



HAL
open science

Ultrametric watersheds: a bijection theorem for hierarchical edge-segmentation

Laurent Najman

► **To cite this version:**

Laurent Najman. Ultrametric watersheds: a bijection theorem for hierarchical edge-segmentation. 2009. hal-00419373v1

HAL Id: hal-00419373

<https://hal.science/hal-00419373v1>

Preprint submitted on 9 Feb 2010 (v1), last revised 17 Dec 2010 (v2)

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.

Ultrametric watersheds: a bijection theorem for hierarchical edge-segmentation

Laurent Najman

Université Paris-Est, Laboratoire d'Informatique Gaspard-Monge, Equipe A3SI,
ESIEE Paris, France

Abstract. We study hierarchical segmentation in the framework of edge-weighted graphs. We define ultrametric watersheds as topological watersheds null on the minima. We prove that there exists a bijection between the set of ultrametric watersheds and the set of hierarchical edge-segmentations. We end this paper by showing how the proposed framework allows to see constrained connectivity as a classical watershed-based morphological scheme, which provides an efficient algorithm to compute the whole hierarchy.


Introduction

This paper¹ is a contribution to a theory of hierarchical (image) segmentation in the framework of edge-weighted graphs. Image segmentation is a process of decomposing an image into regions which are homogeneous according to some criteria. Intuitively, a hierarchical segmentation represents an image at different resolution levels.


In this paper, we introduce a subclass of edge-weighted graphs that we call ultrametric watersheds. Theorem 13 states that there exists a one-to-one correspondence, also called a bijection, between the set of indexed hierarchical edge-segmentations and the set of ultrametric watersheds. In other words, to any hierarchical edge-segmentation (whatever the way the hierarchy is built), it is possible to associate a representation of that hierarchy by an ultrametric watershed. Conversely, from any ultrametric watershed, one can infer a indexed hierarchical edge-segmentation.

This theorem is illustrated on Fig. 1, that is produced using the method proposed in [1]: what is usually done is to compute from an original image (Fig. 1.a) a hierarchical segmentation that can be represented by a dendrogram (Fig. 1.b, see section 4). The borders of the segmentations extracted from the hierarchy (such as the one seen in Fig. 1.c) can be stacked to form a map (Fig. 1.d) that allows for the visual representation of the hierarchical segmentation. Theorem 13 gives a characterization of the class of maps (called ultrametric watersheds) that represent a hierarchical segmentation; more surprisingly, Theorem 13 also states that the dendrogram can be obtained *after* the ultrametric watershed has been computed.


¹ This work was partially supported by ANR grant SURF-NT05-2_45825




(a) Original image



(b) Dendrogram of the hierarchical segmentation



(c) One segmentation extracted from the hierarchy



(d) An ultrametric watershed corresponding to the hierarchical segmentation

Fig. 1. An example of a hierarchical segmentation produced by the method of L. Guigues [1] *et al.*. The classical order for reading the images is (a), (b), (c), (d). But Theorem 13 states that the reading order can also be (a), (d), (c), (b) (see text).

Following [2], we can say that, independently of its theoretical interest, such a bijection theorem is useful in practice. Any hierarchical segmentation problem is a priori heterogeneous: assign to an edge-weighted graph an indexed hierarchy. Theorem 13 allows such classification problem to become homogeneous: assign to an edge-weighted graph a particular edge-weighted graph called ultrametric watershed. Thus, Theorem 13 gives a meaning to questions like: which hierarchy is the closest to a given edge-weighted graph with respect to a given measure or distance?

The paper is organised as follow. Related works are examined in section 1. We introduce segmentation on edges in section 2, and in section 3, we adapt the topological watershed framework from the framework of graphs with discrete weights on the nodes to the one of graphs with real-valued weights on the edges. We then define (section 4) hierarchies and ultrametric distances. In section 5, we introduce hierarchical edge-segmentations and ultrametric watersheds, the main result being the existence of a bijection between these two sets (Th. 13). We then propose (section 6) a convenient way to represent hierarchies as a discrete image. In the last part of the paper, we show how the proposed framework allows to see constrained connectivity [3] as a classical watershed-based morphological scheme, thus providing an efficient algorithm to compute the whole hierarchy.

Apart when otherwise mentionned, and to the best of the author's knowledge, all the properties and theorems formally stated in this paper are new. This paper is an extended version of [4].

1 Related works

1.1 Hierarchical clustering

From its beginning in image processing, hierarchical segmentation has been thought of as a particular instance of hierarchical classification [5]. One of the fundamental theorems for hierarchical clustering states that there exists a one-to-one correspondence between the set of indexed hierarchical classification and a particular subset of dissimilarity measures called ultrametric distances; This theorem is generally attributed to Johnson [6], Jardine *et al.* [7] and Benzécri [5]. Since then, numerous generalisations of that bijection theorem have been proposed (see [2] for a recent review).

Theorem 13 (see below) is an extension to hierarchical edge-segmentation of this fundamental hierarchical clustering theorem. Note that the direction of this extension is different from what is done classically in hierarchical clustering. For example, E. Diday [8] looks for proper dissimilarities that are compatible with the underlying lattice. An ultrametric watershed F is not a proper dissimilarity, *i.e.* $F(x, y) = 0$ does not imply that $x = y$ (see section 4). But F is an ultrametric distance (and thus a proper dissimilarity) on the set of connected components of $\{(x, y) | F(x, y) = 0\}$, those connected components being the regions of a segmentation.

Another point of view on our extension is the following: some authors assimilate classification and segmentation. We advocate that there exists a fun-

damental difference: in classification, we work on the complete graph, *i.e.* the underlying connectivity of the image (like the four-connectivity) is not used, and some points can be put in the same class because for example, their coordinates are correlated in some way with their color; thus a class is not always connected for the underlying graph. In the framework of edge-segmentation, any region of any level of a hierarchy of edge-segmentations is connected for the underlying graph. In other words, our approach yields a constrained classification, the constraint being the four-connectivity of the classes, or more generally any connection defining a graph (for the notion of connection and its links with segmentation, see [9,10].)

1.2 Hierarchical segmentation

There exist many methods for building a hierarchical segmentation [11], which can be divided in three classes: bottom-up, top-down or split-and-merge. A recent review of some of those approaches can be found in [3]. A useful representation of hierarchical segmentations was introduced in [12] under the name of *saliency map*. This representation has been used (under several names) by several authors, for example for visualisation purposes [1] or for comparing hierarchies [13].

In this paper, we show that any saliency map is an ultrametric watershed, and conversely.

1.3 Watersheds

For bottom-up approaches, a generic way to build a hierarchical segmentation is to start from an initial segmentation and progressively merge regions together [14]. Often, this initial segmentation is obtained through a watershed [12, 15, 16]. See [17] for a recent review of these notions in the context of mathematical morphology.

Among many others [18], topological watershed [19] is an original approach to watershed that modifies a map (e.g., a grayscale image) while preserving the connectivity of each lower cross-section. It has been proved [19, 20] that this approach is the only one that preserves altitudes of the passes (named connection values in this paper) between regions of the segmentation. Pass altitudes are fundamental for hierarchical schemes [12]. On the other hand, topological watersheds may be thick. A study of the properties of different kinds of graphs with respect to the thinness of watersheds can be found in [21, 22]. A useful framework is that of edge-weighted graphs, where watersheds are *de facto* thin (*i.e.* of thickness 1); furthermore, in that framework, a subclass of topological watersheds satisfies both the drop of water principle and a property of global optimality [23]. This subclass of topological watersheds can be seen as the limit, when the power of the weights tends to infinity for some specific energy function, of classical algorithms like graph cuts or random walkers [24].

In this paper, we translate topological watersheds from the framework of vertex-weighted-graphs to the one of edge-weighted graphs, and we identify

ultrametric watersheds, a subclass of topological watersheds that is convenient for hierarchical edge-segmentation.

2 Segmentation on edges

This paper is settled in the framework of edge-weighted graphs. Following the notations of [25], we present some basic definitions to handle such kind of graphs.

