



HAL
open science

Prior-based Hierarchical Segmentation Highlighting Structures of Interest

Amin Fehri, Santiago Velasco-Forero, Fernand Meyer

► **To cite this version:**

Amin Fehri, Santiago Velasco-Forero, Fernand Meyer. Prior-based Hierarchical Segmentation Highlighting Structures of Interest. International Symposium on Mathematical Morphology, May 2017, Fontainebleau, France. hal-01484281

HAL Id: hal-01484281

<https://hal.science/hal-01484281>

Submitted on 8 Mar 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.

Prior-based Hierarchical Segmentation Highlighting Structures of Interest

Amin Fehri, Santiago Velasco-Forero and Fernand Meyer
{amin.fehri,santiago.velasco,fernand.meyer@mines-paristech.fr}

MINES ParisTech, PSL Research University, Center for Mathematical Morphology

Abstract. Image segmentation is the process of partitioning an image into a set of meaningful regions according to some criteria. Hierarchical segmentation has emerged as a major trend in this regard as it favors the emergence of important regions at different scales. On the other hand, many methods allow us to have prior information on the position of structures of interest in the images. In this paper, we present a versatile hierarchical segmentation method that takes into account any prior spatial information and outputs a hierarchical segmentation that emphasizes the contours or regions of interest while preserving the important structures in the image. Several applications are presented that illustrate the method versatility and efficiency.

Keywords: Mathematical Morphology, Hierarchies, Segmentation, Prior-based Segmentation, Stochastic Watershed.

1 Introduction

In this paper, we propose a method to take advantage of any prior spatial information previously obtained on an image to get a hierarchical segmentation of this image that emphasizes its regions of interest, allowing us to get more details in the designated regions of interest of an image while still preserving its strong structural information.

Potential applications are numerous. When having a limited storage capacity (for very large images for example), this would allow us to keep details in the regions of interest as a priority. Similarly, in situations of transmission with limited bandwidth, one could first transmit the important information of the image: the details of the face for a video-call, the pitch and the players for a soccer game and so on. One could also use such a tool as a preprocessing one, for example to focus on an individual from one camera view to the next one in video surveillance tasks. Finally, from an artistic point of view, the result is interesting and similar to a combination of focus and cartoon effects. Some of these examples are illustrated in this paper.

Image segmentation has been shown to be inherently a multi-scale problem [10]. That is why hierarchical segmentation has become the major trend in image segmentation and most top-performance segmentation techniques [3][22][23][17]

fall into this category: hierarchical segmentation does not output a single partition of the image pixels into sets but instead a single multi-scale structure that aims at capturing relevant objects at all scales. Researches on this topic are still vivid as differential area profiles [21], robust segmentation of high-dimensional data [9] as well as theoretical aspects regarding the concept of partition lattice [26][24] and optimal partition in a hierarchy [11][12][31]. Our goal in this paper is to develop a hierarchical segmentation algorithm that focuses on certain pre-determined zones of the image. The hierarchical aspect also allows us, for tasks previously described, to very simply tune the level of details wanted depending on the application.

Furthermore, our algorithm is very versatile, as the spatial prior information that it uses can be obtained by any of the numerous learning-based approaches proposed over the last decades to roughly localize objects [20] [13] [25]. In this regard, our work joins an important research point that consists in designing approaches to incorporate prior knowledge in the segmentation, as shape prior on level sets [4], star-shape prior by graph-cut [29], use of a shape prior hierarchical characterization obtained with deep learning [5], or related work making use of stochastic watershed to perform targeted image segmentation [16].

The remainder of the paper is organized as follows. Part 2 explains how we construct and use graph-based hierarchical segmentation. Then part 3 specifies how we use prior information on the image to obtain hierarchies with regionalized fineness. Several examples of applications of this method are described in part 4. Finally, conclusions and perspectives are presented in part 5.

2 Hierarchies and partitions

2.1 Graph-based hierarchical segmentation

Obtaining a suitable segmentation directly from an image is very difficult. This is why it is often make use of hierarchies to organize and propose interesting contours by valuating them. In this section, we remind the reader how to construct and use graph-based hierarchical segmentation.

