



A Metric Approach to Vector-Valued Image Segmentation

PABLO A. ARBELÁEZ AND LAURENT D. COHEN

CEREMADE, UMR CNRS 7534 Université Paris Dauphine, Place du maréchal de Lattre de Tassigny, 75775
Paris cedex 16, France

arbelaez@ceremade.dauphine.fr

cohen@ceremade.dauphine.fr

Received April 15, 2004; Revised November 12, 2004; Accepted May 5, 2005

First online version published in April, 2006

Abstract. We address the issue of low-level segmentation of vector-valued images, focusing on the case of color natural images. The proposed approach relies on the formulation of the problem in the metric framework, as a Voronoi tessellation of the image domain. In this context, a segmentation is determined by a distance transform and a set of sites. Our method consists in dividing the segmentation task in two successive sub-tasks: pre-segmentation and hierarchical representation. We design specific distances for both sub-problems by considering low-level image attributes and, particularly, color and lightness information. Then, the interpretation of the metric formalism in terms of boundaries allows the definition of a soft contour map that has the property of producing a set of closed curves for any threshold. Finally, we evaluate the quality of our results with respect to ground-truth segmentation data.

Keywords: image segmentation, distance transforms, path variation, ultrametrics, vector-valued image, color, boundary detection

1. Introduction

The metric framework for spatial tessellations was first formalized by Dirichlet (1850) and Voronoi (1907), who studied the idea of partitioning the space by considering a finite set of fixed points, called *sites*, and assigning each point to the closest site. The regions defined by this construction are usually called *Voronoi regions* or *Voronoi cells* and the resulting spatial decomposition is known as a *Voronoi tessellation* or *Voronoi diagram*. Figure 1 presents an example of this structure in its original formulation: a rectangle in the plane is partitioned by measuring the Euclidian distance between each point and the four sites shown on the left. In this case, the Voronoi regions are convex polygons.

Since its early introduction, the Voronoi tessellation has found application in a wide range of disciplines and has inspired several generalizations (Aurenhammer and Klein, 2000; Okabe et al., 2002). In the context of image analysis, application of this structure includes compression (Ahuja et al., 1985), texture classification (Tuceryan and Jain, 1990) and shape representation (Mayya and Rajan, 1996).

In this paper, we consider an extension of the Voronoi tessellation to pseudo-metric spaces and we study its application to the segmentation of vector-valued images, focusing on color images. Within this framework, segments are defined as Voronoi regions of a set of sites, relatively to a pseudo-metric. The problem is thus transferred to the definition of a relevant distance transform from the image data and the




Figure 1. Set of sites and Euclidian Voronoi tessellation.

selection of a set of sites. Our approach consists in applying the metric formalism to two successive sub-tasks of segmentation : pre-segmentation and hierarchical representation of color images. For this purpose, we design specific distances by considering two low-level image attributes, namely color and brightness information.

We first present a pre-segmentation technique called the *extrema mosaic*. The pseudo-metric defined in this part, the *path variation*, is a generalization of the one dimensional total variation for vector-valued functions of multiple variables. In the case of color images, we study the Voronoi tessellation obtained by considering the extrema of the lightness channel of the image as sites. This method provides a natural reconstruction of the image that offers a balance between content conservation and simplification.

In the second part, we focus only on those tessellations that are invariant when a site is displaced inside its Voronoi region. This consideration leads to a type of pseudo-metrics called ultrametrics. These distances are useful for a multiscale representation of the image, as their definition amounts to constructing a stratified hierarchy of partitions of the image domain. Starting from the pre-segmented image, we define an ultrametric for color image segmentation. This distance incorporates boundary and inner region information.

Finally, we exploit the boundary-based formulation of the metric formalism to define a soft boundary image, which we call *ultrametric contour map* (UCM). Our region-based approach guarantees that a threshold in this image supplies a set of closed curves. The UCM are then used to evaluate the quality of our results with respect to ground-truth human segmentations.

