Generalized Voronoi Tessellations for Vector-Valued Image Segmentation

Pablo Andrés 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, cohen*}*@ceremade.dauphine.fr*

Abstract

We address the issue of low-level segmentation for vector-valued images, focusing on color images. The proposed approach relies on the formulation of the problem as a generalized Voronoi tessellation of the image domain. In this context, the issue is transferred to the definition of an appropriated pseudo-metric and the selection of a set of sources. Two types of pseudo-metrics are considered; the first one is based on energy minimizing paths and the second is associated to the families of nested partitions of the image domain. We discuss specific applications of our approach to pre-segmentation, edge detection and hierarchical segmentation on color images.

1 Introduction

Spatial tessellations were first studied by Dirichlet [12] and Voronoi [32], who formalized the idea of partitioning the space by considering a set of *source points* and then assigning every point to the closest source. Since its early introduction, the Voronoi tessellation has found application in a wide range of disciplines [2, 29]. In this paper, we consider an extension of this notion to pseudo-metric spaces and we study its application to the segmentation of vector-valued images.

Image segmentation is a fundamental issue in computer vision. Its great complexity lies in the fact that structuring visual information into *meaningful regions* requires a layer of semantic understanding of the image content. However, a first task is the extraction of the image structure provided by the interaction between low-level cues. The present paper addresses the segmentation issue at this pre-cognitive stage of perception for vector-valued images, emphasizing the case of color images. For this purpose, the problem is formulated as a generalized Voronoi tessellation of the image domain.

In this framework, the segmentation issue is transferred to the definition of a relevant pseudo-metric from the image data and the selection of a set of sources. We consider two main types of pseudo-metrics and we study their application to address the segmentation issue at two different levels of analysis of the image.

The first application is a pre-segmentation method, which we call the *extrema mosaic*. The pseudo-metrics considered in this part are based on the study of energy minimizing paths and their use is often appropriated for a local level of analysis of the image. Specifically, we introduce a pseudo-metric, the *path variation*, that is a generalization of the one dimensional total variation for vector-valued functions of multiple variables. Its application provides a natural reconstruction of the image that offers a balance between content conservation and simplification.

In the second part, the type of pseudo-metrics considered are the ultrametrics. These pseudo-metrics are useful for a global level of analysis, since their definition amounts to constructing a stratified hierarchy of partitions of the image domain. Starting from the pre-segmented image, we define a ultrametric that expresses a notion of global contrast in the image. Then, using this measure of contrast, we propose a new model of edges for color images, which we call the *extrema edges*. Our approach guarantees that a threshold in this edge map supplies a set of closed contours where semantically important characteristics of edges are preserved. Finally, we use the contrast measure as the base for the definition of new ultrametrics, where the internal characteristics of the regions are also taken into consideration.

The rest of the paper is organized as follows. In Section 2, the mathematical framework is described. Section 3 presents the path metrics. Section 4 introduces our pre-segmentation method. In Section 5, we first recall the basics of ultrametric geometry and then we construct a measure of global contrast. In Section 6, the new model of edges of color images is proposed. In Section 7, we define ultrametrics for the purpose of hierarchical segmentation. Finally, Section 8 contains some concluding remarks.

2 Generalized Voronoi Tessellations

2.1 Definitions

In this section, we present the mathematical framework for the rest of the paper and we introduce the notations.

Let $\Omega \subset \mathbb{R}^2$ be a compact connected domain in the plane. A **pseudo-metric** on Ω is an application $\psi : \Omega \times \Omega \rightarrow \mathbb{R}^+$ satisfying, for any $x, y, z \in \Omega$, the conditions:

1. *Reflexivity:* $\psi(x,x) = 0$.
2. *Symmetry:* $\psi(x,y) = \psi(y,x)$.
3. *Triangle Inequality:* $\psi(x,y) \leq \psi(x,z) + \psi(z,y)$.

Note that the only difference with the definition of a metric is that the usual *Separation* axiom was replaced by the weaker condition 1. Hence, we consider the equivalence classes $\hat{x}(\psi) = \{ y \in \Omega \mid \psi(x,y) = 0\}$ and we work directly on the quotient space $\hat{\Omega}(\psi) = \{\hat{x}(\psi) \mid x \in \Omega\}$. Thus, the projection of ψ on $\hat{\Omega}(\psi)$, is, by definition, a metric for the quotient space. Note that, if ψ is already a metric, then $\hat{\Omega}(\psi)$ coincides with the domain. Additionally, the existence of geodesics for ψ is assumed. In the sequel, the value of $\psi(x,y)$ will be referred as the *distance* between x and y.