For each image, let us suppose that a fine partition is produced by an initial segmentation (for instance a set of superpixels [1,15] or the basins produced by a classical watershed algorithm [18]) and contains all contours making sense in the image. We define a dissimilarity measure between adjacent tiles of this fine partition. One can then see the image as a graph, the *region adjacency graph* (RAG), in which each node represents a tile of the partition; an edge links two nodes if the corresponding regions are neighbors in the image; the weight of the edge is equal to the dissimilarity between these regions. Working on the RAG is much more efficient than working on the image, as there are far less nodes in the RAG than there are pixels in the image.

Formally, we denote this graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{W})$, where \mathcal{V} corresponds to the image domain or set of pixels/fine regions, $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$ is the set of edges linking neighbour regions, $\mathbf{W} : \mathcal{E} \rightarrow \mathbb{R}^+$ is the dissimilarity measure usually based on local gradient information (or color or texture), for instance $\mathbf{W}(i, j) \propto |\mathbf{I}(v_i) - \mathbf{I}(v_j)|$ with $\mathbf{I} : \mathcal{V} \rightarrow \mathbb{R}$ representing the image intensity.

The edge linking the nodes p and q is designated by e_{pq} . A path is a sequence of nodes and edges: for example the path linking the nodes p and s is the set $\{p, e_{pt}, t, e_{ts}, s\}$. A *connected subgraph* is a subgraph where each pair of nodes is connected by a path. A *cycle* is a path whose extremities coincide. A *tree* is a connected graph without cycle. A *spanning tree* is a tree containing all nodes. A *minimum spanning tree* (MST) $\mathcal{MST}(\mathcal{G})$ of a graph \mathcal{G} is a spanning tree with minimal possible weight, obtained for example using the Boruvka algorithm (the weight of a tree being equal to the sum of the weights of its edges). A *forest* is a collection of trees.

A *partition* π of a set \mathcal{V} is a collection of subsets of \mathcal{V} , such that the whole set \mathcal{V} is the disjoint union of the subsets in the partition, i.e., $\pi = \{R_1, R_2, \dots, R_k\}$, such that $\forall i, R_i \subseteq \mathcal{V} ; \forall i \neq j, R_i \cap R_j = \emptyset ; \bigcup_i R_i = \mathcal{V}$.

Cutting all edges of the $\mathcal{MST}(\mathcal{G})$ having a valuation superior to a threshold λ leads to a minimum spanning forest (MSF) $\mathcal{F}(\mathcal{G})$, i.e. to a partition of the graph. Note that the obtained partition is the same that one would have obtain by cutting edges superior to λ directly on \mathcal{G} [19]. Since working on the $\mathcal{MST}(\mathcal{G})$ is less costly and provides similar results regarding graph-based segmentation, we work only with the $\mathcal{MST}(\mathcal{G})$ in the sequel.

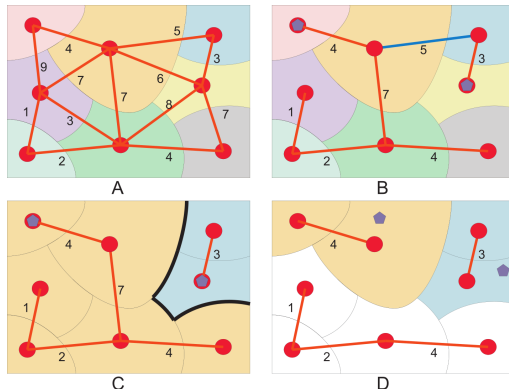


Fig. 1. **A:** a partition represented by an edge-weighted graph; **B:** a minimum spanning tree of the graph, with 2 markers in blue: the highlighted edge in blue is the highest edge on the path linking the two markers; **C:** the segmentation obtained when cutting this edge; **D** blue and orange domain are the domains of variation of the two markers generating the same segmentation.