This paper is organized as follows. Section 2 presents the mathematical framework. Section 3 introduces our pre-segmentation method. Section 4 is dedicated to hierarchical segmentation and ultrametric contour maps. Finally, the evaluation of our results is discussed in Section 5.

2. Voronoi Tessellations

In this section, the classic notion of Voronoi tessellation is extended to the framework of pseudo-metric spaces.

Definition 1. A pseudo-metric space (Kuratowski, 1966; Kelley, 1975) is a pair (Ω, ψ) where Ω is a set and the application $\psi : \Omega^2 \rightarrow \mathbb{R}^+$ satisfies the following axioms:

1. $\psi(x, x) = 0, \forall x \in \Omega$.
2. $\psi(x, y) = \psi(y, x), \forall x, y \in \Omega$.
3. $\psi(x, y) \leq \psi(z, x) + \psi(z, y), \forall x, y, z \in \Omega$.

The number $\psi(x, y)$ is called the *distance* between points x and y .

A pseudo-metric space is *convex* if each pair of points $x, y \in \Omega$ can be joined by a ψ -straight path, i.e., a continuous application $\gamma : [a, b] \rightarrow \Omega$ such that: $\gamma(a) = x, \gamma(b) = y$ and

$$\forall t \in [a, b], \psi(x, y) = \psi(x, \gamma(t)) + \psi(\gamma(t), y).$$

Definition 2. Let (Ω, ψ) be a closed convex pseudo-metric space and $S = \{s_1, \dots, s_n\} \subseteq \Omega$ a set of fixed points called *sites*.

The *Voronoi region*, or *V-region*, of a site $s_i \in S$ is defined as:

$$R_{s_i} = \{x \in \Omega \mid \psi(x, s_i) \leq \psi(x, s_j), \forall j \in \{1, \dots, n\}, j \neq i\}.$$

The *Voronoi tessellation*, or *V-tessellation*, of Ω associated with ψ and S is the set of Voronoi regions:

$$\Pi(\psi, S) = \{R_{s_1}, \dots, R_{s_n}\}.$$

Note that Axiom 1 of Definition 1 allows different points to be at zero distance in a pseudo-metric space. By considering the equivalence classes $\hat{x}(\psi) = \{y \in \Omega \mid \psi(x, y) = 0\}$, one can define a *quotient space* $\hat{\Omega}(\psi)$, which is a metric space. Each element of a V-tessellation is therefore a union of elements in $\hat{\Omega}(\psi)$. The equivalence classes indicate the level of resolution of the pseudo-metric, under which the distance is blind.

Two main differences between our approach and the standard framework of Voronoi tessellations (Aurenhammer and Klein, 2000; Okabe et al., 2002) are worth noting. First, by considering pseudo-metrics

we have access to a class of spaces larger than the metric spaces. Second, since we aim at applying this structure to image analysis, the set Ω corresponds to the domain of definition of a color image and the pseudo-metrics we study depend explicitly on the image data.

Hence, in this context, the segmentation of a color image is determined by the definition of a relevant pseudo-metric and the selection of a set of sites; the rest of the paper is dedicated to these issues. However, as we make no assumption on the image content, we only consider as set of possible contours the discontinuities of the original image. This property can be obtained by defining the pseudo-metric ψ such that the equivalence class of a point $x \in \Omega$ coincides with the connected component of u that contains x . The quotient space $\hat{\Omega}(\psi)$ is then homeomorphic to the space of components of the image. Moreover, for a set of sites S , each element of $\Pi(\psi, S)$ is a union of components of u . Such V-tessellations simplify then the image while preserving its original contour information.

Finally, in order to address the segmentation of color images in the metric framework, a practical problem is the definition of a distance between colors. For this issue, the C.I.E. standard L^*ab (Wyszecki and Stiles, 1982) is adopted in this paper. This color representation ought to approximate a perceptually uniform color space. Though not perfect, it provides two main advantages with respect to the basic RGB system: first, the separation of the color information into a lightness channel L^* and two chromatic channels a^* and b^* ; second, the approximation of the metric in the Riemannian color space by the Euclidian distance. In the sequel, the distance in the color space L^*ab is noted δ^* .