2.1 Basic notions

We define a *graph* as a pair $X = (V, E)$ where V is a finite set and E is composed of unordered pairs of V , *i.e.*, E is a subset of $\{\{x, y\} \subseteq V \mid x \neq y\}$. We denote by $|V|$ the cardinal of V , *i.e.*, the number of elements of V . Each element of V is called a *vertex* or a *point* (of X), and each element of E is called an *edge* (of X). If $V \neq \emptyset$, we say that X is *non-empty*.

As several graphs are considered in this paper, whenever this is necessary, we denote by $V(X)$ and by $E(X)$ the vertex and edge set of a graph X .

A graph X is said *complete* if $E = V(X) \times V(X)$.

Let X be a graph. If $u = \{x, y\}$ is an edge of X , we say that x and y are *adjacent* (for X). Let $\pi = \langle x_0, \dots, x_\ell \rangle$ be an ordered sequence of vertices of X , π is a *path* from x_0 to x_ℓ in X (or in V) if for any $i \in [1, \ell]$, x_i is adjacent to x_{i-1} . In this case, we say that x_0 and x_ℓ are *linked* for X . We say that X is *connected* if any two vertices of X are linked for X .

Let X and Y be two graphs. If $V(Y) \subseteq V(X)$ and $E(Y) \subseteq E(X)$, we say that Y is a *subgraph* of X and we write $Y \subseteq X$. We say that Y is a *connected component* of X , or simply a *component* of X , if Y is a connected subgraph of X which is maximal for this property, *i.e.*, for any connected graph Z , $Y \subseteq Z \subseteq X$ implies $Z = Y$.

Let X be a graph, and let $S \subseteq E(X)$. The *graph induced by S* is the graph whose edge set is S and whose vertex set is made of all points that belong to an edge in S , *i.e.*, $(\{x \in V(X) \mid \exists u \in S, x \in u\}, S)$.

Important remark. Throughout this paper $G = (V, E)$ denotes a connected graph, and the letter V (resp. E) will always refer to the vertex set (resp. the edge set) of G . We will also assume that $E \neq \emptyset$.

Let $S \subset E$. In the following, when no confusion may occur, the graph induced by S is also denoted by S .

Typically, in applications to image segmentation, V is the set of picture elements (pixels) and E is any of the usual adjacency relations, *e.g.*, the 4- or 8-adjacency in 2D [26].

If $S \subset E$, we denote by \overline{S} the *complementary set of S in E* , *i.e.*, $\overline{S} = E \setminus S$.

2.2 Edge-segmentation

A deep insight on our work is that we are working with edges and not with points: the minimal unit which we want to modify is an edge. Indeed, what we

need is a discrete space in which we can draw the border of a segmentation, so that we can represent that segmentation by its border; in other words, we want to be able to obtain the regions from their borders, and conversely. In that context, a desirable property is that the regions of the segmentation are the connected components of the complement of the border.

As illustrated in Fig. 2.b, this is not possible to achieve with the classical definition of a point-cut. Indeed, recall that a *partition* of V is a collection (V_i) of non-empty subsets of V such that any element of V is exactly in one of these subsets, and that a *point-cut* is the set of edges crossing a partition. Even if we add the hypothesis that any $(V_i, (V_i \times V_i) \cap E)$ is a connected graph, a V_i can be reduced to an isolated vertex, as the circled grey-point of Fig. 2.b. In that case, the complement of the point-cut, being a set of edges, does not contain that isolated vertex. The correct space to work with is the one of edges, and this motivates the following definitions.

Definition 1 *A set $C \subset E$ is an (edge-)cut (of G) if each edge of C is adjacent to two different nonempty connected components of \overline{C} . A graph S is called an (edge-)segmentation (of G) if $\overline{E(S)}$ is a cut. Any connected component of a segmentation S is called a region (of S).*




Fig. 2. Illustration of edge-segmentation and edge-cut. (a) A graph X . (b) A subgraph of X which is not a edge-segmentation of X : the circled grey-point is isolated, and if the point-cut D is the set of dotted-lines edges, \overline{D} contains only two connected components, instead of the expected three (see text). (c) An edge-segmentation of X ; the set C of dotted-lines edges is the associated edge-cut of X .

As mentioned above, the previous definitions of cut and segmentation (illustrated on Fig. 2.c) are not the usual ones. One can remark the complement of the complement of a cut is the cut itself, and that any segmentation gives a partition, the converse being false. In particular, Prop. 2.i below states that

there is no isolated point in an edge-segmentation. If we need an isolated point x , it is always possible to replace x with an edge $\{x', y'\}$. An application of the framework of hierarchical edge-segmentation to constrained connectivity (where isolated points are present) is described in section 7.

It is interesting to state the definition of a segmentation from the point of view of vertices of the graph. A graph X is said to be *spanning* (for V) if $V(X) = V$. We denote by ϕ the map that associates, to any $X \subset G$, the graph $\phi(X) = \{V(X), \{\{x, y\} \in E \mid x \in V(X), y \in V(X)\}\}$. We observe that $\phi(X)$ is maximal among all subgraphs of G that are spanning for $V(X)$, it is thus a closing on the lattice of subgraphs of G [27]. We call ϕ the edge-closing.

Property 2 *A graph $S \subseteq G = (V, E)$ is a segmentation of G if and only if*

- (i) *The graph induced by $E(S)$ is S ;*
- (ii) *S is spanning for V ;*
- (iii) *for any connected component X of S , $X = \phi(X)$.*

Proof. Let S be a segmentation of G . Then \overline{S} is a cut, in other word, any edge $v = \{x, y\} \notin E(S)$ is such that x and y are in two different connected components of \overline{S} . As G is connected, that implies that S is spanning for V . Moreover, $E(S)$ is the set of all edges of S , and as S is spanning for G , the graph induced by $E(S)$ is $(V, E(S)) = S$. Let X be a connected component of S , suppose that there exists $v = \{x, y\} \in E$ such that x and y belong to X and $v \notin E(X)$. But then $v \notin E(S)$ and thus x and y are in two different connected components of S , a contradiction.

Conversely, let S be a graph satisfying (i), (ii) and (iii) and let $v = \{x, y\} \notin E(S)$. As, by (ii), S is spanning for V , assertion (iii) implies that x and y are in two different connected components of $\overline{E(S)}$. Assertion (i) implies that there is no isolated points in S , thus \overline{S} is a cut and thus S is a segmentation of G . \square

2.3 Binary watershed

Let X be a subgraph of G . We note $X + u = (V(X) \cup u, E(X) \cup \{u\})$. In other words, $X + u$ is the graph whose vertice-set is composed by the points of $V(X)$ and the points of u , and whose edge-set is composed by the edges of $E(X)$ and u . An edge $u \in \overline{E(X)}$ is said to be *W-simple* (for X) (see [19]) if X has the same number of connected components as $X + u$.

A subgraph X' of G is a *thickening* (of X) if:

- $X' = X$, or if
- there exists a graph X'' which is a thickening of X and there exists an edge u W-simple for X'' and $X' = X'' + u$.

A subgraph X of G such that there does not exist a W-simple edge for X is called a *binary watershed* (of G).

The following property is a consequence of the definitions of segmentation and binary watershed.

Property 3 A graph $X \subseteq G = (V, E)$ is a segmentation of G if and only if X is a binary watershed of G and if X is induced by $E(X)$.

Proof. If X is a segmentation, then $\overline{E(X)}$ is a cut; let $u \in \overline{E(X)}$, u is adjacent to two different non-empty connected components of $E(X)$, in other word u is not W -simple for X . Thus any segmentation is a binary watershed.

Conversely, let X be a binary watershed, any $u \notin E(X)$ is not W -simple for X (and thus u is adjacent to two different connected components of X). If furthermore X is induced by $E(X)$ then $\overline{E(X)}$ is a cut. \square .

Thus, starting from a set of edges X , a segmentation is obtained by iterative thickening steps until idempotence. The next section extends the binary watershed approach to edge-weighted graphs.