The ***energy*** induced by a pseudo-metric ψ, with respect to a source point $s \in \Omega$, is defined as the single variable application $\psi_s : \Omega \rightarrow \mathbb{R}^+$ that measures the distance to s:

$$\psi_s(x) = \psi(s,x), \forall\, x \in \Omega.$$

The *energy* with respect to a set of sources $S = \{s_i\}_{i \in J}$ is given by the minimal individual *energy*:

$$\psi_S(x) = \inf_{s_i \in S} \psi_{s_i}(x), \forall\, x \in \Omega.$$

In the presence of multiple sources, a valuable information is provided by the ***influence zone*** of a source $s_i \in S$, or the set of points that are closer to s_i than to any other source, in the sense of ψ:

$$Z_i = \{x \in \Omega \mid \psi_{s_i}(x) < \psi_{s_j}(x), \forall s_j \in S,\, j \neq i\}.$$

Thus, the influence zone, or the *zone* for short, is a connected subset of the domain. Their union is noted by:

$$Z(\psi, S) = \bigcup_{i \in J} Z_i.$$

The complementary set of $Z(\psi, S)$ is called the ***medial set*** and is denoted by $M(\psi, S)$.

Hence, a pseudo-metric ψ and a set of sources S determine a partition of the domain Ω. In the sequel, this partition will be referred as the ***Voronoi tessellation***, or briefly the *tessellation*, and will be denoted by :

$$\Pi(\psi, S) = \{Z_i\}_{i \in J} \bigcup \{M(\psi, S)\}.$$

Note that every element of a *tessellation* is a union of elements of the quotient space $\hat{\Omega}(\psi)$.

Figure 1 illustrates these definitions with the canonical Euclidean space and a set of four source points, $S = \{s_0, s_1, s_2, s_3\}$. On the left, the Euclidean *energy* is shown and, on the right, the *tessellation* and the sources. In this case, since ψ is a metric, the quotient space $\hat{\Omega}(\psi)$ coincides with the domain. Additionally, the medial set $M(\psi, S)$ (shown in black) corresponds to the well know Voronoi diagram. Finally, the influence zones are in this case convex sets.

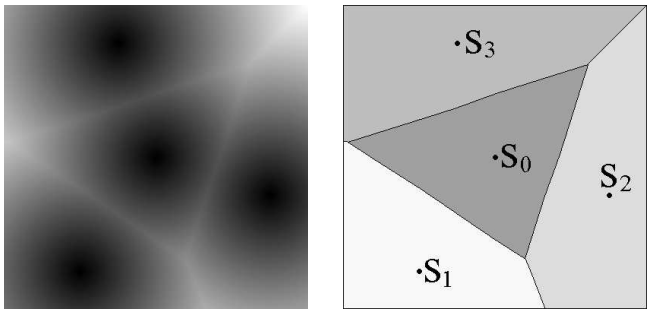

Figure 1: Euclidean *energy* and *tessellation*.

Two main differences of our approach with the standard framework of Voronoi tessellations [2, 29] should nevertheless be noted. First, by considering pseudo-metrics we have access to a larger class of spaces. Last, but not least, since we are interested in the application of these notions to image analysis, we discuss the definition of pseudo-metrics that depend on the image.

2.2 Segmentation of Digital Images

Thus, in this context, the image segmentation problem can be expressed in terms of the definition of a relevant pseudo-metric and the selection of a set of sources. Nonetheless, in practice, digital images are subsampled on the discrete grid. Consequently, important parts of the medial set may fall in the intergrid space. For the purpose of region based segmentation, an alternative to surround this problem is to consider a *tessellation* composed only by the zones. Thus, the elements of the medial set that would fall in the grid are assigned to one of their neighboring influence zones.

Then, an approximation of the image can be constructed by the assignation of a ***model*** to represent each influence zone. The model is determined by the distribution of the image values in the zone and, when it is constant, the valuation of each zone produces a piecewise constant approximation of the image, referred in the sequel as a ***mosaic image***.

In order to address the segmentation of color images in this framework, another practical problem is the definition of a metric on the color space. For this point, the CIE standards L^*ab and L^*uv [33] were adopted in this paper. These color representations ought to approximate a perceptually uniform color space. Though not perfect, their use provides two main advantages with respect to the basic RGB system: first, a separation of the color information into one lightness and two chromatic channels and, second, an approximation of the metric in the Riemannian color space by the Euclidean metric.

3 Path Metrics

A first approach for the definition of a pseudo-metric on an image domain is the study of paths between couples of points.

A **path** between two points $x,\ y \in \Omega$ is an injective continuous function $\gamma\ : [a,b] \rightarrow \Omega$ such that $\gamma(a) = x$ and $\gamma(b) = y$. The image of γ is then a simple curve in Ω. The set of paths from *x* to *y* is noted by Γ_{xy} and the set of paths in Ω is noted by Γ_Ω.