So cutting edges by decreasing valuations gives an *indexed hierarchy of partitions* (\mathcal{H}, λ) , with \mathcal{H} a *hierarchy of partitions* i.e. a chain of nested partitions $\mathcal{H} = \{\pi_0, \pi_1, \dots, \pi_n | \forall j, k, 0 \leq j \leq k \leq n \Rightarrow \pi_j \sqsubseteq \pi_k\}$, with π_n the single-region partition and π_0 the finest partition on the image, and $\lambda : \mathcal{H} \rightarrow \mathbb{R}^+$ being a stratification index verifying $\lambda(\pi) < \lambda(\pi')$ for two nested partitions $\pi \subset \pi'$. This increasing map allows us to value each contour according to the level of the hierarchy for which it disappears: this is the *saliency* of the contour, and we consider that the higher the saliency, the stronger the contour. For a given hierarchy, the image in which each contour takes as value its saliency is called

Ultrametric Contour Map (UCM)[3]. Representing a hierarchy by its UCM is an easy way to get an idea of its effect because thresholding an UCM always provides a set of closed curves and so a partition. In this paper, for better visibility, we represent UCM with inverted contrast.

To get a partition for a given hierarchy, there are several possibilities:

- simply thresholding the highest saliency values,
- marking some nodes as important ones and then computing a partition accordingly, which is known as *marker-based segmentation*,
- smartly editing the graph by finding the partition that minimizes an energetic function.

In a complementary approach, we argue that the quality of the obtained partitions highly depends on the hierarchy that we use, and thus that changing the dissimilarity can lead to more suitable partitions. Indeed, if the dissimilarity reflects only a local contrast as in the hierarchy issued by the RAG, the most salient regions in the image are the small contrasted ones. So instead of departing from a simple and rough dissimilarity such as contrast and then use an sophisticated technique to get a good partition out of it, one can also try to obtain a more informative dissimilarity adapted to the content of the image such that the simplest methods are sufficient to compute interesting partitions. This way, the aforementioned techniques lead to segmentations better suited for further exploitation. How can we construct more pertinent and informative dissimilarities?

2.2 Stochastic watershed hierarchies

The stochastic watershed (SWS), introduced in [2] on a simulation basis and extended with a graph-based approach in [17], is a versatile tool to construct hierarchies. The seminal idea is to operate multiple times marker-based segmentation with random markers and valuate each edge of the \mathcal{MST} by its frequency of appearance in the resulting segmentations.

Indeed, by spreading markers on the RAG \mathcal{G} , one can construct a segmentation as a MSF $\mathcal{F}(\mathcal{G})$ in which each tree takes root in a marked node. Marker-based segmentation directly on the \mathcal{MST} is possible: one must then cut, for each pair of markers, the highest edge on the path linking them. Furthermore, there is a domain of variation in which each marker can move while still leading to the same final segmentation. More details are provided in Figure 1.

Let us then consider on the \mathcal{MST} an edge e_{st} of weight ω_{st} and compute its probability to be cut. We cut all edges of the \mathcal{MST} with a weight superior or equal to ω_{st} , producing two trees T_s and T_t of roots s and t . If at least one marker falls within the domain R_s of T_s nodes and at least one marker falls within the domain R_t of T_t nodes, then e_{st} will be cut in the final segmentation.

Let denote $\mu(R)$ the number of random markers falling in a region R . We want to attribute to e_{st} the following probability value:

$$\begin{aligned} \mathbb{P}[(\mu(R_s) \geq 1) \wedge (\mu(R_t) \geq 1)] &= 1 - \mathbb{P}[(\mu(R_s) = 0) \vee (\mu(R_t) = 0)] \\ &= 1 - \mathbb{P}(\mu(R_s) = 0) - \mathbb{P}(\mu(R_t) = 0) + \mathbb{P}(\mu(R_s \cup R_t) = 0) \end{aligned} \tag{1}$$

If markers are spread following a Poisson distribution, then for a region R :

$$\mathbb{P}(\mu(R) = 0) = \exp^{-\Lambda(R)}, \quad (2)$$

With $\Lambda(R)$ being the expected value (mean value) of the number of markers falling in R . The probability thus becomes:

$$\mathbb{P}(\mu(\mathbf{R}_s) \geq 1 \wedge \mu(\mathbf{R}_t) \geq 1) = 1 - \exp^{-\Lambda(\mathbf{R}_s)} - \exp^{-\Lambda(\mathbf{R}_t)} + \exp^{-\Lambda(\mathbf{R}_s \cup \mathbf{R}_t)} \quad (3)$$

