, each labeling contribution is weighted by the spatial distance between the central superpixels barycenters $c_i \in A$ and $c_j \in T$:

$$\omega(A_i, T_j) = \exp \left(1 - \left(\frac{D(\mathbf{A}_i, \mathbf{T}_j)}{h(A_i)^2} + \frac{\|c_i - c_j\|_2}{\beta^2} \right) \right), \quad (6)$$

where $h(A_i)^2 = \alpha^2 \min_{T_j \in \mathcal{K}_i^m} (D(\mathbf{A}_i, \mathbf{T}_j) + \epsilon)$, with $\epsilon \rightarrow 0$, and α and β are scaling parameters. With the function $h(A_i)$, the distance of the current contribution is divided by the minimal distance among all k -ANN contributions. For each superpixel A_i , the final labeling map $\mathcal{L}(A_i)$ is obtained with the label of highest probability:

$$\mathcal{L}(A_i) = \underset{m \in \{1, \dots, M\}}{\operatorname{argmax}} L_m(A_i). \quad (7)$$

The relation (7) gives a superpixel-wise decision that may have some irregularities. As in [11], we can use (5) as a multi-label data term and consider the following regularization problem, that consists in minimizing the energy J , defined on the graph built from adjacent superpixels:

$$J(\mathcal{L}) = \sum_{i=1}^{|A|} \left(\sum_{i' \in \mathcal{I}_i^A} \exp \left(-\frac{d(F_i, F_{i'})}{\gamma} \right) \delta_{i, i'} + 1 - L_{\mathcal{L}(A_i)}(A_i) \right), \quad (8)$$

where γ is a regularization parameter, the data term $1 - L_{\mathcal{L}(A_i)}(A_i)$ is close to 0 (respectively 1) when the probability of label $\mathcal{L}(A_i)$ is high (respectively low), and $\delta_{i, i'} = 1$ when $\mathcal{L}(A_i) \neq \mathcal{L}(A_{i'})$ and 0 otherwise.

B. Face Labeling Experiments

Face segmentation and labeling are challenging tasks due to several issues such as the diversity of hair styles, background, color skins, or occlusions. We evaluate the proposed SPM approach for face labeling on the funneled version of the Labeled Faces in the Wild (LFW) dataset [21]. The dataset contains 2927 images of size 250×250 pixels, that have been coarsely aligned [51], and segmented into 225 to 250 superpixels. LFW is a widely used database for validating new methods based on superpixels since it contains decompositions with associated superpixel-wise ground truths, and comparisons with state-of-the-art methods are not biased by the ground truth superpixel decomposition one would have to compute.

1) *Parameter settings*: SPM was implemented with MATLAB using C-MEX code. Our experiments are performed on a standard Linux server of 16 cores at 2.6 GHz with 100 GB of RAM. To compare to [18], [19], we use the same 1 500 training images, and the same 927 images for testing. Nevertheless, we could use all images in a leave-one-out procedure since our method does not need any training step.

The number of SPM iterations is set to 5, as in [2]. We only use a ℓ_2 -norm between histogram of oriented gradients (HoG) [22] as distance d in (1). In Eq. (2), since the images are $h \times w = 250 \times 250$ pixels, and decomposed into approximately $K = 250$ superpixels, σ_1 is set to 8. In Eq. (6), parameters α and β are respectively set to 2 and 4. Finally, we set γ to 0.5 and use the α -expansion algorithm [52] to minimize (8). The reported times for SPM in Fig. 8 (b) include k -ANN searches, label fusion and the complete labeling with regularization.

2) *Influence of the superpatch size*: We first investigate the influence of the superpatch size and number of ANN. Fig. 8 represents the superpixel-wise labeling accuracy and computational time. The labeling accuracy is increased with our superpatch structure. Best results are obtained with $R=50$ pixels (95.08% with $k=50$ ANN). Such superpatch size corresponds in average to the capture of the three neighboring rings of superpixels, since superpixels are approximately of size 16×16 pixels. Fig. 8 also represents the corresponding ROC curves obtained with $k=50$ ANN for the three classes (face, background, hair). Without the superpatch structure, *i.e.*, only computing the distance on central superpixels ($R=0$ pixels), worse ANN are found, decreasing the labeling accuracy




Fig. 8. Influence of the superpatch size on superpixel-wise labeling accuracy (a) with corresponding computational time (b). ROC curves with area under curves (AUC) for several superpatch sizes ($k=50$) for face (c), background (d) and hair (e). $R = 50$ pixels gives the best results for all labels.