A **length structure** for Ω [16] is a map $e : \Gamma_\Omega \rightarrow \mathbb{R}^+$ that satisfies the following conditions:

1. If γ is constant, then $e(\gamma) = 0$.
2. If γ is the juxtaposition of γ_1 and γ_2, then $e(\gamma) = e(\gamma_1) + e(\gamma_2)$.
3. e is invariant under changes of parameterization.

A length structure e can be used to define a pseudo-metric, which we call the ***path metric*** associated to e, by considering the minimal value of e along all the paths joining two points:

$$\psi(x,y) = \inf_{\gamma \in \Gamma_{xy}} e(\gamma),\ \forall\, x, y \in \Omega.$$

A particularly interesting type of path metrics occurs when e can be expressed as the integral of a *potential function* $P : \Omega \rightarrow \mathbb{R}^+$:

$$e(\gamma) = \int_0^L P(\gamma(l))\, dl,$$

where l denotes the arc-length parameter. In this case, $\hat{x}(\psi)$, the equivalence class of a point x, corresponds to the largest connected set with null potential that contains x. Thus, if the potential is strictly positive, the quotient space $\hat{\Omega}(\psi)$ coincides with the domain. The metrics of this type are usually known as **weighted distance transforms**. Additionally, the relation between the *energy* and the potential is given in this case by the *eikonal equation* [7].

Weighted distances are widely used in computer vision, where the issue becomes the definition of a relevant potential from the image data for a particular problem. Examples of applications include shape from shading [18], continuous scale morphology [19], shape recovery [22], active contour models [7], differential morphology [23], tubular shape extraction [10] and perceptual grouping [9].

4 The Path Variation

In this section, we discuss the application of the path metric obtained by considering the variation on the paths as the length structure.

4.1 Definition

Let $[a,b] \subset \mathbb{R}$ be an interval and (X,d) a metric space.

Consider a function $f : [a,b] \rightarrow X$, a finite partition of $[a,b]$, $\sigma = \{t_0, ..., t_n\}$, such that $a = t_0 < t_1 < ... < t_n = b$ and denote by Φ the set of such partitions. The *variation* of f is defined as the (possibly infinite) number given by the formula:

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

Note that, if $X = \mathbb{R}$, then $v(f)$ corresponds to the *total variation* of f, the well known functional introduced by Jordan [17].

In the case of two variable functions, we consider the path metric induced by the variation, i.e., the minimal variation of the function on all the paths between two points:

The ***path variation*** of a function $u : \Omega \rightarrow X$, is defined as:

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

Note that, in contrast to the usual notion of total variation for functions of multiple variables [30], the path variation is defined pointwise. For further details about the path variation for real-valued functions, the reader is referred to [1]

The **component** of u containing x, designates the maximal connected subset of points $y \in \Omega$ such that $u(y) = u(x)$. By definition, the component containing x coincides with the equivalence class $\hat{x}(V^u)$. Thus, the quotient space $\hat{\Omega}(V^u)$ is the space of components of the function. Moreover, for a set of sources S, each element of the *tessellation* $\Pi(V^u, S)$ is a union of components of u. Hence, the operator that associates $\Pi(V^u, S)$ to the function u is connected [31] and its application simplifies the image while preserving its contour information. Therefore, the path variation presents a particular interest for image analysis.

4.2 Implementation for Color Images

For color images, we have $X = \mathbb{R}^3$, and d corresponds to the distance in the color space.

In a discrete domain, the choice of a digital connectivity (usually 4, 6 or 8 connectivity) determines a notion of component and of vicinity. Thus, the component space of the function u can be represented by an adjacency graph G, where the nodes correspond to discrete components and each link joins two neighboring components. Since the quotient space $\hat{\Omega}(V^u)$ is the space of components, we propose to construct the discrete path variation directly on G.

Hence, in the case of color images, the implementation of the path variation is reduced to finding a path of minimal cost on G, when the nodes of the graph are weighted by distance between the two neighboring components. For the examples presented in this paper, the color difference in the spaces L^*ab or L^*uv was used. The problem can then be solved using a greedy algorithm [11, 21]. The complexity of this straightforward implementation for the path variation is then $O(Nlog(N))$, where N denotes the total number of discrete components of the image.

4.3 The Extrema Mosaic

The path variation is an interesting pseudo-metric for a local level of analysis of the image. Indeed, since its definition (1) is based on an sum along the paths, the *tessellations* $\Pi(V^u, S)$ are very sensitive to the location of the sources. Therefore, in order to construct such *tessellations*, the set

Figure 2: Original image and extrema mosaic.

of sources S must be selected with care. First, the sources should be physically representative of the image content. Second, each significant feature should contain at least one of them. In the case of color images, the lightness extrema appear as natural candidates for the sources.