When the Poisson distribution has an homogeneous density λ :

$$\Lambda(R) = \mathbf{area}(R)\lambda, \quad (4)$$

When the Poisson distribution has a non-uniform density λ :

$$\Lambda(R) = \int_{(x,y) \in R} \lambda(x,y) \, dx dy \quad (5)$$

The output of the SWS algorithm thus depends on the departure \mathcal{MST} (structure and edges valuations) and of the probabilistic law governing the markers distribution. Furthermore, SWS hierarchies can be chained, leading to a wide exploratory space that can be used in a segmentation workflow [8].

Because of its versatility and good performance, SWS represents a good departure algorithm to modify in order to inject prior information. Indeed, when having a prior information about the image, is it possible to use it in order to have more details in some parts rather than others?

3 Hierarchies highlighting structures of interest using prior information

3.1 Hierarchy with Regionalized Fineness (HRF)

In the original SWS, a uniform distribution of markers is used (whatever size or form they may have). In order to have stronger contours in a specific region of the image, we adapt the model so that more markers are spread in this region.

Let E be an object or class of interest, for example $E = \text{“face of a person”}$, and \mathbf{I} be the studied image. We denote by θ_E the probability density function (PDF) associated with E obtained separately, and defined on the domain D of \mathbf{I} , and by $\text{PM}(\mathbf{I}, \theta_E)$ the probabilistic map associated, in which each pixel $p(x, y)$ of \mathbf{I} takes as value $\theta_E(x, y)$ its probability to be part of E . Given such an information on the position of an event in an image, we obtain a hierarchical segmentation focused on this region by modulating the distribution of markers.

If λ is a density defined on D to distribute markers (uniform or not), we set $\theta_E \lambda$ as a new density, thus favoring the emergence of contours within the regions of interest.

Considering a region R of the image, the mean number of markers falling within R is then:

$$\Lambda_E(R) = \int_{(x,y) \in R} \theta_E(x,y) \lambda(x,y) \, dx dy \quad (6)$$

Note that if we want N markers to fall in average within the domain D , we work with a slightly modified density:

$$\hat{\lambda} = \frac{N}{\mu(D)}\lambda \quad (7)$$

Furthermore, this approach can be easily extended to the case where we want to take advantage of information from multiple sources. Indeed, if θ_{E_1} and θ_{E_2} are the PDF associated with two events E_1 and E_2 , we can combine those two sources by using as a new density $(\theta_{E_1} + \theta_{E_2})\lambda$.

3.2 Methodology

We present here the steps to compute a HRF for an event E given a probabilistic map $\text{PM}(\mathbf{I}, \theta_E)$ providing spatial prior information on an image \mathbf{I} :

- compute a fine partition π_0 of the image, define a dissimilarity measure between adjacent regions and compute the RAG \mathcal{G} , and then the $\text{MST}(\mathcal{G})$ to easily work with graphs,
- compute a probabilistic map $\pi_\mu = \pi_\mu(\pi_0, \text{PM}(\mathbf{I}, \theta_E))$ with each region of the fine partition π_0 taking as a new value the mean value of $\text{PM}(\mathbf{I}, \theta_E)$ in this region,
- compute new values of edges by a bottom-up approach as described in section 3.1, where for each region \mathbf{R}_i of π_0 , $A(\mathbf{R}_i)$ corresponds to the mean value taken by pixels of the region \mathbf{R}_i in π_μ . Note that this approach allows a highly efficient implementation using dynamic programming on graphs.

3.3 Modulating the HRF depending on the couple of regions considered

If we want to favor certain contours to the detriment of others, we can modulate the density of markers in each region by taking into account the strength of the contour separating them but also the relative position of both regions.

We use the same example and notations as in section 3.1, and thus want to modulate the distribution of markers relatively to \mathbf{R}_s , \mathbf{R}_t and their frontier. For example, to stress the strength of the gradient separating both regions we can locally spread markers following the distribution $\chi(\mathbf{R}_s, \mathbf{R}_t)\lambda$, with $\chi(\mathbf{R}_s, \mathbf{R}_t) = \omega_{st}$. This corresponds to the classical volume-based SWS, which allows to obtain a hierarchy that takes into account both surfaces of regions and contrast between them.