3 Topological watershed

3.1 Edge-weighted graphs

We denote by \mathcal{F} the set of all maps from E to \mathbb{R}^+ . Given any $F \in \mathcal{F}$, the positive numbers $F(u)$ for $u \in E$ are called the *weights* and the pair (G, F) an *edge-weighted graph*. Whenever no confusion can occur, we will denote the edge-weighted graph (G, F) by F .

For applications to image segmentation, we take for weight $F(u)$, where $u = \{x, y\}$ is an edge between two pixels x and y , a dissimilarity measure between x and y (e.g., $F(u)$ equals the absolute difference of intensity between x and y ; see [28] for a more complete discussion on different ways to set the map F for image segmentation). Thus, we suppose that the salient contours are located on the highest edges of (G, F) .

Let $\lambda \in \mathbb{R}^+$ and $F \in \mathcal{F}$, we define $F[\lambda] = \{v \in E \mid F(v) \leq \lambda\}$. The graph (induced by) $F[\lambda]$ is called a *(cross)-section* of F . A connected component of a section $F[\lambda]$ is called a *component of F (at level λ)*.

We define $\mathcal{C}(F)$ as the set composed of all the pairs $[\lambda, C]$, where $\lambda \in \mathbb{R}^+$ and C is a component of the graph $F[\lambda]$. We call *altitude of $[\lambda, C]$* the number λ . We note that one can reconstruct F from $\mathcal{C}(F)$; more precisely, we have:

$$F(v) = \min\{\lambda \mid [\lambda, C] \in \mathcal{C}(F), v \in E(C)\}$$

For any component C of F , we set $h(C) = \min\{\lambda \mid [\lambda, C] \in \mathcal{C}(F)\}$. We define $\mathcal{C}^*(F)$ as the set composed by all $[h(C), C]$ where C is a component of F . The set $\mathcal{C}^*(F)$, called the *component tree of F* [29,30], is a finite subset of $\mathcal{C}(F)$ that is widely used in practice for image filtering. Note that the previous equation also holds for $\mathcal{C}^*(F)$:

$$F(v) = \min\{\lambda \mid [\lambda, C] \in \mathcal{C}^*(F), v \in E(C)\}$$

We will make use of the component tree in the proof of Pr. 12.

A *(regional) minimum* of F is a component X of the graph $F[\lambda]$ such that for all $\lambda_1 < \lambda$, $F[\lambda_1] \cap E(X) = \emptyset$. We remark that a minimum of F is a subgraph of G and not a subset of vertices of G ; we also remark that any minimum X of F is such that $|V(X)| > 1$.

We denote by $\mathcal{M}(F)$ the graph whose vertex set and edge set are, respectively, the union of the vertex sets and edge sets of all minima of F . In Fig. 3, $\mathcal{M}(F)$ is in bold. Note that $\mathcal{M}(F)$ is induced by $E(\mathcal{M}(F))$. As a convenient notation, and when no confusion can occur, we will sometimes write $X \in \mathcal{M}(F)$ if X is a connected component of $\mathcal{M}(F)$.

3.2 Topological watersheds on edge-weighted graphs

In that section, we extend the definition of topological watershed [19] to edge-weighted graphs, and we give an original characterization of topological watersheds in that framework (Th. 7).

Let $F \in \mathcal{F}$. An edge u such that $F(u) = \lambda$ is said to be *W-destructible* (for F) with lowest value λ_0 if there exists λ_0 such that, for all λ_1 , $\lambda_0 < \lambda_1 \leq \lambda$, u is W-simple for $F[\lambda_1]$ and if u is not W-simple for $F[\lambda_0]$.

A *topological watershed* (on G) is a map that contains no W-destructible edges.

A map F' is a *topological thinning* (of F) if:

- $F' = F$, or if
- there exists a map F'' which is a topological thinning of F and there exists an edge u W-destructible for F'' with lowest value λ such that $\forall v \neq u, F'(v) = F''(v)$ and $F'(v) = \lambda_0$, with $\lambda \leq \lambda_0 < F''(v)$.

An illustration of a topological watershed can be found in Fig. 3.

The *connection value* between $x \in V$ and $y \in V$ is the number

$$F(x, y) = \min\{\lambda \mid [\lambda, C] \in \mathcal{C}(F), x \in V(C), y \in V(C)\} \quad (1)$$

In other words, $F(x, y)$ is the altitude of the lowest element $[\lambda, C]$ of $\mathcal{C}(F)$ such that x and y belong to C (rule of the least common ancestor).

The connection value is a practical way to know if an edge is W-destructible. The following property is a translation of prop. 2 in [31] to the framework of edge-weighted graphs.

Property 4 (Prop. 2 in [31]) *Let $F \in \mathcal{F}$. An edge $v = \{x, y\} \in E$ is W-destructible for F with lowest value λ if and only if $\lambda = F(x, y) < F(v)$.*

Two points x and y are *separated* (for F) if $F(x, y) > \max\{\lambda_1, \lambda_2\}$, where λ_1 (resp. λ_2) is the altitude of the lowest element $[\lambda_1, c_1]$ (resp. $[\lambda_2, c_2]$) of $\mathcal{C}(F)$ such that $x \in c_1$ (resp. $y \in c_2$). The points x and y are *λ -separated* (for F) if they are separated and $\lambda = F(x, y)$.

The map F' is a *separation* of F if, whenever two points are λ -separated for F , they are λ -separated for F' .

If X and Y are two subgraphs of G , we set $F(X, Y) = \min\{F(x, y) \mid x \in X, y \in Y\}$.




Fig. 3. Illustration of topological watershed. (a) An edge-weighted graph F . (b) A topological watershed of F . The minima of (a) are $(\{m, i\}, \{p, l\}, \{g, h\}, \{c, d\}, \{g, c\}, \{h, d\})$ and are in bold in (a).

Theorem 5 (Restriction to minima [19]). Let $F' \leq F$ be two elements of \mathcal{F} . The map F' is a separation of F if and only if, for all distinct minima X and Y of $\mathcal{M}(F)$, we have $F'(X, Y) = F(X, Y)$.

A graph X is flat (for F) if for all $u, v \in E(X)$, $F(u) = F(v)$. If X is flat, the altitude of X is the number $F(X)$ such that $F(X) = F(v)$ for any $v \in E(X)$.

We say that F' is a strong separation of F if F' is a separation of F and if, for each $X' \in \mathcal{M}(F')$, there exists $X \in \mathcal{M}(F)$ such that $X \subseteq X'$ and $F(X) = F(X')$.

Theorem 6 (strong separation [19]). Let F and F' in \mathcal{F} with $F' \leq F$. Then F' is a topological thinning of F if and only if F' is a strong separation of F .

In other words, topological thinnings are the only way to obtain a watershed that preserves connection values.

In the framework of edge-weighted graphs, topological watersheds allows for a simple characterization.

Theorem 7. A map F is a topological watershed if and only if:

- (i) $\mathcal{M}(F)$ is a segmentation of G ;
- (ii) for any edge $v = \{x, y\}$, if there exist X and Y in $\mathcal{M}(F)$, $X \neq Y$, such that $x \in V(X)$ and $y \in V(Y)$, then $F(v) = F(X, Y)$.

Proof. Let F be a topological watershed. Thus there does not exist any edge W -destructible for F .

- Suppose that $\mathcal{M}(F)$ is not a segmentation of G . That means that there exists an edge $u = \{x, y\} \in \overline{E(\mathcal{M}(F))}$ such that x and y belongs to the

same connected component X of $\mathcal{M}(F)$. That implies that $F(u) > F(X) = F(x, y)$. By Pr. 4, that implies that the edge u is W-destructible for F , a contradiction. Thus $\mathcal{M}(F)$ is a segmentation of G .

- As F is a topological watershed, we have by Pr. 4 that for any $v = \{x, y\} \in E$, $F(x, y) = F(v)$. In particular, if there exist X and Y in $\mathcal{M}(F)$, $X \neq Y$, such that $x \in V(X)$ and $y \in V(Y)$, then $F(v) = F(X, Y)$.