The *extrema tessellation* of an image u is defined as the *Voronoi tessellation* $\Pi(V^u, ext(u))$, where $ext(u)$ denotes the set of extremal components of the lightness channel L^* of u. A mosaic image determined by this *tessellation* is called an ***extrema mosaic*** of u.

Figure 2 shows the extrema mosaic of a natural image. The original image is on the left and the mosaic image, with the color at the source as the model, is on the right. This example illustrates four properties of the method. First, the blur in the original image is reduced, as can be observed on the background. This effect is due to the low number of sources in these regions and the fact that components belonging to blurred contours and transition zones are not extremal; consequently, they are absorbed by one neighboring zone. Second, as shown on the wolves' fur, the texture information is preserved in the simplified image because of the high density of extrema on these regions. Third, the contrast of the image is enhanced. Finally, note how the contour information is preserved in the simplified image.

In summary, the choice of the path variation as the pseudo-metric and the spatial distribution of the sources determine a *tessellation* where a balance between simplification and content conservation is obtained. The extrema mosaic is a natural reconstruction of the image that can be seen as a first level of abstraction for the image information. Its application as a parameter-free pre-segmentation method is illustrated in the next sections.

5 Ultrametrics

When a *Voronoi tessellation* with a small number of zones is required, the sensitivity of the path variation to the location of the sources may become a drawback. For this reason, in the sequel, the extrema mosaic is used as the starting point for the construction of a different type of pseudo-metrics, called the ultrametrics. These pseudo-metrics are more appropriated for coarse level of analysis of the image, as they are closely related to the families of nested partitions of the domain.

5.1 Definitions

In this paragraph, the basic properties of ultrametric geometry are recalled.

A **ultrametric** is a special type of metric for which the usual *Triangle Inequality* is replaced by the stronger relation:

$$\psi(x,y) \leq max\{\psi(x,z), \psi(z,y))\}, \ \forall x,y,z \in \Omega. \quad (2)$$

From a geometric point of view, the previous inequality can be interpreted as follows: all the triangles in a ultrametric space are either isosceles or equilateral.

Furthermore, as a consequence of (2), the structure of neighborhoods differs significantly from the usual Euclidean space. First, all the points in a ball of center x and radius r can be considered as the center. Second, two non-disjoint ultrametric balls are always concentric. Thus, the set of all the balls of a fixed radius r determines a partition of the domain. Hence, the sets of ultrametric balls of radius r, as r increases, produce a family of nested partitions of the domain. Furthermore, the radii of the balls determine a **stratification index** for the family of partitions and the resulting structure is called a **stratified hierarchy** of partitions. Conversely, every stratified hierarchy of partitions determines a ultrametric on the domain. For further details, the reader is referred to [4].

Therefore, in our case, when the *tessellation* is determined by a ultrametric, the previous properties imply that replacing a source $s_i \in S$ by another point $s_i' \in Z_i$ does not modify the *tessellation*. Moreover, the problem of selecting a set of sources can be addressed in this case through the choice of a radius r.

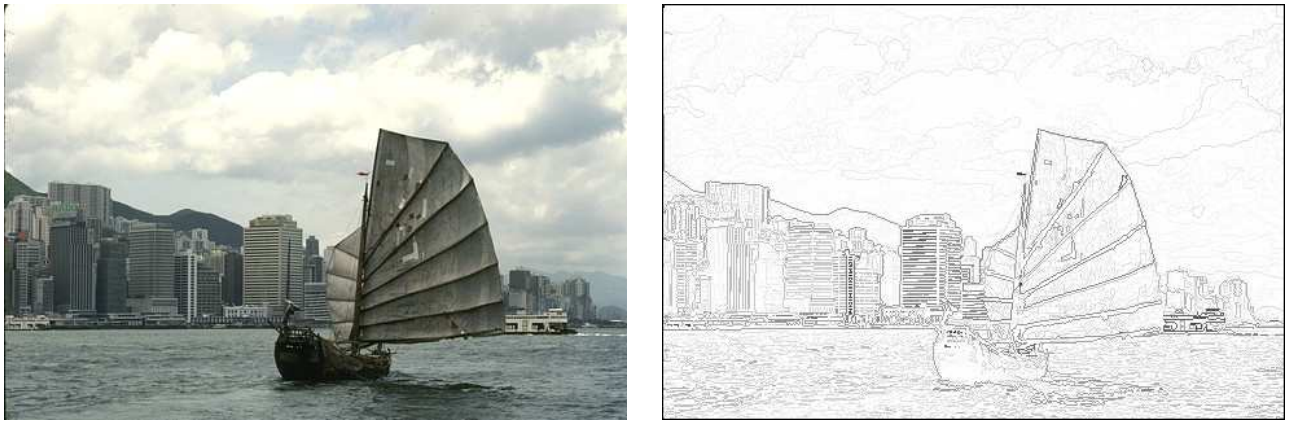

Figure 3: Original image and extrema edges.