To go further, one can use any prior information in a similar way. Indeed, while using prior information to influence the output of the segmentation workflow, one might also want to choose whether the relevant information to emphasize in resulting segmentations is the foreground, the background or the transitions between them.

For example, having more details in the transition regions between background and foreground allows us to have more precision where the limit between

foreground and background is actually unclear. As a matter of fact, the prior information often only provides rough positions of the foreground object with blurry contours, and such a process would allow to get precise contours of this object from the image.

Let us consider this case and define for each couple of regions (R_s, R_t) a suitable $\chi(R_s, R_t)$. We then want $\chi(R_s, R_t)$ to be low if R_s and R_t both are in the background or the foreground, and high if R_s is the background and R_t in the foreground (or the opposite). We use:

$$\begin{cases} \tilde{\lambda} = \chi\lambda \\ \chi(R_s, R_t) = \frac{\max(m(R_s), m(R_t))(1 - \min(m(R_s), m(R_t)))}{0.01 + \sigma(R_s)\sigma(R_t)}, \end{cases} \quad (8)$$

$m(R)$ (resp. $\sigma(R)$) being the normalized mean (resp. normalized standard deviation) of pixels values in the region R of $PM(\mathbf{I})$. Thus the number of markers spread will be higher when the contrast between adjacent regions is high (numerator term) and when these regions are coherent (denominator term).

Then for each edge, its new probability to be cut is :

$$\begin{aligned} \mathbb{P}[(\mu(R_p) \geq 1) \wedge (\mu(R_q) \geq 1)] &= 1 - \exp^{-\chi(R_s, R_t)A(R_p)} - \exp^{-\chi(R_s, R_t)A(R_q)} \\ &+ \exp^{-\chi(R_s, R_t)A(R_p \cup R_q)} \end{aligned} \quad (9)$$

In the spirit of [6], this mechanism provides us with a way to “realign” the hierarchy with respect to the relevant prior information to get more details where the information is blurry. Similar adaptations can be thought of to emphasize details of background or foreground regions.

In the following, we illustrate the methodology exposed with some applications.




Fig. 2. Hierarchical segmentation of faces. (a) Original image (b) Prior : Probabilistic map obtained thanks to face detection algorithm and (c) Volume-based SWS Hierarchy UCM (d) Volume-based HRF UCM with face position as prior (e)(f)(g) Examples of segmentations obtained with HRF - 10,100,1000 regions.

4 Application examples

4.1 Scalable transmission favoring regions of interest

Let us consider a situation where one emitter wants to transmit an image through a channel with a limited bandwidth, e.g. for a videoconference call. In such a case, the more important informations to transmit are details on the face of the person on the image. Besides, we nowadays have highly efficient face detectors, using for example Haar-wavelets as features in a learning-based vision approach [30]. Considering that for an image in entry, the face can be easily detected, we can use this information to produce a hierarchical segmentation of the image that accentuates the details around the face while giving a good sketch of the image elsewhere. Depending on the bandwidth available, we can then choose the level of the hierarchy to select and obtain the associated partition to transmit, ensuring us to convey the face with as much details as possible. Some results are presented in Figure 2, with notably a comparison between a classical volume-based SWS UCM and a volume-based HRF UCM.