TABLE I
LABELING ACCURACY ON LFW

Method	Superpixel accuracy	Pixel accuracy	Computational time
PatchMatch	87.73%	87.02%	3.940s
Spatial CRF [18]	93.95%	<i>not reported</i>	<i>not reported</i>
CRBM [18]	94.10%	<i>not reported</i>	<i>not reported</i>
GLOC [18]	94.95%	<i>not reported</i>	0.323s
DCNN [19]	<i>not reported</i>	95.24%	<i>not reported</i>
SuperPatchMatch	95.08%	95.43%	0.255s

Computational times are given per subject. SPM results are obtained with $k=50$ ANN, and $R=50$ pixels. The presented values are the published results, therefore, some evaluation metrics could not be reported.

(93.29% with $k=50$ ANN). The superpatch size must be large enough to capture the information contained within the superpixel neighborhood. However, with too large superpatches, *i.e.*, $R > 50$ pixels, too many neighboring superpixels contribute, leading to less relevant ANN and less accurate labeling. Note that we propose in (6) a slight improvement of the label fusion step to take into account the LFW database registration but we obtain very comparable results (95.00% instead of 95.08%) without any position a priori, *i.e.*, $\beta=\infty$ in (6).

Fig. 9 illustrates the regularization process. Labeling probabilities (5) obtained from SPM are displayed for each label (Fig. 9(d)). The spatial regularization (8) gives more consistent results (Fig. 9(f)) than taking the label of highest probability (Fig. 9(e)). Finally, Fig. 10 shows the superpatch influence on labeling for various examples. Labeling failures are mostly due to high similarity between hair and background, or inaccurate superpixel segmentation.

3) *Comparison with the state-of-the-art methods*: SPM is compared to the recent methods applied to the LFW database in Table I. In [18], the GLOC (GLOBal and LOCal) method uses a restricted Boltzmann machine as complement of a conditional random field labeling [53]. This combination reduces the error in face labeling of single models which do not use global shape priors, at the expense of a higher computational cost. In [19], a method based on a deep convolutional neural network (DCNN) is proposed. For all compared methods, learning steps, that can be up to several hours, are necessary to train the models. Moreover, they consider priors learned from semantic information into the process, *e.g.*, hair label should be on top of face label in the segmentation. We also provide the results of a pixel-wise PM applied with the same framework, where a SSD between patches of size 9×9 pixels in RGB color space is used as distance.

To compare to all methods, we provide in Table I superpixel and pixel-wise accuracy results. The presented values are the results published by the authors, therefore, all the evaluation metrics could not be reported. SPM superpixel-wise labeling accuracy outperforms the ones of the compared methods (95.08%), while being performed on basic features, and faster (0.255s per subject) than the best compared method with reported computational time. The pixel-wise accuracy of SPM (95.43%) also outperforms the reported result of the DCNN architecture [19], that has been optimized to perform on the LFW dataset. Note that the increase of pixel accuracy over superpixel accuracy demonstrates that our method mostly fails at labeling small and stretched superpixels. This comes from the initial LFW segmentation that may produce inaccurate color clustering and allows irregular superpixel shapes.

The global computational time is another important comparison point. SPM outperforms the compared methods in term of labeling accuracy without any training step. Contrary to other methods, with SPM, computational efforts needed for learning are canceled, and new training images are directly considered in the library. To illustrate this point, for each processed image, we add the remaining test ones to the library. This way, SPM reaches 95.26% of superpixel-wise labeling accuracy. This result highlights the impact of the image diversity within the database, which leads to find more accurate ANN. Moreover, results are obtained with no computational time increase, since the algorithm complexity only depends on the test image size. Hence, SPM easily integrates new images in the database, and provides very competitive results in limited computational time, without model or shape *priors*.

4) *Robustness to superpixel decomposition method*: To emphasize the robustness of our method to the used superpixel method, we have segmented the test images with another method [1] that produces more regular superpixels (see an example in Fig. 11). The new decompositions are computed with respect to the ground truth label mask of each image. Hence, they are still constrained by the initial segmentations provided with LFW but only on the edges of each class (hair, face, background). Even with test and training decompositions computed with different methods, we get similar superpixel-wise labeling accuracy (95.05%), showing that our method can compare superpixel neighborhoods of various shapes.