3. Pre-segmentation

This section presents an application of the metric formalism to the pre-segmentation of color images. For this purpose, we begin by defining a specific distance.

Let (X, ψ) be a pseudo-metric space. Consider a finite partition of an interval $[a, b]$, $\sigma = \{t_0, \dots, t_n\}$, such that $a = t_0 < t_1 < \dots < t_n = b$ and note Φ the set of such partitions.

The *variation* at order p of a function $f : [a, b] \rightarrow (X, \psi)$, is given by:

$$v^p(f) = \sup_{\sigma \in \Phi} \sum_{i=1}^n \psi(f(t_i), f(t_{i-1}))^p.$$

If (X, ψ) is the set of real numbers with the usual metric, then $v^1(f)$ corresponds to the *total variation* of f , the

well known functional introduced by Jordan. In our case, (X, ψ) is the color space (L^*ab, δ^*) .

For multiple variable functions, we propose to measure the minimal variation on all the paths between two points:

The *path variation* at order p of a function $u : \Omega \rightarrow (X, \psi)$, is defined as:

$$V_u^p(x, y) = \inf_{\gamma \in \Gamma_{xy}} v^p(u \circ \gamma), \quad \forall x, y \in \Omega. \quad (1)$$

where γ is a continuous path between x and y and Γ_{xy} the set of such paths.

Note that, in contrast to the usual notion of *total variation* for functions of multiple variables (Rudin et al., 1992), the path variation is defined for each couple of points of Ω . By definition, the equivalence class $\hat{x}(V_u^p)$ coincides with the connected component of u containing x . A discrete definition of this distance for color images can be found in Arbeláez and Cohen (2003a).

The path variation is an interesting pseudo-metric for a local level of analysis, as it quantifies the minimal variation of color between pixels. In order to apply this distance to pre-segmentation, we now turn to the selection of an appropriate set of sites.

The purpose of pre-segmentation is to decompose the image into local entities that preserve its geometric structure. Therefore, the spatial distribution of sites should be physically representative of the image content and each significant feature should contain at least one site. In the case of natural images, the lightness extrema satisfy these two properties, as shown in Fig. 2.

Thus, we consider the Voronoi tessellation $\Pi(V_u^p, ext(u))$, where $ext(u)$ denotes the set of extremal components of the lightness channel L^* of a color image u . An *extrema mosaic* of u is the piecewise constant reconstruction of the image obtained by assigning a color to each V-region of $\Pi(V_u^p, ext(u))$.

Figure 2 illustrates the method. The central column presents the sites in black and the extrema mosaics are shown on the right, with the V-regions depicted on their mean color. The density of the lightness extrema is high on focused or textured regions and low in blurred or homogeneous zones. Consequently, the blur is reduced in the reconstructed images, while the geometric structure, including notably contour information, is preserved.

Hence, the choice of the path variation as the pseudo-metric and the spatial distribution of the sites determine a V-tessellation where a balance between simplification and content conservation is obtained. In the sequel, the




Figure 2. From left to right: Original images, sites and extrema mosaics (images in color in the electronic version).

extrema mosaic is used as a pre-segmentation method. This approach may also be combined with a non-linear diffusion filtering in order to reduce the number of extrema and regularize the *tessellations* (Arbeláez and Cohen, 2003b).

4. Hierarchical Segmentation

In the framework of Voronoi tessellations, segments are modeled as Voronoi regions. It is then desirable that displacing the site inside a V-region does not modify its boundary. Unfortunately, the path variation does not satisfy this property. The rest of this paper is dedicated to the study of a particular type of pseudo-metrics exhibiting this invariance requirement.

4.1. Stratified Hierarchies and Ultrametrics

Ultrametrics are a standard tool in data analysis (Benzécri, 1984). They are often used for clustering because they determine a particular type of strong hierarchies. These distances are thus naturally suited for a multiscale representation of the image.