Conversely, suppose that F satisfies (i) and (ii). By Pr. 4, for any edge $v = \{x, y\} \in E(\mathcal{M}(F))$, $F(v) = F(x, y) = F(X)$, and thus $\mathcal{M}(F)$ does not contain any edge W-destructible for F . As, by (i), $\mathcal{M}(F)$ is a segmentation, any edge $v \notin E(\mathcal{M}(F))$ satisfies (ii). By Pr. 4, such an edge v is not W-destructible. Thus F contains no W-destructible edge and is a topological watershed. \square

Note that if F is a topological watershed, then for any edge $v = \{x, y\}$ such that there exists $X \in \mathcal{M}(F)$ with $x \in V(X)$ and $y \in V(X)$, we have $F(v) = F(X)$.

4 Hierarchies and ultrametric distances

Let Ω be a finite set. A *hierarchy* H on Ω is a set of parts of Ω such that

- (i) $\Omega \in H$
- (ii) for every $\omega \in \Omega$, $\{\omega\} \in H$
- (iii) for each pair $(h, h') \in H^2$, $h \cap h' \neq \emptyset \implies h \subset h'$ or $h' \subset h$.

The (iii) can be expressed by saying that two elements of a hierarchy are either disjoint or nested.

An *indexed hierarchy* on Ω is a pair (H, μ) , where H denotes a given hierarchy on Ω and μ is a positive function, defined on H and satisfying the following conditions:

- (i) $\mu(h) = 0$ if and only if h is reduced to a singleton of Ω ;
- (ii) if $h \subset h'$, then $\mu(h) < \mu(h')$.

Hierarchy are usually represented using a special type of tree called *dendrograms* (Fig. 4). The leafs of the tree are the data that are to be classified, while the branching point (the junctions) are the agglomeration of all the data that are below that point. In that sense, one can see that, for a given h , $\mu(h)$ corresponds to the “level” of aggregation, where the elements of h have been aggregated for the first time.

Recall that a *dissimilarity on Ω* is a map d from the Cartesian product $\Omega \times \Omega$ to the set \mathbb{R} of real numbers such that: $d(\omega_1, \omega_2) = d(\omega_2, \omega_1)$, $d(\omega_1, \omega_1) = 0$ and $d(\omega_1, \omega_2) \geq 0$ for all $\omega_1, \omega_2, \omega_3 \in \Omega$. The dissimilarity d is said to be proper whenever $d(\omega_1, \omega_2) = 0$ implies $\omega_1 = \omega_2$.

A *distance d (on Ω)* is a proper dissimilarity that obeys the triangular inequality $d(\omega_1, \omega_2) \leq d(\omega_1, \omega_3) + d(\omega_3, \omega_2)$ where ω_1, ω_2 and ω_3 are any three points of the space.




Fig. 4. Hierarchical trees. We have $\lambda_1 < \lambda_3 < \lambda_4$ and $\lambda_2 < \lambda_4$.

The ultrametric inequality [32] is stronger than the triangular inequality. An *ultrametric distance* (on Ω) is a proper dissimilarity such that $d(\omega_1, \omega_2) \leq \max(d(\omega_1, \omega_3), d(\omega_2, \omega_3))$ for all $\omega_1, \omega_2, \omega_3 \in \Omega$.

Note that any given partition (Ω_i) of the set Ω induces a large number of trivial ultrametric distances: $d(\omega_1, \omega_1) = 0$, $d(\omega_1, \omega_2) = 1$ if $\omega_1 \in \Omega_i, \omega_2 \in \Omega_j, i \neq j$, and $d(\omega_1, \omega_2) = a$ if $i = j, 0 < a < 1$. The general connection between indexed hierarchies and ultrametric distances goes back to Benzécri [5] and Johnson [6]. They proved there is a bijection between indexed hierarchies and ultrametric distances, both defined on the same set. Indeed, associated with each indexed hierarchy (H, μ) on Ω is the following ultrametric distance:

$$d(\omega_1, \omega_2) = \min\{\mu(h) \mid h \in H, \omega_1 \in h, \omega_2 \in h\}. \quad (2)$$

In other words, the distance $d(\omega_1, \omega_2)$ between two elements ω_1 and ω_2 in Ω is given by the smallest element in H which contains both ω_1 and ω_2 . Conversely, each ultrametric distance d is associated with one and only one indexed hierarchy.

Observe the similarity between Eq. 2 and Eq. 1. Indeed, connection value is an ultrametric distance on V whenever $F > 0$. More precisely, we can state the following property, whose proof is a simple consequence of Eq.2 and Eq. 1.

Property 8 *Let $F \in \mathcal{F}$. Then $F(X, Y)$ is an ultrametric distance on $\mathcal{M}(F)$. If furthermore, $F > 0$, then $F(x, y)$ is an ultrametric distance on V .*

Let Ψ be the application that associates to any $F \in \mathcal{F}$ the map $\Psi(F)$ such that for any edge $\{x, y\} \in E$, $\Psi(F)(\{x, y\}) = F(x, y)$. It is straightforward to see that $\Psi(F) \leq F$, that $\Psi(\Psi(F)) = \Psi(F)$ and that if $F' \leq F$, $\Psi(F') \leq \Psi(F)$. Thus Ψ is an opening on the lattice (\mathcal{F}, \leq) [33]. We observe that the subset of strictly positive maps that are defined on the complete graph $(V, V \times V)$ and that are open with respect to Ψ is the set of ultrametric distances on V . The mapping Ψ is known under several names, including “subdominant ultrametric” and “ultrametric opening”. It is well known that Ψ is associated to the simplest method for hierarchical classification called single linkage clustering [7, 34], closely related to Kruskal’s algorithm [35] for computing a minimum spanning tree.

Thanks to Th. 7, we observe that if F is a topological watershed, then $\Psi(F) = F$. However, an ultrametric distance d may have plateaus, and thus the weighted complete graph $(V, V \times V, d)$ is not always a topological watershed. Nevertheless, those results underline that topological watersheds are related to hierarchical classification, but not yet to hierarchical edge-segmentation; the study of such relations is the subject of the rest of the paper.

5 Hierarchical edge-segmentations, saliency and ultrametric watersheds

Informally, a hierarchical segmentation is a hierarchy of connected regions. However, in our framework, if a segmentation is a partition, the converse is not true (see Pr. 2); thus, as the union of two disjoint connected subgraphs of G is not a connected subgraph of G , the formal definition is slightly more involved.

A *hierarchical (edge-)segmentation (on G)* is an indexed hierarchy (H, μ) on the set of regions of a segmentation S of G , such that for any $h \in H$, $\phi(\cup_{X \in h} X)$ is connected (ϕ being the edge-closing defined in section 2).

For any $\lambda \geq 0$, we denote by $H[\lambda]$ the graph induced by $\{\phi(\cup_{X \in h} X) | h \in H, \mu(h) \leq \lambda\}$. The following property is an easy consequence of the definition of a hierarchical segmentation.

Property 9 *Let (H, μ) be a hierarchical segmentation. Then for any $\lambda \geq 0$, the graph $H[\lambda]$ is a segmentation of G .*

Proof. Let (H, μ) be a hierarchical segmentation, and let $\lambda \geq 0$. Suppose that $H[\lambda]$ is not a segmentation, i.e. that $\overline{H[\lambda]}$ is not a cut. Then there exist a connected component X of $H[\lambda]$ and $v = \{x, y\} \in \overline{H[\lambda]}$ such that $x \in X$ and $y \in X$. That implies that $\phi(X) \neq X$, a contradiction with the definition of a hierarchical segmentation. \square

Prop. 8 implies that the connection value defines a hierarchy on the set of minima of F . If F is a topological watershed, then by Th. 7, $\mathcal{M}(F)$ is a segmentation of G , and thus from any topological watershed, one can infer a hierarchical segmentation. However, $F[\lambda]$ is not always a segmentation: if there exists a minimum X of F such that $F(X) = \lambda_0 > 0$, for any $\lambda_1 < \lambda_0$, $F[\lambda_1]$ contains at least two connected components X_1 and X_2 such that $|V(X_1)| = |V(X_2)| = 1$. Note that the value of F on the minima of F is not related to the position of the divide nor to the associated hierarchy of minima/segmentations. This leads us to introduce the following definition.