C. Non-Registered MRI Segmentation Experiments

To demonstrate the robustness of the superpatch structure and the proposed framework, we apply SPM to brain tu-




Fig. 9. Labeling results with SPM. (a) Image. (b) Superpixel decomposition. (c) Associated ground truth. (d) Label fusion maps for the 3 classes (face, background, hair). (e) Labeling with the highest probability from (7). (f) Labeling after regularization.




Fig. 10. Labeling examples obtained with SPM for $R=0$, and $R=50$ pixels, *i.e.*, using superpatches, with superpixel-wise labeling accuracy.




Fig. 11. Re-segmentation of the LFW dataset. (a) Initial image. (b) Ground truth labels. (c) LFW initial decomposition. (d) Decomposition using [1].

mor segmentation on multi-modal non-registered Magnetic Resonance Images (MRI). Classical patch-based and multi-atlas structure segmentation methods are based on registered subjects. Consequently, they cannot be efficiently applied in this non-registered context, due to the substantial variation in tumor shape and locations. Superpixels enable to better capture the tumor geometry, thus increasing the segmentation accuracy. Superpixel and supervoxel-based approaches have been applied to tumor segmentation [54]. However, in this work, the neighborhood is not considered and the ANN search is exhaustive, and computed on a large multi-modal histogram descriptor, leading to prohibitive computational time.

SPM can be efficiently applied to tumor segmentation since it quickly finds good correspondences without image registration, and uses the superpixel neighborhood to improve the matching. In this application, the segmentation is computed

from a superpixel decomposition [1], then each region (tumor or background) is labeled with SPM.

We present results obtained on the MICCAI multi-modal Brain Tumor Segmentation (BRATS) dataset [49]. This challenging dataset contains real and simulated patient data, with overall poor resolution and large variation of tumor shape and position. For both types, high grade (HG) and low grade (LG) tumors are provided with four modalities: T1, contrast enhanced T1 (T1C), T2, and FLAIR. Overall, there are 20 and 10 real patient data with respectively HG and LG tumors, and 25 images for both HG and LG simulated tumor data. We use the same SPM parameters as in Section V-B, taking a multi-modal histogram, containing the levels of gray intensity on all MRI modalities as descriptor for superpatch matching, and performing the regularization (8) at the pixel scale to compare with pixel-wise ground truths. Each subject is segmented by the remaining of its type in a leave-one-out procedure.

In Fig. 12, we show several tumor segmentation results for all data types. In Table II, we compare results obtained using different descriptor structures: patch-based [48], superpixel-based [54], and superpatch-based ($R=25$ pixels). We use the Dice coefficient [55] as evaluation metric, measuring the overlap between the automatically segmented structure and the ground truth. The superpixel-based approach appears very limited since it fails at capturing the tumor context and their




Fig. 12. Examples of patch, superpixel and superpatch-based tumor segmentation results. The four modalities are displayed for simulated HG data.

TABLE II

DICE COEFFICIENT AND COMPUTATIONAL TIME RESULTS FOR DIFFERENT STRUCTURE DESCRIPTORS

Method	Simulated Data		Real Data		Computational time
	HG	LG	HG	LG	
Superpixel-based	73.95%	41.95%	44.63%	38.07%	0.33s
Patch-based	69.11%	49.55%	52.59%	68.92%	2.69s
Superpatch-based	90.75%	82.40%	61.87%	73.34%	0.94s

location in other images. Regular patches are also limited in this context, due to the variations in the structure shapes. Superpatches provide a robust descriptor, since they follow image intensities and capture the superpixel neighborhood, leading to more accurate segmentation. These experiments demonstrate that superpatches within the SPM framework provide fast and accurate segmentation results even on non-registered multi-modal images with poor resolution.

VI. CONCLUSION AND PERSPECTIVES

In this paper, we propose a new structure based on patches of superpixels that can use irregular and non stable image decompositions. These superpatches include neighborhood information and lead to more accurate matching. We also introduce SuperPatchMatch, a general and novel correspondence algorithm of superpatches.

We have demonstrated the interest of our framework by obtaining state-of-the-art results for face labeling and tumor segmentation on non-registered MRI. SuperPatchMatch does not need any learning phase, that can be up to several hours for many methods of the literature. By including spatial consistency, superpatches are able to reach the accuracy of highly tuned approaches, and provide more reliable descriptors than single superpixels.