A *stratified hierarchy* \mathcal{H} is the set of regions of a family of nested partitions, together with a function $st : \mathcal{H} \rightarrow \mathbb{R}$, called a stratification index, such that: $\forall a, b \in \mathcal{H} : a \subset b \Rightarrow st(a) < st(b)$.

An *ultrametric space* (Ω, ψ) is a pseudo-metric space for which Axiom 3 of Definition 1 is replaced by the stronger relation: $\psi(x, y) \leq \max\{\psi(z, x), \psi(z, y)\}$.

The geometry of an ultrametric space differs significantly from the usual Euclidian case. In particular, two ultrametric balls can only be disjoint or nested.

Consequently, the set of all the closed balls for a fixed radius r determines a special type of Voronoi tessellation, noted $\Pi(\psi, r)$. The set \mathcal{H} of all the ultrametric balls of positive radius is then a stratified hierarchy of Ω . A stratification index for \mathcal{H} is given by the function that assigns to each ultrametric ball its radius. Conversely, each stratified hierarchy defines an ultrametric distance.

If a Voronoi tessellation is determined by an ultrametric, the previous properties imply that replacing a site $s_i \in S$ by another point s'_i in the interior of its Voronoi region R_{s_i} does not modify the V-tessellation. These pseudo-metrics satisfy then the invariance requirement mentioned at the beginning of the section. Moreover, the problem of selecting a set of sites can be addressed in this case through the choice of a radius r for the ultrametric V-regions. In the sequel, the ultrametrics are normalized in order to assign the value of 1 to the radius of the smallest V-region that contains the whole domain.

4.2. An Ultrametric for Segmentation

In this subsection, we define a specific ultrametric for the segmentation of color images. Its construction is derived from the characterization of this type of distances as stratified hierarchies.

A family of nested partitions can be constructed by a graph-based region merging strategy. Such a clustering approach consists in progressively merging regions of an initial partition according to a dissimilarity measure, a real valued function defined for each pair of adjacent subsets of the domain (Garrido et al., 1998; Forsyth and Ponce, 2003).

However, in order to define a stratified hierarchy \mathcal{H} , the dissimilarity d must be compatible with the hierarchical order:

$$\begin{aligned} a \subset a' \wedge b \subset b' &\Rightarrow d(a, b) < d(a', b'), \\ \forall a, a', b, b' &\in \mathcal{H}. \end{aligned} \quad (2)$$

In our case, the ultrametric is constructed in two steps. First, we use the color information to quantify the contrast between V-regions. For this purpose, a *contrast dissimilarity*, noted d^c , is defined as:

$$d^c(R_1, R_2) = \frac{\sum \delta^*(u(p_1), u(p_2))}{\text{length}(\partial(R_1, R_2))}$$

where $\partial(R_1, R_2)$ denotes the common boundary of two connected regions R_1 and R_2 , and the sum is calculated on all the adjacent pixels of the initial partition such that $p_1 \in R_1$ and $p_2 \in R_2$. Thus, this dissimilarity measures the average color difference in the common boundary of the regions, on the extrema mosaic. Note that d^c satisfies Eq. (2).

In order to complement the boundary information supplied by d^c , we measure an internal *attribute*, a positive real valued function \mathcal{A} , in each V-region. The attribute is required to be increasing with the inclusion order. Thus, starting from d^c , a new dissimilarity d^α is defined by the formula:

$$d^\alpha(R_1, R_2) = d^c(R_1, R_2) \cdot \min\{\mathcal{A}(R_1), \mathcal{A}(R_2)\}^\alpha.$$

Since \mathcal{A} is increasing and d^c is compatible with the hierarchical order, so is d^α . The associated ultrametric, noted Υ^α , incorporates boundary as well as internal information of the V-regions. For the examples presented in this paper, the attribute is the size of the V-region. The parameter $\alpha \geq 0$ weights the balance between contrast and area. Note that, as for the path variation, the equivalence class of a point x , $\hat{x}(\Upsilon^\alpha)$, coincides with the connected component of the image that contains x . Thus, the quotient space $\widehat{\Omega}(\Upsilon^\alpha)$ is also homeomorphic to the space of components of the image.