Definition 10 *A map $F \in \mathcal{F}$ is an ultrametric watershed if F is a topological watershed, and if furthermore, for any $X \in \mathcal{M}(F)$, $F(X) = 0$.*

Definition 10 directly yields to the nice following property, illustrated in Fig. 5, that states that any level of an ultrametric watershed is a segmentation and conversely.

Property 11 *A map F is an ultrametric watershed if and only if for all $\lambda \geq 0$, $F[\lambda]$ is a segmentation of G .*

Proof. Suppose that F is an ultrametric watershed, then it is a topological watershed, and by Th. 7.(i), $\mathcal{M}(F)$ is a segmentation of G . But as the value of F on its minima is null, then any cross-section of F is a segmentation of G .

Conversely, if for any $\lambda \geq 0$, $F[\lambda]$ is a segmentation of G , then F contains no W -destructible edge for F . Indeed, suppose that there exists an edge v W -destructible for F , let $\lambda = F(v)$, then v is W -simple for $F[\lambda]$. In other words, adding v to $F[\lambda]$ does not change the number of connected components of $F[\lambda]$. This is a contradiction with the definition of a segmentation. Hence F is a topological watershed. Furthermore, as $F[\lambda]$ is a segmentation for any $\lambda \geq 0$, the value of F on its minima is null, hence F is an ultrametric watershed. \square




Fig. 5. An example of an ultrametric watershed F and a cross-section of F .

By definition of a hierarchy, two elements of H are either disjoint or nested. If furthermore (H, μ) is a hierarchical segmentation, the graphs $\overline{E(H[\lambda])}$ can be stacked to form a map. We call *saliency map* [12] the result of such a stacking, *i.e.* a saliency map is a map F such that there exists (H, μ) a hierarchical segmentation with $F(v) = \min\{\lambda \mid v \in E(H[\lambda])\}$.

Property 12 *A map F is a saliency map if and only if F is an ultrametric watershed.*

Proof. If F is a saliency map, then there exists (H, μ) a hierarchical segmentation such that $F(v) = \min\{\lambda \mid v \in E(H[\lambda])\}$. But $F[\lambda] = \{v \mid F(v) \leq \lambda\} = \{v \mid \min\{\lambda \mid v \in E(H[\lambda])\} \leq \lambda\} = H[\lambda]$ and thus by Pr. 9, for any $\lambda \geq 0$, $F[\lambda]$ is a segmentation. By Pr. 11, F is an ultrametric watershed.

Conversely, let F be an ultrametric watershed, and let $\mathcal{C}^*(F)$ be the component tree of F . We build the pair (H, μ) in the following way: $h \in H$ if and only if there exists $[\lambda, C] \in \mathcal{C}^*(F)$ such that $h = \{X_i \mid X_i \in \mathcal{M}(F) \text{ and } X_i \subset C\}$; in that case, we set $\mu(h) = \lambda$.

Then (H, μ) is a hierarchical edge-segmentation. Indeed, let h and h' two elements of H such that there exists $[\lambda, X]$ and $[\lambda', Y]$ in $\mathcal{C}^*(F)$ with $h = \{X_0, \dots, X_p \mid X_i \in \mathcal{M}(F) \text{ and } X_i \subset X\}$ and with $h' = \{Y_0, \dots, Y_n \mid Y_i \in \mathcal{M}(F) \text{ and } Y_i \subset Y\}$.

- by Th. 7, $\mathcal{M}(F)$ is a segmentation,
- we set $\lambda_{\max} = \max\{F(v) \mid v \in E\}$, it is easy to see that $[\lambda_{\max}, (V, E)] \in \mathcal{C}^*(F)$, thus $\{\cup_{X \in \mathcal{M}(F)} X\} \in H$;
- any minimum X of F is such that $[0, X]$ belongs to $\mathcal{C}^*(F)$, thus $\{X\} \in H$;
- furthermore, h and h' are either disjoint or nested:
 - either disjoint: suppose that $X \cap Y = \emptyset$, in that case h and h' are also disjoint;
 - or nested: suppose that $X \cap Y \neq \emptyset$, then as X and Y are two connected components of the cross-sections of F , either $X \subset Y$ or $Y \subset X$; suppose that $X \subset Y$; by reordering the X_i and the Y_i , that means that $X_i = Y_i$ for $i = 0, \dots, p$, $p < n$. In other words, $h \subset h'$.
- by construction, $\mu(h) = 0$ if and only if there exists $X \in \mathcal{M}(F)$ such that $h = \{X\}$;
- If $h \subset h'$, then $\mu(h) < \mu(h')$, because in that case, $X \subset Y$ and thus $\lambda < \lambda'$.

Thus (H, μ) is an indexed hierarchy on $\mathcal{M}(F)$.

Furthermore, $\phi(\cup_{X_i \in h} X_i)$ is connected: more precisely, as $\mathcal{M}(F)$ is a segmentation, and as X is a connected component of the cross-sections of F , $\phi(\cup_{X_i \in h} X_i) = X$. Thus (H, μ) is a hierarchical edge segmentation. \square

The following theorem, a corollary of Prop. 12, states the equivalence between hierarchical segmentations and ultrametric watersheds. It is the main result of this paper.

Theorem 13. *There exists a bijection between the set of hierarchical edge-segmentations on G and the set of ultrametric watersheds on G .*

Proof. By Pr. 12, any ultrametric watershed is a saliency map, thus for any ultrametric watershed, there exists an associated hierarchical edge-segmentation.

Conversely, for any hierarchical edge-segmentation, there exists a unique saliency map, thus by Pr. 12, a unique ultrametric watershed. \square

Th. 13 states that any hierarchical segmentation can be represented by an ultrametric watershed. Such a representation can easily be built by stacking the border of the regions of the hierarchy (see Pr. 9 and 12, but also [1,12,13]). More interestingly, Th. 13 also states that any ultrametric watershed yields a hierarchical segmentation. As the definition of topological watershed is constructive, this is an incentive to searching for algorithmic schemes that directly compute

the whole hierarchy. An exemple of such an application of Th. 13 is developed in section 7.

As there exists a one-to-one correspondence between the set of indexed hierarchies and the set of ultrametric distances, it is interesting to search if there exists a similar property for the set of hierarchical segmentations. Let d be the ultrametric distance associated to a hierarchical segmentation (H, μ) . We call *ultrametric contour map* (associated to (H, μ)) the map d_E such that:

1. for any edge $v \in E(H[0])$, then $d_E(v) = 0$;
2. for any edge $v = \{x, y\} \in \overline{E(H[0])}$, $d_E(v) = d(X, Y)$ where X (resp. Y) is the connected component of $H[0]$ that contains x (resp. y).

Property 14 *A map F is an ultrametric watershed if and only if F is the ultrametric contour map associated to a hierarchical segmentation.*

Proof. Let F be an ultrametric watershed. By Pr. 8, $F(X, Y)$ is an ultrametric distance on $\mathcal{M}(F)$. By Pr. 11, F is a saliency map, hence there exists a hierarchical segmentation (H, μ) such that $F(v) = \min\{\lambda \mid v \in E(H[\lambda])\}$. In particular,

1. for any edge $v \in E(H[0])$, then $F(v) = 0$;
2. for any edge $v = \{x, y\} \in \overline{E(H[0])}$, $F(v) = F(X, Y)$ where X (resp. Y) is the connected component of $H[0]$ that contains x (resp. y).

Hence F is an ultrametric contour map associated to a hierarchical segmentation.

Conversely, let d_E be an ultrametric contour map associated to a hierarchical segmentation (H, μ) . Then by Th. 7, d_E is a topological watershed. Indeed, as H is a hierarchical segmentation, $H[0] = \mathcal{M}(d_E)$ is a segmentation of G , and furthermore for any edge $v = \{x, y\}$, if there exist X and Y in $\mathcal{M}(d_E)$, $X \neq Y$, such that $x \in V(X)$ and $y \in V(Y)$, then $d_E(v) = d(X, Y) = d_E(X, Y)$.