Our work opens new insights for future adaptations to superpixel-based methods, *e.g.*, segmentation [34], [43], labeling [14], saliency detection [11], or color and style transfer [10]. For instance, SuperPatchMatch can be considered for defining good ANN initializations at the pixel level, when the size of the database is too large. A possible application is the optical flow initialization, instead of multi-resolution schemes, to better capture large displacements of small objects.

REFERENCES

- [1] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Süsstrunk, "SLIC superpixels compared to state-of-the-art superpixel methods," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 34, no. 11, pp. 2274–2282, 2012.
- [2] C. Barnes, E. Shechtman, A. Finkelstein, and D. Goldman, "PatchMatch: A randomized correspondence algorithm for structural image editing," *ACM Trans. Graph.*, vol. 28, no. 3, 2009.
- [3] M. Muja and D. G. Lowe, "Fast approximate nearest neighbors with automatic algorithm configuration," vol. 2, 2009, pp. 331–340.
- [4] S. Korman and S. Avidan, "Coherency sensitive hashing," in *Proc. IEEE ICCV*, 2011, pp. 1607–1614.
- [5] I. Olonetsky and S. Avidan, "TreeCANN - k-d tree coherence approximate nearest neighbor algorithm," in *Proc. ECCV*, 2012, pp. 602–615.
- [6] C. Barnes, E. Shechtman, D. B. Goldman, and A. Finkelstein, "The generalized PatchMatch correspondence algorithm," in *Proc. ECCV*, 2010, pp. 29–43.
- [7] S. Wang, H. Lu, F. Yang, and M. H. Yang, "Superpixel tracking," in *Proc. IEEE ICCV*, 2011, pp. 1323–1330.
- [8] M. Reso, J. Jachalsky, B. Rosenhahn, and J. Ostermann, "Temporally consistent superpixels," in *Proc. IEEE ICCV*, 2013, pp. 385–392.
- [9] J. Rabin, S. Ferradans, and N. Papadakis, "Adaptive color transfer with relaxed optimal transport," in *Proc. IEEE ICIP*, 2014, pp. 4852–4856.
- [10] J. Liu, W. Yang, X. Sun, and W. Zeng, "Photo stylistic brush: Robust style transfer via superpixel-based bipartite graph," *arXiv preprint arXiv:1606.03871*, 2016.
- [11] S.-C. Pei, W.-W. Chang, and C.-T. Shen, "Saliency detection using superpixel belief propagation," in *Proc. IEEE ICIP*, 2014, pp. 1135–1139.
- [12] R. Sawhney, F. Li, and H. I. Christensen, "GASP: Geometric association with surface patches," in *Proc. 3DV*, 2014, pp. 107–114.
- [13] P. Buysens, M. Toutain, A. Elmoataz, and O. Lézoray, "Eikonal-based vertices growing and iterative seeding for efficient graph-based segmentation," in *Proc. IEEE ICIP*, 2014, pp. 4368–4372.
- [14] S. Gould, J. Zhao, X. He, and Y. Zhang, "Superpixel graph label transfer with learned distance metric," in *Proc. ECCV*, 2014, pp. 632–647.
- [15] J. Zheng and Z. Li, "Superpixel based patch match for differently exposed images with moving objects and camera movements," in *IEEE ICIP*, 2015, pp. 4516–4520.
- [16] J. Lu, H. Yang, D. Min, and M. N. Do, "PatchMatch filter: Efficient edge-aware filtering meets randomized search for fast correspondence field estimation," in *Proc. IEEE CVPR*, 2013, pp. 1854–1861.
- [17] X. He, R. Zemel, and M. Carreira-Perpiñán, "Multiscale conditional random fields for image labeling," in *Proc. IEEE CVPR*, vol. 2, 2004.
- [18] A. Kae, K. Sohn, H. Lee, and E. Learned-Miller, "Augmenting CRFs with Boltzmann machine shape priors for image labeling," in *Proc. IEEE CVPR*, 2013, pp. 2019–2026.
- [19] S. Liu, J. Yang, C. Huang, and M. Yang, "Multi-objective convolutional learning for face labeling," in *Proc. IEEE CVPR*, 2015, pp. 3451–3459.
- [20] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *Proc. IEEE CVPR*, 2015, pp. 3431–3440.
- [21] G. Huang, M. Ramesh, T. Berg, and E. Learned-Miller, "Labeled faces in the wild: A database for studying face recognition in unconstrained environments," *Tech. Rep. 07-49, Univ. of Massachusetts, Amherst*, vol. 1, no. 2, 2007.