Figure 3 illustrates the influence of α in the ultrametric. The central and right images show the V-tessellations associated to the same normalized radius $r = 0.35$, for the ultrametries Υ^0 and $\Upsilon^{0.2}$ respectively. Since the first distance is determined only by the contrast, small and contrasted regions, as the letters, are extracted. When $\alpha = 0.2$, these regions are eliminated

from the V-tessellation. Thus, the choice of α allows the ultrametric to adapt to the image content or a particular application.

4.3. Ultrametric Contour Maps

Many segmentation approaches provide only binary boundary maps, while soft edge maps generated by local edge detectors often require contour completion techniques in order to obtain closed curves. The formulation of the segmentation problem in the metric framework and the use of ultrametric distances provide a natural way to fill this gap. For this purpose, the following notion is central.

The *saliency* of a point x in an ultrametric space (Ω, Υ) is defined as the highest radius λ such that x belongs to a boundary of the V-tessellation $\Pi(\Upsilon, \lambda)$. The valuation of each point by its saliency determines a real-valued image, called in the sequel a *ultrametric contour map* (UCM). This single real-valued image is a compact representation of the ultrametric space: a threshold in the UCM supplies the set of boundaries of the corresponding V-tessellation. This idea was first used in Najman and Schmitt (1996) to valuate the watershed arcs of a gradient image.

The interest of this soft boundary map is that it combines the strong points of both region-based segmentation methods and local edge detectors. On one hand, the saliency of each point is given by the hierarchical structure and is therefore not restricted to the information of a local neighborhood. On the other hand, a threshold of an ultrametric contour map provides by definition a set of closed curves. Figure 4 shows an example of UCM, with dark intensities representing high saliency. For comparison, the right image presents the optimal threshold of a state of the art local edge detector (Martin et al., 2004).

5. Evaluation of Results

In order to measure the quality of our results, we used as ground-truth the Berkeley Segmentation Dataset and Benchmark (BSDB), a database of images representing natural scenes, manually segmented by humans (Martin et al., 2001). A methodology for evaluating the performance of boundary detectors with this database was developed in Martin et al. (2004). It is based on the comparison of detected edge points with respect to human-marked boundaries, using the Precision-Recall




Figure 3. From left to right : Original image, tessellations $\Pi(\Upsilon^0, 0.35)$ and $\Pi(\Upsilon^{0.2}, 0.35)$.




Figure 4. From left to right : Original image, UCM, thresholded UCM and thresholded local edge detector.