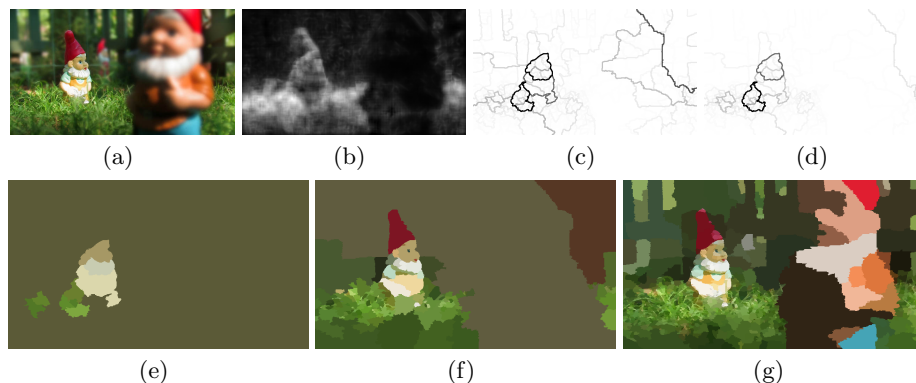


Fig. 3. Hierarchical segmentation of non-blur objects. (a) Original image (b) Prior image obtained with non-blur zones detection algorithm (c) Volume-based SWS Hierarchy UCM (d) Volume-based HRF UCM with non-blur zones as prior (e)(f)(g) Examples of segmentations obtained with HRF - 10,200,2000 regions.

4.2 Artistic aspect: focus and cartoon effect

The same method can also be used for artistic purposes. For example, when taking as prior the result of a blur detector [28], we can accentuate the focus effect wanted by the photograph and turn it into a cartoon effect as well - see Figure 3 for an illustration of the results.

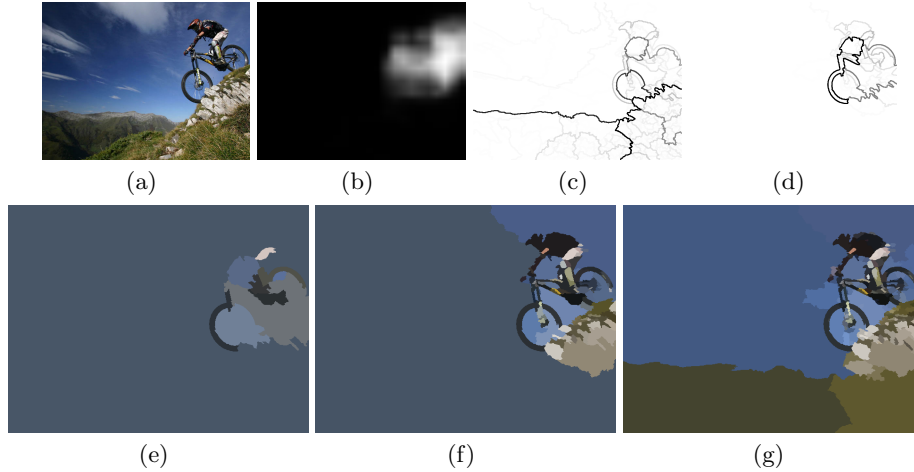


Fig. 4. Hierarchical segmentation of the main class (class "bike") in an image with heatmap issued of a CNN-based method as input. (a) Original image, (b) Heatmap issued by the CNN-based localization method, (c) Volume-based Watershed Hierarchy UCM and (d) Hierarchy with prior UCM. (e)(f)(g) Examples of segmentations obtained with HRF - 10,100,200 regions.

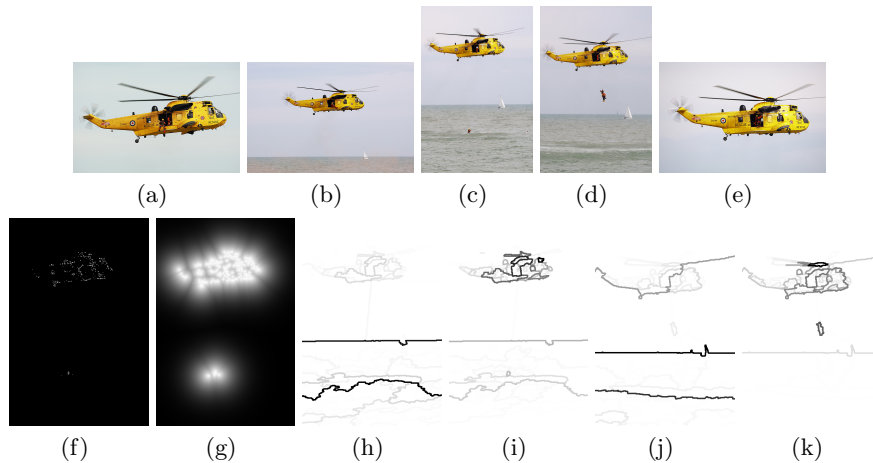


Fig. 5. Hierarchical co-segmentation of matched objects. (a)-(e) Images I_1 to I_5 (f) All matched key-points of I_3 with other images key-points (g) Prior for I_3 : probability map generated thanks to morphological distance function (h)(i) Volume-based SWS Hierarchy UCM for I_3 and I_4 (j)(k) Volume-based HRF UCM for I_3 and I_4 with matched key-points as prior.