- [22] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *Proc. IEEE CVPR*, 2005, pp. 886–893.
- [23] L. Vincent and P. Soille, "Watersheds in digital spaces: An efficient algorithm based on immersion simulations," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 13, no. 6, pp. 583–598, 1991.
- [24] A. Vedaldi and S. Soatto, "Quick shift and kernel methods for mode seeking," in *Proc. ECCV*, 2008, pp. 705–718.
- [25] Z. Li and J. Chen, "Superpixel segmentation using linear spectral clustering," in *Proc. IEEE CVPR*, 2015, pp. 1356–1363.
- [26] V. Machairas, M. Faessel, D. Cárdenas-Peña, T. Chabardes, T. Walter, and E. Decencière, "Waterpixels," *IEEE Trans. Image Process.*, vol. 24, no. 11, pp. 3707–3716, 2015.
- [27] R. Giraud, V.-T. Ta, and N. Papadakis, "SCALP: Superpixels with contour adherence using linear path," in *Proc. ICPR*, 2016, pp. 2374–2379.
- [28] Z. Ban, J. Liu, and J. Fouriaux, "GLSC: LSC superpixels at over 130 FPS," *J. Real-Time Image Process.*, pp. 1–12, 2016.
- [29] S. Gould, J. Rodgers, D. Cohen, G. Elidan, and D. Koller, "Multi-class segmentation with relative location prior," *Int. J. Comput. Vis.*, vol. 80, no. 3, pp. 300–316, 2008.
- [30] J. Tighe and S. Lazebnik, "SuperParsing: Scalable nonparametric image parsing with superpixels," in *Proc. ECCV*, 2010, pp. 352–365.
- [31] Y. Yang, S. Hallman, D. Ramanan, and C. Fowlkes, "Layered object detection for multi-class segmentation," in *Proc. IEEE CVPR*, 2010, pp. 3113–3120.
- [32] M. Mostajabi, P. Yadollahpour, and G. Shakhnarovich, "Feedforward semantic segmentation with zoom-out features," in *Proc. IEEE CVPR*, 2015, pp. 3376–3385.
- [33] G. Mori, "Guiding model search using segmentation," in *Proc. IEEE ICCV*, 2005, pp. 1417–1423.
- [34] B. Fulkerson, A. Vedaldi, and S. Soatto, "Class segmentation and object localization with superpixel neighborhoods," in *Proc. IEEE ICCV*, 2009, pp. 670–677.
- [35] P. Arbeláez, M. Maire, C. Fowlkes, and J. Malik, "Contour detection and hierarchical image segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 33, no. 5, pp. 898–916, 2011.
- [36] A. Efros and T. Leung, "Texture synthesis by non-parametric sampling," in *Proc. IEEE ICCV*, 1999, pp. 1033–1038.
- [37] A. Buades, B. Coll, and J.-M. Morel, "A non-local algorithm for image denoising," in *Proc. IEEE CVPR*, 2005, pp. 60–65.
- [38] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *Int. J. Comput. Vis.*, vol. 60, no. 2, pp. 91–110, 2004.
- [39] H. Bay, T. Tuytelaars, and L. Van Gool, "SURF: Speeded up robust features," in *Proc. ECCV*, 2006, pp. 404–417.
- [40] M. Garnier, T. Hurtut, and L. Wendling, "Object description based on spatial relations between level-sets," in *Proc. DICTA*, 2012, pp. 1–7.
- [41] M. Clément, M. Garnier, C. Kurtz, and L. Wendling, "Color object recognition based on spatial relations between image layers," in *Proc. VISAPP*, 2015, pp. 427–434.
- [42] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan, "Object detection with discriminatively trained part based models," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 32, no. 9, pp. 1627–1645, 2010.
- [43] E. Trulls, S. Tsogkas, I. Kokkinos, A. Sanfeliu, and F. Moreno-Noguer, "Segmentation-aware deformable part models," in *Proc. IEEE CVPR*, 2014, pp. 168–175.
- [44] G. Sharma, F. Jurie, and C. Schmid, "Expanded parts model for human attribute and action recognition in still images," in *Proc. IEEE CVPR*, 2013, pp. 652–659.
- [45] I. Bloch, "Fuzzy spatial relationships for image processing and interpretation: A review," *Image and Vision Comp.*, vol. 23, no. 2, pp. 89–110, 2005.
- [46] W. Freeman, T. Jones, and E. Pasztor, "Example-based super-resolution," *IEEE Trans. Comp. Graph. App.*, vol. 22, no. 2, pp. 56–65, 2002.
- [47] W. Shi, J. Caballero, C. Ledig, X. Zang, W. Bai, K. Bhatia, A. Marvao, T. Dawes, D. O'Regan, and D. Rueckert, "Cardiac image super-resolution with global correspondence using multi-atlas PatchMatch," in *Proc. MICCAI*, 2013, pp. 9–16.
- [48] R. Giraud, V.-T. Ta, N. Papadakis, J. V. Manjón, D. L. Collins, P. Coupé, and the Alzheimer's Disease Neuroimaging Initiative, "An optimized PatchMatch for multi-scale and multi-feature label fusion," *NeuroImage*, vol. 124, pp. 770–782, 2016.
- [49] B. H. Menze, A. Jakab, S. Bauer, J. Kalpathy-Cramer, K. Farahani, J. Kirby, Y. Burren, N. Porz, J. Slotboom, R. Wiest *et al.*, "The multimodal brain tumor image segmentation benchmark (BRATS)," *IEEE Trans. Med. Imaging*, vol. 34, no. 10, pp. 1993–2024, 2015.
- [50] P. Coupé, J. V. Manjón, V. Fonov, J. Pruessner, M. Robles, and D. L. Collins, "Patch-based segmentation using expert priors: application to hippocampus and ventricle segmentation," *NeuroImage*, vol. 54, no. 2, pp. 940–954, 2011.
- [51] G. Huang, V. Jain, and E. Learned-Miller, "Unsupervised joint alignment of complex images," in *Proc. IEEE ICCV*, 2007, pp. 1–8.
- [52] Y. Boykov, O. Veksler, and R. Zabih, "Fast approximate energy minimization via graph cuts," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, no. 11, pp. 1222–1239, 2001.
- [53] J. Lafferty, A. McCallum, and F. Pereira, "Conditional random fields: Probabilistic models for segmenting and labeling sequence data," in *Proc. ICML*, 2001, pp. 282–289.
- [54] H. Wang and P. A. Yushkevich, "Multi-atlas segmentation without registration: A supervoxel-based approach," in *Proc. MICCAI*, 2013, pp. 535–542.
- [55] A. Zijdenbos, B. Dawant, R. Margolin, and A. Palmer, "Morphometric analysis of white matter lesions in MR images: method and validation," *IEEE Trans. Med. Imaging*, vol. 13, no. 4, pp. 716–724, 1994.