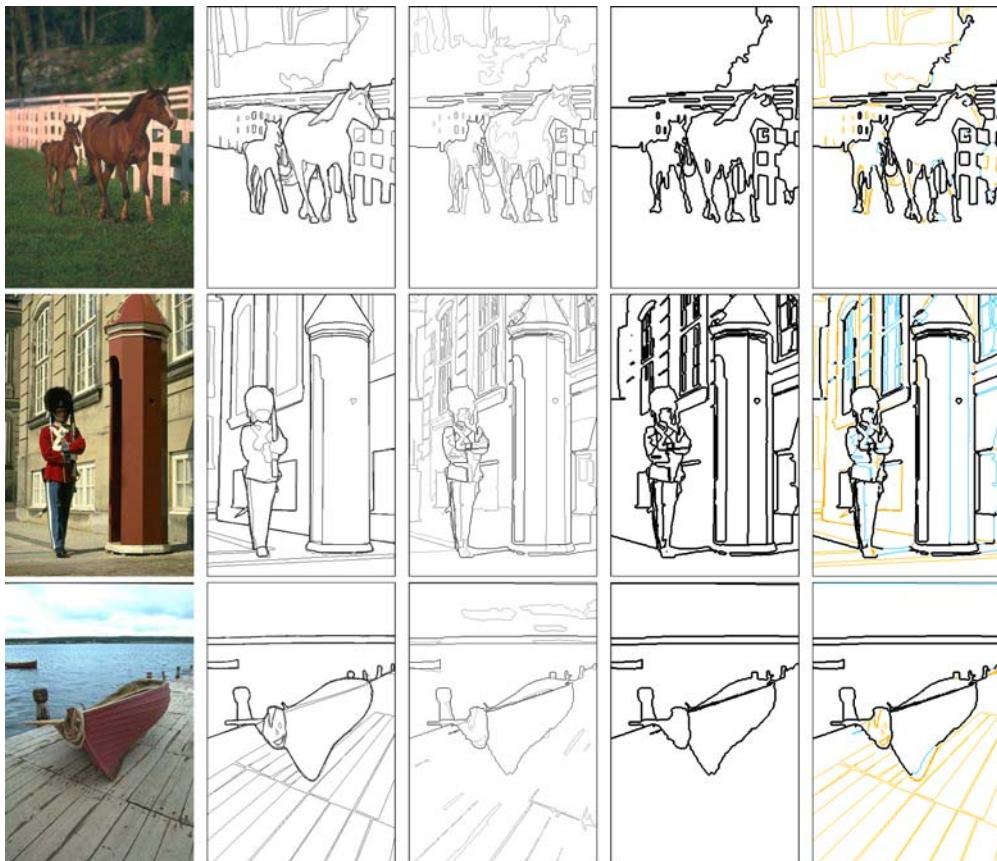


Figure 5. From left to right: original images, human segmentations, UCM, thresholded UCM and correspondence with ground-truth (see text).

framework (van Rijsbergen, 1979). Precisely, two quality descriptors are considered, *Recall* (R), given by the fraction of detections among true boundaries and *Precision* (P), defined as the fraction of true boundaries

among detections. Measuring these quantities for different thresholds of a detector on a given set of images provides a parametric *Precision-Recall curve*. The two quantities can be combined in a single descriptor,