In the same spirit, various methods now exist to automatically roughly localize the principal object in an image. We inspire ourselves from [20] to do so. Using the state-of-the-art convolution neural network (CNN) classifier VGG19 [27] trained on the 1000 classes ImageNet database [7], we first determine what is the main class in the image. Note that this CNN takes as input only images of size 224×224 pixels. Once it is known, we can then, by rescaling the image by a factor $s \in \{0.5, 0.7, 1.0, 1.4, 2.0, 2.8\}$, compute for sub-windows of size 224×224

of the image the probability of appearance of the main class. By simply superimposing the results for all sub-windows, we thus obtain a probabilistic map of the main class for each rescaling factor. By max-pooling, we keep in memory the result of the scale for which the probability is the highest. The *heatmap* thus produced can then be used to feed our algorithm. This way, we have at our disposal an automatized way to focus on the principal class in the scene. Some results are presented in Figure 4.




Fig. 6. Hierarchical segmentation of faces. (a) Original image (b)(c)(d) UCM of HRF depending of the couples of regions (section 3.3), of the regions only (section 3.1), and both combined. The rest of the images are segmentations examples with 4,10,25 regions, the hierarchies being presented in the same order.

4.3 Hierarchical co-segmentation

Another potential application is to co-segment with the same fineness level an object appearing in several different images. For example, when given a list of images of the same object taken from different perspectives/for different conditions, we can follow the state-of-the-art matching procedure [14]: (i) compute

all key-points in both images, (ii) compute local descriptors at these key-points, (iii) match those key-points using a spatial coherency algorithm as RANSAC. Once it is done we retain these matched key-points for both images, and generate probability maps of the appearance of the matched objects using a morphological distance function to the matched key-points.

These probability maps can then feed our algorithm, given as result a hierarchical co-segmentation that emphasizes the matched zones of the image. Some results are presented in Figure 5.

4.4 Example of the effect of the HRF highlighting transitions between foreground and background

We illustrate here the HRF highlighting transitions between foreground and background presented in section 3.3, by presenting its effect in the face detection example presented in figure 6.

5 Conclusions and perspectives

In this paper we have proposed a novel and efficient hierarchical segmentation algorithm that emphasizes the regions of interest in the image by using spatial exogenous information on it. The wide variety of sources for this exogenous information makes our method extremely versatile and its potential applications numerous, as shown by the examples developed in the last section. To go further, we could find a way to efficiently extend this work to videos. One could also imagine a semantic segmentation method that would go back and forth between localization algorithm and HRF to progressively refine the contours of the main objects in the image.

References

1. Achanta, R., Shanji, A., Smith, K., Lucchi, A., Fua, P., Süsstrunk, S.: SLIC super pixels compared to state-of-the-art super pixel methods. *IEEE PAMI* (2012) 2
2. Angulo, J., Jeulin, D.: Stochastic watershed segmentation. *ISMM* 1, 265–276 (2007) 4
3. Arbelaez, P., Maire, M., Fowlkes, C., Malik, J.: Contour detection and hierarchical image segmentation. *IEEE PAMI* 33(5), 898–916 (2011) 1, 4
4. Chan, T., Zhu, W.: Level set based shape prior segmentation. In: *CVPR*. vol. 2, pp. 1164–1170 (2005) 2
5. Chen, F., Yu, H., Hu, R., Zeng, X.: Deep learning shape priors for object segmentation. In: *CVPR*. pp. 1870–1877 (2013) 2
6. Chen, Y., Dai, D., Pont-Tuset, J., Van Gool, L.: Scale-aware alignment of hierarchical image segmentation. In: *CVPR*. pp. 0–0 (2016) 7
7. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: A Large-Scale Hierarchical Image Database. In: *CVPR* (2009) 9
8. Fehri, A., Velasco, S., Meyer, F.: Automatic selection of stochastic watershed hierarchies. In: *EUSIPCO* (August 2016) 5