5.2 A Measure of Contrast for Color Images

In this paragraph, we construct a ultrametric that expresses a global notion of contrast on a color image. For this purpose, the bijection with the class of stratified hierarchies is fundamental, as it provides a constructive definition for this type of pseudo-metrics.

Indeed, because of the properties presented in the previous paragraph, the distance between any two points $x_i \in Z_i$ and $x_j \in Z_j$ in a ultrametric *tessellation* can be expressed as a **dissimilarity measure**, noted by d, between the zones Z_i and Z_j. In the discrete space, this remark allows to construct the pseudo-metric through a region merging strategy.

The idea of progressively merging regions of an initial partition has been used since the early days of computer vision to address segmentation problems [6]. In general, this type of methods, often called *bottom-up* approaches, can be implemented efficiently using a region adjacency graph (RAG) [14]. A RAG is an undirected graph where the nodes correspond to connected regions of the domain. The links encode the vicinity relation and are weighted by the dissimilarity.

Therefore, in this context, the choice of an initial partition and the definition of a dissimilarity measure determine an order for the merging. Then, removing the links of the RAG for increasing values of the dissimilarity and merging the corresponding regions produces a family of nested partitions.

Typically, the dissimilarity expresses a notion of resemblance between neighboring regions and many examples have been proposed in the vast literature on the subject. A simple case is the difference of the average color (or gray level) in the regions [6, 8, 20], noted by d^a. However, other authors consider also factors as the variance and the size of the regions [3, 27], the orientation and the texture [34].

Nevertheless, it should be noted that, in order to produce a stratified hierarchy of partitions, the dissimilarity must be increasing with the order of merging. Unfortunately, this condition is seldom satisfied in the examples found in the literature. When the dissimilarity is not increasing, a stratification index for the hierarchy of partitions can still be defined by considering any increasing function of the merging order. However, in this case, the resulting ultrametric is no longer directly related to the dissimilarity.

The goal in this paragraph was to construct a ultrametric expressing the global contrast of the original image. A natural candidate is the dissimilarity d^a. This option suffers nonetheless from two drawbacks: first, d^a is not increasing and, second, since its definition uses all the information in the zones, its value may not reflect the real contrast. As a consequence, a merging process governed by this dissimilarity can create artificial contours when the color inside the regions varies gradually.

Thus, in our case, the dissimilarity was constructed using only boundary information and was measured directly on the initial partition. For the examples presented in this paper, the dissimilarity, noted by d^c, was defined as the average color difference in the common boundary of the zones, measured in the extrema mosaic. As a consequence of this choice, d^c is increasing with the merging order and the corresponding ultrametric, noted by ψ^c, is strongly related to the contrast information provided by the original image.

A classical example of a stratified hierarchy comes from the construction of the watershed transform [5]. Intuitively, this method can be summarized as follows: the image, considered as a topographical surface, is flooded from its regional minima. The water forms lakes in the valleys and, when two lakes meet, they are merged. Thus, increasing levels of water produce coarser partitions. When the image is the modulus of a gradient, the resulting hierarchy is known as the *dynamics* [15]. In terms of a region merging process, the initial partition is composed by the catchment basins of the minima and the dissimilarity is defined as the height of the lowest pass point between two adjacent lakes, i.e., the minimal value of the gradient in the common border of the regions [26]. Therefore, the dynamics hierarchy also induces a ultrametric. However, since its definition is based on a gradient image, the result depends on the choice of a the discrete approximation for the gradient.