Moreover, for any $v \in \mathcal{M}(d_E)$, $d_E(v) = 0$, hence d_E is an ultrametric watershed. \square

6 Representations of hierarchical edge-segmentations in Khalimski grids

As we mentionned in section 2.2, one of the motivations of this work is to be able to imbed the hierarchical segmentation in a discrete space in a way that can be represented. Until now, we have used the classical representation of a graph for all of our examples.

For the purpose of visualisation, it is enough to represent the image by a grid of double resolution. For example, with the usual four connectivity in 2D, each pixel will be the center of a 3x3 neighborhood, and if two pixels share an edge, the two corresponding neighborhoods will share 3 elements corresponding to that edge. The representation of an ultrametric watershed with double resolution can be seen in Fig. 6.a.

A convenient interpretation of the doubling of the resolution can be given in the framework of cubical complexes, that have been popularized in computer

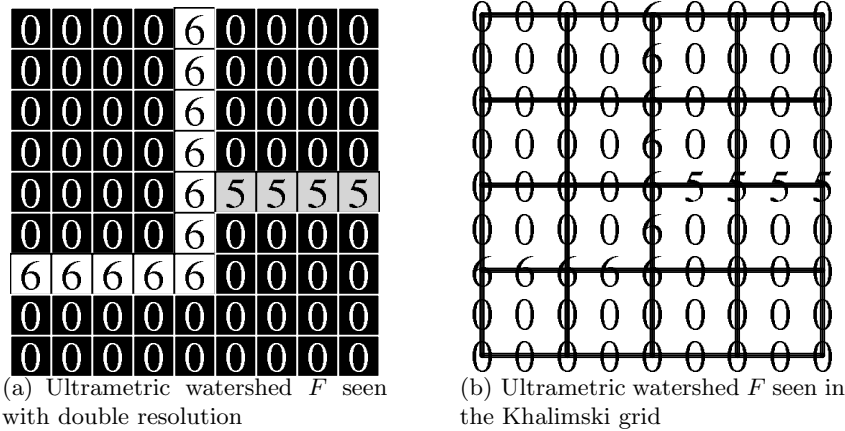


Fig. 6. Two possible representations of the ultrametric watershed F of Fig. 5.a.

vision by E. Khalimski [36], but can be found earlier in the literature, originally in the work of P.S. Alexandroff [37, 38].

Intuitively, a cubical complex can be seen as a set of elements of various dimensions (cubes, squares, segments and points) with specific rules between those elements. The traditional vision of a numerical image as being composed of pixels (elementary squares) in 2D or voxels (elementary cubes) in 3D leads to a natural link between numerical images and complexes. The representation of an ultrametric watershed in the Khalimski grid can be seen in Fig. 6.b.

The framework of complexes is useful in the study of topological properties [39]. It is indeed possible to provide a formal treatment of watersheds in complexes, which we will not do in this paper. The interested reader can have a look at [40].

7 Ultrametric watershed for constrained connectivity

Let us illustrate the usefulness of the proposed framework by providing an original way of revisiting constrained connectivity hierarchical segmentations [3], which leads to efficient algorithms. This section is meant as an illustration of our framework, and, although it is self-sufficient, technical details can be somewhat difficult to grasp for someone not familiar with the watershed-based segmentation framework of mathematical morphology [17]. We plan to provide more information in an extended version of that section that will appear in a journal paper.

In this section, we propose to compute an ultrametric watershed that corresponds to the constrained connectivity hierarchy of a given image. We show that, in the framework of edge-weighted segmentations, constrained connectivity can be thought as a classical morphological scheme, that consists of:

- computing a gradient;
- filtering this gradient by attribute filtering;
- computing a watershed of the filtered gradient.

We first give the formal definition of constrained connectivity, and then we move on to using ultrametric watersheds for computing such a hierarchy.

7.1 Constrained connectivity

This section is a reminder of P. Soille’s approach [3], using the same notations.

Let f be an application from V to \mathbb{R} , *i.e.* an image with values on the points. For any set of points $U \subseteq V$, we set

$$R_f(U) = \sup\{f(x) - f(y) \mid x, y \in U\}.$$

The number $R_f(U)$ is called the *range* of U (for f).

For any $x \in V$, and for any $\alpha \geq 0$, define [41] the α -connected component $\alpha\text{-CC}(x)$ as the set:

$$\alpha\text{-CC}(x) = \{x\} \cup \{y \in V \mid \text{there exists a path } \pi = \{x_0, \dots, x_n\}, n > 0, \\ \text{such that } R_f(\{x_i, x_{i+1}\}) \leq \alpha, \text{ for all } 0 \leq i < n\}$$

An essential property of the α -connected components of a point x is that they form an ordered sequence (*i.e.* a hierarchy) when increasing the value of α :

$$\alpha\text{-CC}(x) \subseteq \beta\text{-CC}(x)$$

whenever $\beta \geq \alpha$.

We now define the (α, ω) -connected component of an arbitrary point x as the largest α -connected component of x whose range is lower than ω ; more precisely,

$$(\alpha, \omega)\text{-CC}(x) = \sup\{\beta\text{-CC}(x) \mid \beta \leq \alpha \text{ and} \\ R_f(\beta\text{-CC}(x)) \leq \omega\} \quad (3)$$

The (α, ω) -CCs also define a hierarchy, that is called a *constrained connectivity hierarchy*. We have:

$$(\alpha, \omega)\text{-CC}(x) \subseteq (\alpha', \omega')\text{-CC}(x)$$

whenever $\alpha' \geq \alpha$ and $\forall \omega' \geq \omega$. In practice [3], we are interested in this hierarchy for $\alpha = \omega$, *i.e.*, for any $x \in V$ and any $\lambda \geq 0$, we are looking for $(\lambda, \lambda)\text{-CC}(x)$.

7.2 Ultrametric watershed for constrained connectivity

In that section, we show how to build a weighted graph on which the ultrametric watershed corresponding to the hierarchy of constrained connectivity can be computed. Intuitively, this weighted graph can be seen as the gradient of the

original image. We compute an ultrametric watershed for the hierarchy of α -connected components. We filter that watershed to obtain the family of (α, ω) -connected components. We then show how to directly compute the ultrametric watershed corresponding to the hierarchy of (α, ω) -connected components.

Constrained connectivity is a hierarchy of flat zones of f , in the sense where the 0-connected components of f are the zones of f where the intensity of f does not change. In a continuous world, such zones would be the ones where the gradient is null, *i.e.* $\nabla f = 0$. However, the space we are working with is discrete, and a flat zone of f can consist in a single point. In general, it is not possible to compute a gradient on the points or on the edges such that this gradient is null on the flat zones. To compute a gradient on the edges such that the gradient is null on the flat zones, we need to “double” the graph, for example we can do that by doubling the number of points of V and adding one edge between each new point and the old one (see Fig. 7).




Fig. 7. Doubling the graph. (a) Original graph with weights f on the vertices. (b) Double graph, with weights f on the vertices and the gradient F on the edges (see text).

More precisely, if we denote the points of V by $V = \{x_0, \dots, x_n\}$, we set $V' = \{x'_0, \dots, x'_n\}$ (with $V \cap V' = \emptyset$), and $E' = \{\{x_i, x'_i\} \mid 0 \leq i \leq n\}$. We then set $V_1 = V \cup V'$ and $E_1 = E \cup E'$. By construction, as $G = (V, E)$ is a connected graph, the graph $G_1 = (V_1, E_1)$ is a connected graph. We also extend f to V' , by setting, for any $x' \in V'$, $f(x') = f(x)$, where $\{x, x'\} \in E'$.

Let (V_1, E_1, F) be the weighted graph obtained from f by setting, for any $\{x, y\} \in E_1$, $F(\{x, y\}) = |f(x) - f(y)|$. The map F can be seen as the “natural gradient” of f [42]. It is easy to see that the flat zones of f , *i.e.* the 0-connected components of f are (in bijection with) the connected components of the set $\{v = \{x, y\} \in E_1 \mid F(\{x, y\}) = 0\}$.