9. Gueguen, L., Velasco-Forero, S., Soille, P.: Local mutual information for dissimilarity-based image segmentation. *JMIV* pp. 1–20 (2013) 2
10. Guigues, L., Cocquerez, J.P., Le Men, H.: Scale-sets image analysis. *IJCV* 68(3), 289–317 (2006) 1
11. Kiran, B.R., Serra, J.: Ground truth energies for hierarchies of segmentations. In: *ISMM*, pp. 123–134 (2013) 2
12. Kiran, B.R., Serra, J.: Global-local optimizations by hierarchical cuts and climbing energies. *Pattern Recognition* 47(1), 12–24 (2014) 2
13. Lampert, C., Blaschko, M., Hofmann, T.: Beyond sliding windows: Object localization by efficient subwindow search. In: *CVPR*. pp. 1–8 (2008) 2
14. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *IJCV* 60(2), 91–110 (2004) 10
15. Machairas, V., Faessel, M., Cárdenas-Peña, D., Chabardes, T., Walter, T., Decencière, E.: Waterpixels. *IEEE TIP* 24(11), 3707–3716 (2015) 2
16. Malmberg, F., Hendriks, C.L.L., Strand, R.: Exact evaluation of targeted stochastic watershed cuts. *Discrete Applied Mathematics* 216, 449–460 (2017) 2
17. Meyer, F.: Stochastic watershed hierarchies. In: *ICAPR*. pp. 1–8 (2015) 1, 4
18. Meyer, F., Beucher, S.: Morphological segmentation. *Journal of visual communication and image representation* 1(1), 21–46 (1990) 2
19. Najman, L., Cousty, J., Perret, B.: Playing with kruskal: algorithms for morphological trees in edge-weighted graphs. In: *Mathematical Morphology and Its Applications to Signal and Image Processing*, pp. 135–146. Springer (2013) 3
20. Oquab, M., Bottou, L., Laptev, I., Sivic, J.: Is object localization for free? weakly-supervised learning with convolutional neural networks. In: *CVPR* (2015) 2, 9
21. Ouzounis, G.K., Pesaresi, M., Soille, P.: Differential area profiles: decomposition properties and efficient computation. *IEEE PAMI* 34(8), 1533–1548 (2012) 2
22. Pont-Tuset, J., Arbelaez, P., Barron, J., Marques, F., Malik, J.: Multiscale combinatorial grouping for image segmentation and object proposal generation. *PAMI* (2016) 1
23. Ren, Z., Shakhnarovich, G.: Image segmentation by cascaded region agglomeration. In: *CVPR*. pp. 2011–2018 (2013) 1
24. Ronse, C.: Ordering partial partitions for image segmentation and filtering: Merging, creating and inflating blocks. *JMIV* pp. 1–32 (2013) 2
25. Sermanet, P., Eigen, D., Zhang, X., Mathieu, M., Fergus, R., LeCun, Y.: Overfeat: Integrated recognition, localization and detection using convolutional networks. In: *ICLR* (2013) 2
26. Serra, J.: Tutorial on connective morphology. *Selected Topics in Signal Processing, IEEE Journal of* 6(7), 739–752 (Nov 2012) 2
27. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *arXiv 1409.1556* (2014) 9
28. Su, B., Lu, S., Tan, C.L.: Blurred image region detection and classification. In: *ACM International Conference on Multimedia*. pp. 1397–1400. *MM '11* (2011) 8
29. Veksler, O.: Star shape prior for graph-cut image segmentation. In: *ECCV*, pp. 454–467. Springer (2008) 2
30. Viola, P., Jones, M.: Rapid object detection using a boosted cascade of simple features. In: *CVPR*. pp. 511–518 (2001) 8
31. Xu, Y., Géraud, T., Najman, L.: Hierarchical image simplification and segmentation based on mumford-shah-salient level line selection. *Pattern Recognition Letters* (2016) 2