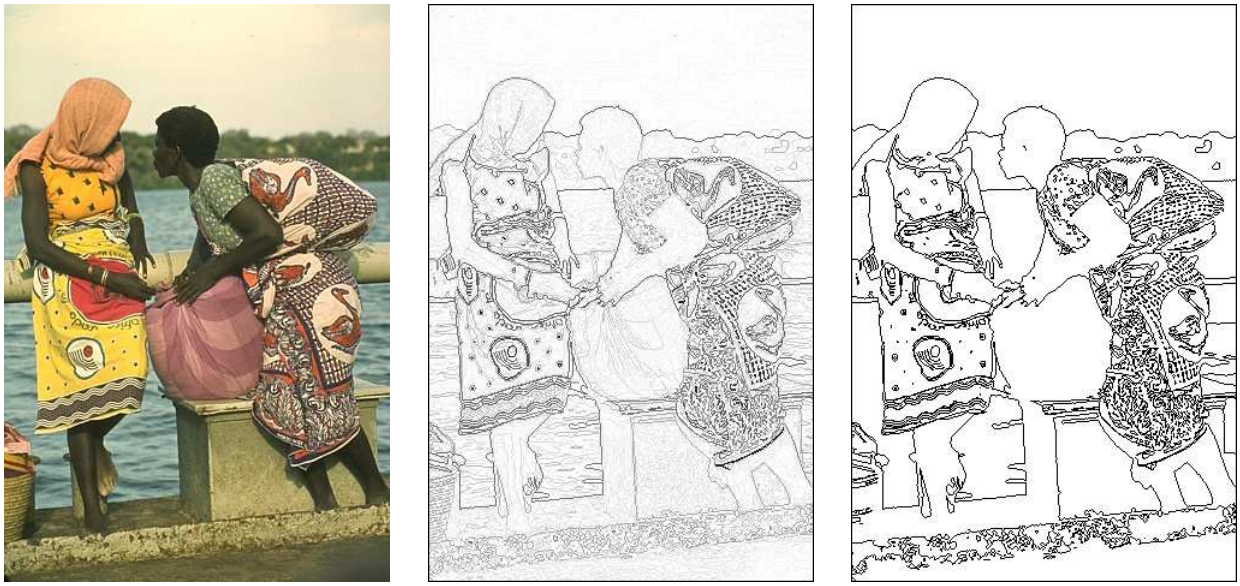

Figure 4: Original image, extrema edges and threshold.

6 The Extrema Edges

The perception of discontinuities seems to play a essential role for the interpretation of visual information in humans. Consequently, edge detection is a topic of great interest in the field of computer vision. Many edge models and detection techniques have been proposed through the years. Typically, the edge detection process is divided in two steps: a selection of possible edge points and the estimation of their relevance. Among the models proposed for the edges, one can cite the zero-crossings of the Laplacian, the maxima in the gradient direction and the crest lines of the gradient's modulus. However, in spite of their diversity, the strategy in many edge detection methods consists in a differential approach and the use of local image information to measure the relevance of the edge points [13]. In this section, the contrast measure ψ^c is used to model the edges of color images.

In addition to the features described in Sect. 3, a remarkable property of the extrema mosaic is the preservation and enhancement of the contour information. Indeed, thanks to the use of the path variation, the boundaries of the zones describe accurately the contours. Therefore, they constitute a sound set to look for edges in the image.

Once a set of candidates for possible edge points has been determined, the next problem is their valuation. In order to address this issue, the following notion presents a particular interest:

Consider a stratified hierarchy of partitions $\{\mathcal{P}_\lambda\}_\lambda$, the *saliency* of a point is defined as the highest index λ for which the point belongs to a boundary of $\mathcal{P}_\lambda$. The valuation of each point by its saliency determines a *saliency image*. The saliency image provides a compact description of the hierarchy: a threshold λ in this image supplies the set of boundaries of the corresponding partition $\mathcal{P}_\lambda$. This approach was first used in [28] to valuate the watershed lines of a gradient image on monochrome images.

In our case, the saliency image associated to the contrast ultrametric ψ^c was used for the valuation of the edges. The corresponding saliency image will be called the *extrema edges*.

Our model of edges presents the following advantages with respect to classical edge detectors. First, in order to well pose differentiation, these methods often perform a linear smoothing step; nevertheless, blurring implies a loss of resolution and the displacement of edges. In contrast, the extrema edges are precisely located and they preserve semantically important characteristics of contours such as corners and junctions. Second, our valuation method takes into account global contrast information. Last, but not least, a threshold in the extrema edges always provides a set of closed contours.

Figures 3 and 4 illustrate the properties of our edge model. The extrema edges are displayed with a high contrast represented by a low intensity. Note that the edges model accurately the contours in the scene and their value translates effectively the perceived contrast. The right image of Fig. 4 presents the closed contours corresponding to a threshold of 29% of the maximal contrast.

7 Derived Ultrametrics

The previous section showed the application of the ultrametric ψ^c to the extraction of the contrast information in a color image. However, contrast is just one among many factors considered in high-level vision tasks. In this section, ψ^c is used as the base to define new ultrametrics. For this purpose, other perceptually important characteristics of the zones, as their size, are used to complement the boundary information

Figure 5: From left to right: original images, *tessellations* and mosaic images with median color. Top: 15 zones and $\alpha = 0.23$. Bottom: 10 zones and $\alpha = 0.3$.

supplied by ψ^c.

Precisely, an *attribute*, a positive real valued function $\mathcal{A}$, is defined for every zone. The attribute is required to be increasing with the inclusion order. In general, $\mathcal{A}$ can be calculated using the internal information of the zone; the simplest example of an increasing attribute is the area of the zone. Then, starting from d^c, the dissimilarity associated to ψ^c, a new dissimilarity d' can be defined by the formula:

$$d'(Z_1, Z_2) = d^c(Z_1, Z_2) \cdot min\{\mathcal{A}(Z_1), \mathcal{A}(Z_2)\}.$$

Since $\mathcal{A}$ and d^c are increasing, d' induces a ultrametric where the internal information is also considered. For the examples presented in this paper, the attribute was defined as : $\mathcal{A}(Z) = A(Z)^\alpha$, where A denotes the area of the zone Z and the parameter $\alpha \geq 0$ weights the balance between contrast and area. Thus, for a fixed number of zones, the choice of α can be seen as the introduction of higher-level information, allowing the ultrametric to adapt to the image content.