Rémi Giraud received the M.Sc. in telecommunications at ENSEIRB-MATMECA School of Engineers, and the M.Sc. in signal and image processing from the University of Bordeaux, France, in 2014. Since, he is pursuing his Ph.D. at the Laboratoire Bordelais de Recherche en Informatique in the field of image processing. His research areas mainly include computer vision and image processing applications with patch-based and superpixel methods.



Vinh-Thong Ta received the M.Sc. and Doctoral degrees in computer science from the University of Caen Basse-Normandie, France, in 2004 and 2009, respectively. From 2009 to 2010, he was an Assistant Professor in computer science with the School of Engineers of Caen, France. Since 2010, he is an Associate Professor with the Computer Science Department, School of Engineers ENSEIRB-MATMECA. His research mainly concerns image and data processing.



Aurélie Bugeau received the Ph.D. degree in signal processing from the University of Rennes, France, in 2007. She is an Associate Professor with the University of Bordeaux and was involved in research with the Laboratoire Bordelais de Recherche en Informatique. Her main research interests include patch-based and graph-based methods for image and video processing, inpainting, and editing.



Pierrick Coupé received a M. Eng degree in Biomedical engineering from the University of Technology of Compiègne in 2003, a M.Sc. in Image and Signal processing degree from the University of Rennes I in 2004, and obtained a Ph.D. degree in Image processing from the University of Rennes I in 2008. Since 2011, he is a full time CNRS researcher at the Laboratoire Bordelais de Recherche en Informatique. His research mainly focuses on medical image enhancement and analysis, and computer-aided diagnosis.



Nicolas Papadakis received the Graduate degree in applied mathematics from the National Institute of Applied Sciences, Rouen, France, in 2004, and the Ph.D. degree in applied mathematics from the University of Rennes, France, in 2007. He is currently a Researcher from the Centre National de la Recherche Scientifique, Institut de Mathématiques de Bordeaux, France. His main research interests include tracking, motion estimation, and optimal transportation for image processing.