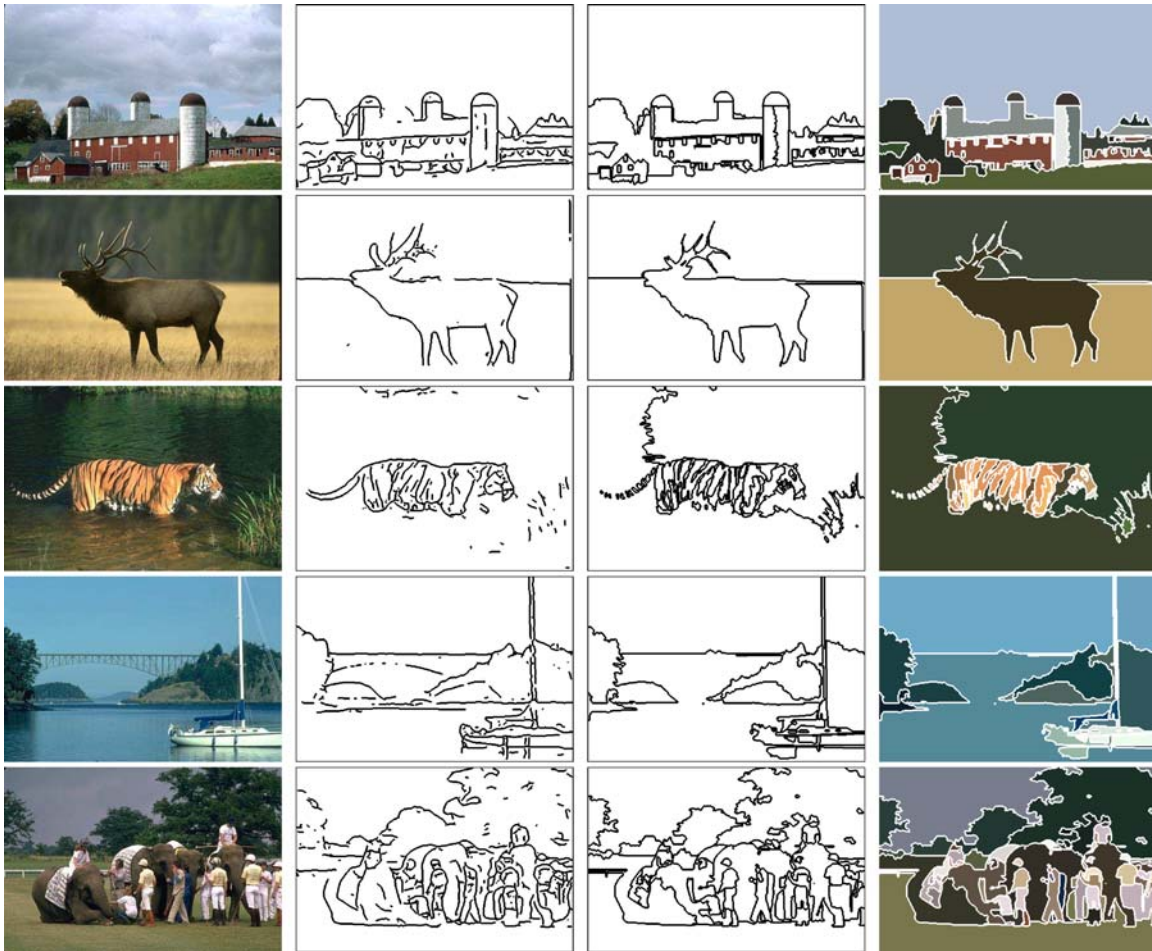


Figure 6. From left to right: Original images, thresholded local detector, thresholded UCM and associated reconstruction (with segments represented by their mean color).

the F -measure, defined as the harmonic mean of Precision and Recall: $F(P, R) = 2PR/(P + R)$. The maximal F -measure on the curve is then used as a summary statistic for the quality of the detector on the set of images.

Our system has a number of free parameters, intended to allow its adaptation to a particular application or type of images. In the pre-segmentation stage, these are the order of the path variation, the color model to represent the V-regions and the number of iterations of the extrema mosaic. The hierarchical representations defined in Section 4 depend also on the parameter α . Additionally, we consider a weighting factor between chrominance and lightness in the color space.

Using the ultrametric contour maps, all the free parameters were tuned with respect to the set of 200 train images of the BSDB. The optimal threshold of the UCM provided then the best radius λ for the ultrametric V-tessellations $\Pi(\Upsilon^\alpha, \lambda)$. The set of optimal train parameters was next used to evaluate our method on the independent test set of 100 images of the BSDB, leading to an overall score of $F(0.62, 0.64) = 0.63$.

Figure 5 illustrates the type of description of the image geometry provided by our ultrametric contour maps. The results were obtained using the set of optimal train parameters and the corresponding optimal threshold. It can be observed that the contours of the main structures are accurately extracted by the

UCM. The thresholded UCM show the best overall compromise between data fidelity and noise for our system. The right column presents the thresholded UCM superimposed on the human segmentations. True positives are depicted in thick black, false positives in thick blue and missed detections in yellow. Although the system was tuned by considering only boundary information, note that the segments obtained often coincide with physical objects or their connected parts.

The Precision-Recall framework permitted also the comparison of our method with the local boundary detector proposed by the authors of the BSDB (Martin et al., 2004) and noted in the sequel MFM. This approach relies on the measure at each pixel of discontinuities in feature channels like brightness, color and texture, over a range of scales and orientations. The cues are combined and optimized with respect to the BSDB. Using the evaluation methodology previously described, Martin et al. show that the MFM outperforms all classical edge detection techniques, its overall score being $F(0.62, 0.68) = 0.65$ on the test set.

Hence, for the optimal threshold, our system and the MFM present a similar amount of noise, while the latter detects more ground-truth. However, it should be noted that the two approaches are essentially different, as illustrated in Fig. 6. The edge points provided by a thresholded version of a local detector like the MFM do not necessarily form closed curves. In contrast, our region-based approach guarantees that any threshold on the UCM determines a set of closed curves, the boundaries of a segmentation. Hence, the lower overall score is compensated in our case by a higher representation level, the region instead of the point. Qualitatively, the MFM often performs better than our UCM in textured images, while our system is usually superior on images representing contrasted objects of relatively uniform color. Since both methods use only low-level image information, they show poor performances when semantic knowledge is determinant for human segmentation.