Figures 5 and 6 show examples of segmentations obtained with this method. It can be observed that, in spite of the low number of zones and the simplicity of the attributes used in the definition of the ultrametric, the main features of the scene are recovered in the segmentations.

This approach can be seen as an extension of the *flooding hierarchies* from morphological segmentation where, during the watershed flooding, the area, the depth or the volume of the lakes is measured in order to provide a hierarchy for the gradient's minima [25, 26].

8 Conclusion and Perspectives

We formulated the problem of color image segmentation as a generalized Voronoi tessellation of the image domain.

The examples of pseudo-metrics presented in this paper can describe accurately the image structure when the scene is composed by relatively homogeneous objects. However, in order to apply this pseudo-metrics to the segmentation of highly textured or noisy images, a pre-processing step should be considered.

Present work includes the evaluation of the method with respect to other segmentation algorithms and the definition of pseudo-metrics where the information about the texture and the regularity of the contours is also taken into consideration.

Acknowledgements

The images used in this paper are part of the Berkeley Segmentation Dataset. We would like to thank the authors of [24] for their effort in providing ground-truth segmentation data.

References

[1] P. A. Arbeláez and L. D. Cohen. Energy partitions and image segmentation. Technical report, CEREMADE, 2003.

[2] F. Aurenhammer and R. Klein. *Handbook of Computational Geometry*, chapter 5: Voronoi Diagrams, pages 201–290. Elsevier Science Publishing, 2000.

[3] J.M. Beaulieu and M. Goldberg. Hierarchy in picture segmentation: a stepwise optimization approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(2):150–163, February 1989.

[4] J. P. Benzécri. *L'Analyse des Données. Tome I: La Taxinomie*. Dunod, Paris, 4 edition, 1984.

[5] S. Beucher and F. Meyer. *Mathematical Morphology in Image Processing*, chapter 12: The Morphological Approach to Segmentation: The Watershed Transformation, pages 433–481. Marcel Dekker, 1992.

[6] C. R. Brice and C. L. Fenema. Scene analysis using regions. *Artificial Intelligence*, 1:205–226, 1970.

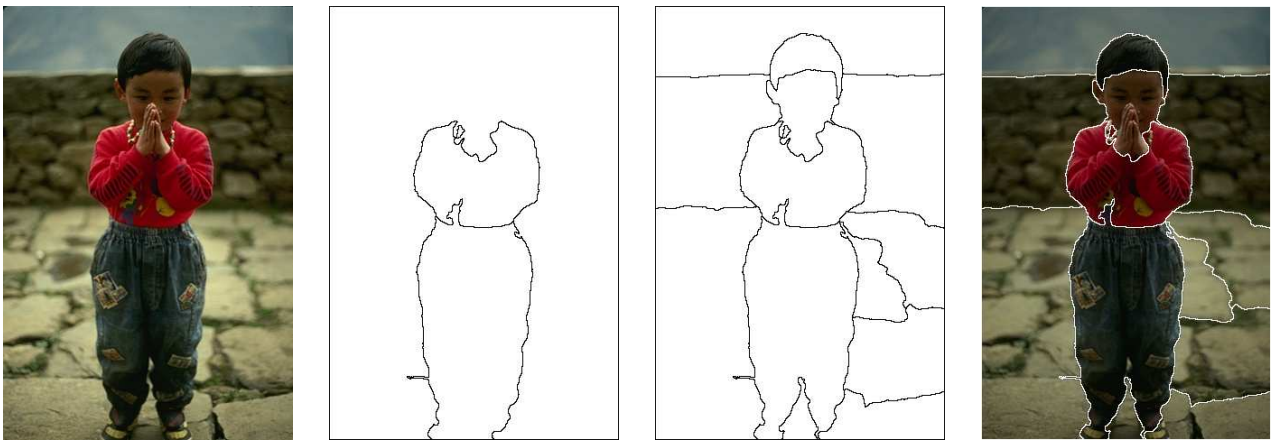

(a) (b) (c) (d)

Figure 6: (a): original image (b): *tessellation* with 3 zones and $\alpha = 0.6$ (c): *tessellation* with 13 zones and $\alpha = 0.6$ (d): contours of (c) on the original image.