Let W^1 be a topological watershed of F . From Th. 7 and Eq. 1, if $W^1(\{x, y\}) = \lambda$, there exists a path $\pi = \{x_0 = x, \dots, x_n = y\}$ linking x to y such that the altitude of any edge along π is below λ , *i.e.* we have, for any $0 \leq i < n$,

$F(\{x_i, x_{i+1}\}) = |f(x_i) - f(x_{i+1})| \leq \lambda$. The following property, the proof of which is left to the reader, states that the hierarchy of α -connected components is given by W^1 .

Property 15 *We have*

- W^1 is an ultrametric watershed;
- W^1 is uniquely defined (if W' is a topological watershed of F , then $W' = W^1$);
- let $\lambda \geq 0$ and let X be a connected component of the cross-section $W^1[\lambda]$; then for any $x \in V(X) \setminus V'$, λ -CC(x) = $V(X) \setminus V'$.




Fig. 8. Soille's (α, ω) -constrained connectivity hierarchy. (a) Original image. (b) Ultrametric watershed W^1 for the α -connectivity. (c) Ultrametric watershed W^2 for the constrained connectivity. (d) Ultrametric watersheds corresponding to one of the possible hierarchies of area-filterings on W^2 .

One can notice that R_f is increasing on 2^V , i.e. $R_f(X) \subset R_f(Y)$ whenever $X \subseteq Y$. Thus R_f is increasing on $\mathcal{C}(W^1)$, and by removing the connected components of $\mathcal{C}(W^1)$ that are below a threshold ω for R_f , we have an attribute

filtering which is idempotent (the values on the points do not change), thus it is a closing. More precisely, we denote by $(R_\lambda)_{\lambda \geq 0}$ the family of maps obtained by applying this closing on W^1 for varying λ , *i.e.*, for any $\lambda \geq 0$ and any $\{x, y\} \in E_1$, we set

$$R_\lambda(\{x, y\}) = \min\{\lambda' \mid [\lambda', C] \in \mathcal{C}(W^1), x \in V(C), y \in V(C), R_f(V(C)) \geq \lambda\}$$

In other words, the altitude for R_λ of the edge $\{x, y\}$ is the altitude of the lowest component of $\mathcal{C}(W^1)$ that contains both x and y and such that the range of that component is greater than λ .

The family $(R_\lambda)_{\lambda \geq 0}$ allows us to retrieve the (α, ω) -CCs of f : surprisingly, it can be shown that any R_λ is a topological watershed, and thus $\mathcal{M}(R_\lambda)$ is a segmentation from which it is easy to extract the (λ, λ) -connected component of a point, as the minimum of $\mathcal{M}(R_\lambda)$ that contains that point (See Pr. 16 below for a more formal setting).

Moreover, one can directly compute the ultrametric watershed associated to the hierarchy of (α, ω) -constrained connectivity. We set:

$$W^2(\{x, y\}) = \min\{R_f(V(C)) \mid [\lambda, C] \in \mathcal{C}(W^1), x \in V(C), y \in V(C)\} \quad (4)$$

In other words, the altitude for W^2 of the edge $\{x, y\}$ is the range of the lowest component of $\mathcal{C}(W^1)$ that contains both x and y . One can remark that Eq. 4 corresponds to Eq. 3 for the framework of edge-segmentation.

The following property, the proof of which is left to the reader, states that the hierarchy of (α, ω) -connected components is given by W^2 .

Property 16 *We have*

- $\forall \lambda \geq 0$, R_λ is a topological watershed;
- $\forall \lambda \geq 0$, $W^2[\lambda] = \mathcal{M}(R_\lambda)$;
- W^2 is an ultrametric watershed;
- W^2 is uniquely defined;
- let $\lambda \geq 0$ and let X be a connected component of the cross-section $W^2[\lambda]$; then for any $x \in V(X) \setminus V'$, (λ, λ) -CC(x) = $V(X) \setminus V'$.

Prop. 16, illustrated on Fig. 8.c, thus gives an efficient algorithm to compute the hierarchy of (α, ω) -constrained connectivity. Indeed, Eq. 4 can be computed in constant time [43] on $\mathcal{C}(W^1)$, which itself can be computed in quasi-linear time [30]. Such an algorithm is much faster than the one proposed in [3], that computes only one level of the hierarchy.

In a forthcoming paper, we will propose various data structures, including but not limited to component tree, that allows an efficient computation of hierarchical edge-segmentations. We will also study how to extend Prop. 16 in order to compute any granulometry of operators (strong hierarchies in the sense of [9]).

Visualising W^2 allows to assess some of the qualities of the hierarchy of constrained connectivity. One can notice in Fig. 8.c a large number of transition regions (small undesirable regions that persist in the hierarchy), and this

problem is known [44]. As W^2 is an image, a number of classical morphological schemes (*e.g.*, area filtering) can be used to remove those transition zones (see Fig. 8.d for an example). Studying the usefulness of such schemes is the subject of future research.

8 Conclusion

Fig. 9 is an illustration of the application of the framework developed in this paper to a classical hierarchical segmentation scheme based on attribute opening [12, 17, 29]. Fig. 10 shows some of the differences between applying such scheme and applying a classical watershed-based segmentation scheme, *e.g.* attribute opening followed by a watershed [15]. As watershed algorithms generally place watershed lines in the middle of plateaus, the two schemes give quite different results.

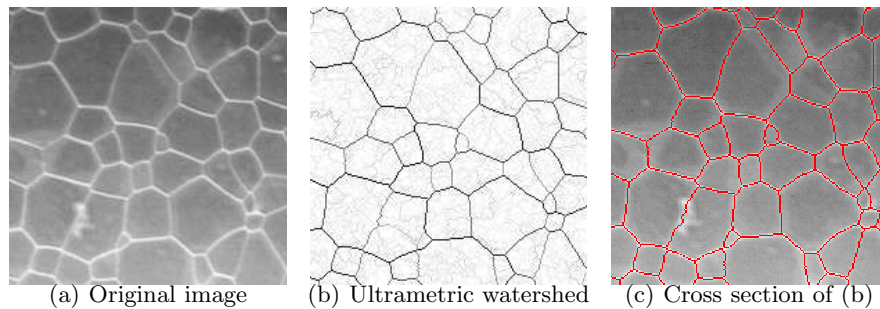


Fig. 9. Example of ultrametric watershed.

It is important to note that most of the algorithms proposed in the literature to compute saliency maps are not correct, often because they rely on wrong connection values or because they rely on thick watersheds where merging regions is difficult. Future papers will propose novel algorithms (based on the topological watershed algorithm [31]) to compute ultrametric watersheds, with proof of correctness.

On a more theoretical level, this work can be pursued in several directions.

- We will study lattices of watersheds [45] and will bring to that framework recent approaches like scale-sets [1] and other metric approaches to segmentation [13]. For example, scale-sets theory considers a rather general formulation of the partitioning problem which involves minimizing a two-term-based energy, of the form $\lambda C + D$, where D is a goodness-of-fit term and C is a regularization term, and proposes an algorithm to compute the hierarchical




Fig. 10. Zoom on a comparison between two watersheds of a filtered version of the image 9.a. Morphological filtering tends to create large plateaus, and both watersheds (a) and (b) are possible, but only (a) is a subset of a watershed of 9.a. No hierarchical scheme will ever give a result as (b).

segmentation we obtain by varying the λ parameter. As in the case of constrained connectivity (see section 7 above), we can hope that the topological watershed algorithm [31] can be used on a specific energy function to directly obtain the hierarchy.

- Subdominant theory (mentioned at the end of section 4) links hierarchical classification and optimisation. In particular, the subdominant ultrametric d' of a dissimilarity d is the solution to the following optimisation problem for $p < \infty$:

$$\min\{\|d - d'\|_p^p \mid d' \text{ is an ultrametric distance and } d' \leq d\}$$

It is certainly of interest to search if topological watersheds can be solutions of similar optimisation problems.