Finally, the overall F -measure for the human segmentations, when compared among them, is $F(0.90, 0.70) = 0.79$ on the test set. This score quantifies the human performance for the task and serves as the ultimate goal for machine segmentation. The gap between the two tested methods and human performance is mainly due to Precision (almost 0.3). Thus, efforts should be put on reducing the fraction of false positives in the segmentation results. In our case, future work

includes the definition of pseudo-metrics for which the texture information is also taken into consideration.

References

- Ahuja, N., An, B., and Schachter, B. 1985. Image representation using Voronoi tessellation. *CVGIP*, 29(3):286–295.
- Arbeláez, P.A. and Cohen, L.D. 2003a. Generalized Voronoi tessellations for vector-valued image segmentation. In *Proc. 2nd IEEE Workshop on Variational, Geometric and Level Set Methods in Computer Vision (VLSM'03)*, Nice, France, pp. 49–56.
- Arbeláez, P.A. and Cohen, L.D. 2003b. Path variation and image segmentation. In *Proc. EMMCVPR'03*, Lisbon, Portugal, pp. 246–260.
- Aurenhammer, F. and Klein, R. 2000. *Handbook of Computational Geometry*, Chapt. 5: Voronoi Diagrams, Elsevier Science Publishing, pp. 201–290.
- Benzécri, J.P. 1984. *L'Analyse des Données. Tome I: La Taxinomie*, 4th edition. Paris: Dunod.
- Dirichlet, P.G.L. 1850. Über die Reduction der positiven quadratischen Formen mit drei unbestimmten ganzen Zahlen. *J. Reine Angew. Mathematik*, 40:209–227.
- Forsyth, D.A. and Ponce, J. 2003. *Computer Vision: A Modern Approach*. Prentice Hall.
- Garrido, L., Salembier, P., and Garcia, D. 1998. Extensive operators in partition lattices for image sequence analysis. *IEEE Trans. on Signal Processing*, 66(2):157–180. Special Issue on Video Sequence Segmentation.
- Kelley, J.L. 1975. *General Topology*: Springer.
- Kuratowski, K. 1966. *Topology*, vol. I. Academic Press.
- Martin, D., Fowlkes, C., and Malik, J. 2004. Learning to detect natural image boundaries using local brightness, color and texture cues. *IEEE Trans. on PAMI*, 26(5):530–549.
- Martin, D., Fowlkes, C., Tal, D., and Malik, J. 2001. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proc. ICCV'01*, Vol. II, Vancouver, Canada, pp. 416–423.
- Mayya, N. and Rajan, V. 1996. Voronoi diagrams of polygons: A framework for shape representation. *Journal of Mathematical Imaging and Vision*, 6(4):355–378.
- Najman, L. and Schmitt, M. 1996. Geodesic saliency of watershed contours and hierarchical segmentation. *IEEE Trans. on PAMI*, 18(12):1163–1173.
- Okabe, A., Boots, B., Sugihara, K., and Chiu, S.N. 2002. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*. 2nd edition, Wiley.
- Rudin, L., Osher, S., and Fatemi, E. 1992. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268.
- Tuceryan, M. and Jain, A. 1990. Texture segmentation using Voronoi polygons. *IEEE Trans. on PAMI*, 12(2):211–216.
- van Rijsbergen, V. 1979. *Information Retrieval*. Dept. of Comp. Science, Univ. of Glasgow.
- Voronoi, G.M. 1907. Nouvelles applications des paramètres continus à la théorie des formes quadratiques. Premier Mémoire: Sur quelques propriétés des formes quadratiques positives parfaites. *Journal für die Reine und Angewandte Mathematik*, 133:97–178.
- Wyszecki, G. and Stiles, W.S. 1982. *Color Science: Concepts and Methods, Quantitative Data and Formulas*. J. Wiley and Sons.