[7] L. D. Cohen and R. Kimmel. Global minimum for active contour models: A minimal path approach. *International Journal of Computer Vision*, 24(1):57–78, August 1997.

[8] L. D. Cohen, L. Vinet, P. Sander, and A. Gagalowicz. Hierarchical region based stereo matching. In *Proc. 6th Scandinavian Conference on Image Analysis*, Oulu, Finland, 1989.

[9] L. D. Cohen. Multiple contour finding and perceptual grouping using minimal paths. *Journal of Mathematical Imaging and Vision*, 14(3):225–236, 2001.

[10] T. Deschamps and L. D. Cohen. Fast extraction of minimal paths in 3d images and applications to virtual endoscopy. *Medical Image Analysis*, 5(4):281–299, 2001.

[11] E. W. Dijkstra. A note on two problems in connection with graphs. *Numerische Mathemetic*, 1:269–271, 1959.

[12] P. G. L. Dirichlet. Uber die reduction der positiven quadratischen formen mit drei unbestimmten ganzen zalhen. *J. Reine Angew. Mathematik*, 40:209–227, 1850.

[13] D. A. Forsyth and J. Ponce. *Computer Vision: A Modern Approach*. Prentice Hall, 2003.

[14] L. Garrido, P. Salembier, and D. Garcia. Extensive operators in partition lattices for image sequence analysis. *Signal Processing*, 66(2):157–180, April 1998. Special Issue on Video Sequence Segmentation.

[15] M. Grimaud. New measure of contrast: Dynamics. In *Image Algebra and Morphological Processing III*, SPIE, San Diego, USA, 1992.

[16] M. Gromov. *Metric Structures for Riemannian and Non-Riemannian Spaces*. Birkhauser, Boston, 1999.

[17] C. Jordan. Sur la série de fourier. *Comptes Rendus de l'Académie des Sciences. Série Mathématique.*, 92(5):228–230, 1881.

[18] R. Kimmel and A. M. Bruckstein. Global shape from shading. *Computer Vision and Image Understanding*, 62(3):360–369, 1995.

[19] R. Kimmel, N. Kiryati, and A. M. Bruckstein. Distance maps and weighted distance transforms. *Journal of Mathematical Imaging and Vision*, 6:223–233, May 1996. Special Issue on Topology and Geometry in Computer Vision.

[20] G. Koepfler, C. Lopez, and J. M. Morel. A multiscale algorithm for image segmentation by variational method. *SIAM Journal on Numerical Analysis*, 31(1):282–299, 1994.

[21] R. Kruse and A. Ryba. *Data structures and program design in C++*. Prentice Hall, New York, 1999.

[22] R. Malladi and J.A. Sethian. A unified approach to noise removal, image-enhancement, and shape recovery. *IEEE Transactions on Image Processing*, 5(11):1554–1568, November 1996.

[23] P. Maragos and M. A. Butt. Curve evolution, differential morphology and distance transforms applied to multiscale and eikonal problems. *Fundamenta Informaticae*, 41:91–129, 2000.

[24] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proc. ICCV'01*, volume II, pages 416–423, Vancouver, Canada, 2001.

[25] F. Meyer, A. Oliveras, P. Salembier, and C. Vachier. Morphological tools for segmentation: Connected filters and watersheds. *Annals of Telecommunications*, 52(7-8):367–379, 1997.

[26] F. Meyer. Hierarchies of partitions and morphological segmentation. In Michael Kerckhove, editor, *Scale Space and Morphology in Computer Vision*, pages 161–182, 2001.

[27] P.F.M. Nacken. Image segmentation by connectivity preserving relinking in hierarchical graph structures. *PR*, 28(6):907–920, June 1995.

[28] L. Najman and M. Schmitt. Geodesic saliency of watershed contours and hierarchical segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(12):1163–1173, 1996.

[29] A. Okabe, B. Boots, K. Sugihara, and S. N. Chiu. *Spatial Tesselations: Concepts and Applications of Voronoi Diagrams*. Wiley, 2 edition, 2002.

[30] L.I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992.

[31] J. Serra and P. Salembier. Connected operators and pyramids. In SPIE, editor, *Image Algebra and Mathematical Morphology*, volume 2030, pages 65–76, San Diego CA., July 1993.

[32] G. M. Voronoi. Nouvelles applications des paramètres continus à la théorie des formes quadratiques. deuxième mémoire: Recherches sur les parallélo̊edres primitifs. *J. Reine Angew. Mathematik*, 134:198–287, 1908.

[33] G. Wyszecki and W. S. Stiles. *Color Science: Concepts and Methods, Quantitative Data and Formulas*. J. Wiley and Sons, 1982.

[34] W. Yu, J. Fritts, and F. Sun. A hierarchical image segmentation algorithm. In *Proc. ICME'02*, pages pp. 221–224, August 2002.