- Several generalisations of hierarchical clustering have been proposed in the literature [2]. An interesting direction of research is to see how to extend in the same way the topological watershed approach, for example for allowing regions to overlap.
- Last, but not least, the links of hierarchical edge-segmentation with connective segmentation [9] have to be studied.

Acknowledgements

The author would like to thank P. Soille for the permission to use Fig. 8.a (that was shot at ISMM'09 by the author), as well as J.-P. Coquerez for the permission to use Fig. 1.

For numerous discussions, the author feels indebted to (by alphabetical order): Gilles Bertrand, Michel Couprie, Jean Cousty, Christian Ronse, Jean Serra and Hugues Talbot.

References

1. Guigues, L., Cocquerez, J.P., Men, H.L.: Scale-sets image analysis. *International Journal of Computer Vision* **68**(3) (2006) 289–317
2. Barthélemy, J.P., Brucker, F., Osswald, C.: Combinatorial optimization and hierarchical classifications. *4OR: A Quarterly Journal of Operations Research* **2**(3) (2004) 179–219
3. Soille, P.: Constrained connectivity for hierarchical image decomposition and simplification. *IEEE Trans. Pattern Anal. Mach. Intell.* **30**(7) (July 2008) 1132–1145
4. Najman, L.: Ultrametric watersheds. In Springer, ed.: ISMM 09. Number 5720 in LNCS (2009) 181–192
5. Benzécri, J.: *L'Analyse des données: la Taxinomie*. Volume 1. Dunod (1973)
6. Johnson, S.: Hierarchical clustering schemes. *Psychometrika* **32** (1967) 241–254.
7. Jardine, N., Sibson, R.: *Mathematical taxonomy*. Wiley (1971)
8. Diday, E.: Spatial classification. *Discrete Appl. Math.* **156**(8) (2008) 1271–1294
9. Serra, J.: A lattice approach to image segmentation. *J. Math. Imaging Vis.* **24**(1) (2006) 83–130
10. Ronse, C.: Partial partitions, partial connections and connective segmentation. *J. Math. Imaging Vis.* **32**(2) (oct. 2008) 97–105
11. Pavlidis, T.: Hierarchies in structural pattern recognition. *Proceedings of the IEEE* **67**(5) (May 1979) 737–744
12. Najman, L., Schmitt, M.: Geodesic saliency of watershed contours and hierarchical segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* **18**(12) (December 1996) 1163–1173
13. Arbeláez, P.A., Cohen, L.D.: A metric approach to vector-valued image segmentation. *International Journal of Computer Vision* **69**(1) (2006) 119–126
14. Pavlidis, T. In: *Structural Pattern Recognition*. Volume 1 of Springer Series in Electrophysics. Springer (1977) 90–123 segmentation techniques, chapter 4–5.
15. Meyer, F., Beucher, S.: Morphological segmentation. *Journal of Visual Communication and Image Representation* **1**(1) (September 1990) 21–46
16. Meyer, F.: Morphological segmentation revisited. In: *Space, Structure and Randomness*. Springer (2005) 315–347
17. Meyer, F., Najman, L.: Segmentation, arbre de poids minimum et hiérarchies. In Najman, L., Talbot, H., eds.: *Morphologie mathématique 1: approches déterministes*. Lavoisier, Paris (2008) 201–233
18. Roerdink, J.B.T.M., Meijster, A.: The watershed transform: Definitions, algorithms and parallelization strategies. *Fundamenta Informaticae* **41**(1-2) (2001) 187–228
19. Bertrand, G.: On topological watersheds. *J. Math. Imaging Vis.* **22**(2-3) (May 2005) 217–230
20. Najman, L., Couprie, M., Bertrand, G.: Watersheds, mosaics and the emergence paradigm. *Discrete Appl. Math.* **147**(2-3) (2005) 301–324
21. Cousty, J., Bertrand, G., Couprie, M., Najman, L.: Fusion graphs: merging properties and watersheds. *J. Math. Imaging Vis.* **30**(1) (January 2008) 87–104
22. Cousty, J., Najman, L., Bertrand, G., Couprie, M.: Weighted fusion graphs: merging properties and watersheds. *Discrete Appl. Math.* **156**(15) (August 2008) 3011–3027
23. Cousty, J., Bertrand, G., Najman, L., Couprie, M.: Watershed cuts: minimum spanning forests, and the drop of water principle. *IEEE Trans. Pattern Anal. Mach. Intell.* (2009) To appear.

24. Couprie, C., Grady, L., Najman, L., Talbot, H.: Power watersheds: A new image segmentation framework extending graph cuts, random walker and optimal spanning forest. In: International Conference on Computer Vision (ICCV'09), Kyoto, Japan, IEEE (october 2009) To appear.
25. Diestel, R.: Graph Theory. Graduate Texts in Mathematics. Springer (1997)
26. Kong, T., Rosenfeld, A.: Digital topology: Introduction and survey. *Comput. Vision Graph. Image Process.* **48**(3) (1989) 357–393
27. Cousty, J., Najman, L., Serra, J.: Some morphological operators in graph spaces. In Springer, ed.: ISMM 09. Number 5720 in LNCS (2009) 149–160
28. Cousty, J., Bertrand, G., Najman, L., Couprie, M.: Watershed cuts: thinnings, shortest-path forests and topological watersheds. *IEEE Trans. Pattern Anal. Mach. Intell.* (2009) to appear.
29. Salembier, P., Oliveras, A., Garrido, L.: Anti-extensive connected operators for image and sequence processing. *IEEE Transactions on Image Processing* **7**(4) (April 1998) 555–570
30. Najman, L., Couprie, M.: Building the component tree in quasi-linear time. *IEEE Transactions on Image Processing* **15**(11) (2006) 3531–3539
31. Couprie, M., Najman, L., Bertrand, G.: Quasi-linear algorithms for the topological watershed. *J. Math. Imaging Vis.* **22**(2-3) (2005) 231–249
32. Krasner, M.: Espaces ultramétriques. *C.R. Acad. Sci. Paris* **219** (1944) 433–435
33. Leclerc, B.: Description combinatoire des ultramétriques. *Mathématique et sciences humaines* **73** (1981) 5–37
34. Gower, J., Ross, G.: Minimum spanning tree and single linkage cluster analysis. *Appl. Stats.* **18** (1969) 54–64
35. Kruskal, J.B.: On the shortest spanning subtree of a graph and the traveling salesman problem. *Proceedings of the American Mathematical Society* **7** (February 1956) 48–50
36. Khalimsky, E., Kopperman, R., Meyer, P.: Computer graphics and connected topologies on finite ordered sets. *Topology and its Applications* **36** (1990) 1–17
37. Alexandroff, P., Hopf, H.: *Topology*. Springer Verlag (1937)
38. Alexandroff, P.: Diskrete Räume. *Math. Sbornik* **2**(3) (1937) 501–518
39. Bertrand, G.: On critical kernels. *Comptes Rendus de l'Académie des Sciences **Série Math.** I*(345) (2007) 363–367
40. Cousty, J., Bertrand, G., Couprie, M., Najman, L.: Collapses and watersheds in pseudomanifolds. In: IWCIA. (2009)
41. Nagao, M., Matsuyama, T., Ikeda, Y.: Region extraction and shape analysis in aerial photographs. *Computer Graphics and Image Processing* **10**(3) (July 1979) 195–223
42. Mattiussi, C.: The Finite Volume, Finite Difference, and Finite Elements Methods as Numerical Methods for Physical Field Problems. *Advances in Imaging and Electron Physics* **113** (2000) 1–146
43. Bender, M., Farach-Colton, M.: The lca problem revisited. In: Latin American Theoretical INformatics. (2000) 88–94
44. Soille, P., Grazzini, J.: Constrained connectivity and transition regions. In Springer, ed.: ISMM 09. Number 5720 in LNCS (2009) 59–69
45. Cousty, J., Najman, L., Serra, J.: Raising in watershed lattices. In: 15th IEEE ICIP'08, San Diego, USA (October 2008) 2196–2199