Scale-sets Image Analysis

Laurent Guigues, Jean Pierre Cocquerez, Hervé Le Men

► To cite this version:

Laurent Guigues, Jean Pierre Cocquerez, Hervé Le Men. Scale-sets Image Analysis. International Journal of Computer Vision, 2006, 68 (3), pp.289-317. hal-00705364

HAL Id: hal-00705364

<https://hal.science/hal-00705364>

Submitted on 7 Jun 2012

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.

Scale-Sets Image Analysis

Laurent Guigues, Jean Pierre Cocquerez, Hervé Le Men

Université de Technologie de Compiègne(UTC), Institut Géographique National
(cocquerez@hds.utc.fr)

Abstract

This paper introduces a multi-scale theory of piecewise image modelling, called the scale-sets theory, and which can be regarded as a region-oriented scale-space theory. The first part of the paper studies the general structure of a geometrically unbiased region-oriented multi-scale image description and introduces the scale-sets representation, a representation which allows to handle such a description exactly. The second part of the paper deals with the way scale-sets image analyses can be built according to an energy minimization principle. We consider 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. We describe the way such energies arise from basic principles of approximate modelling and we relate them to operational rate/distorsion problems involved in lossy compression problems. We then show that an important subset of these energies constitutes a class of multi-scale energies in that the minimal cut of a hierarchy gets coarser and coarser as parameter λ increases. This allows us to devise a fast dynamic-programming procedure to find the complete scale-sets representation of this family of minimal cuts. Considering then the construction of the hierarchy from which the minimal cuts are extracted, we end up with an exact and parameter-free algorithm to build scale-sets image descriptions whose sections constitute a monotone sequence of upward global minima of a multi-scale energy, which is called the "scale climbing" algorithm. This algorithm can be viewed as a continuation method along the scale dimension or as a minimum pursuit along the operational rate/distorsion curve. Furthermore, the solution verifies a linear scale invariance property which allows to completely postpone the tuning of the scale parameter to a subsequent stage. For computational reasons, the scale climbing algorithm is approximated by a pair-wise region merging scheme: however the principal properties of the solutions are kept. Some results obtained with Mumford-Shah's piece-wise constant model and a variant are provided and different applications of the proposed multi-scale analyses are finally sketched.

1 Introduction

1.1 The need for multi-scale low level image descriptions

Following David Marr's computational theory of vision [33], a number of image analysis systems are based on a bottom-up architecture, made up of two global stages. A *low level* analysis of the incoming signal is first performed, which builds a low level description of the data in terms of primitive structures (characteristic points, contours, regions, optical flow, disparity map...). Based on this low level description, *high level* vision tasks are then issued (object recognition, scene modelling, interpretation...). In its most general — and radical — form, Marr's theory assumes a complete independence between the low and the high level stages of vision. The low level processes are supposed to output a general-purpose description of the image that is a description independent from any specific high level task.

Now, the structures which can be useful to high level reasoning can be located at arbitrary positions in an image, can have any orientation and can be of any size. Also, useful structures can have arbitrary image contrasts, that is they can be very salient as well as very poorly contrasted. Hence, a low level image description should be *uncommitted* in terms of position, orientation, size and contrast [31]. Size refers to the characteristic spatial extension of a phenomenon while contrast refers to its characteristic extension in the image's values space. Size and contrast can thus be thought of as two different measures of the "scale" of a structure along the two basic dimensions of an image: the spatial dimension and the

values dimension. Hence, low-level image analyses should be multi-scale analyses both in the spatial and in the value sense.

1.2 From scale-space...

Marr proposed to start image analyses with the elaboration of a low level description of the image in terms of *contours* – the *raw primal sketch* – and put forward the idea to look for transitions of different spatial extensions. Indeed, a progressive transition in an image content becomes more and more local as the distance of observation or as the size of an analyzing tool increases. “The critical observation” by Marr [33, p. 69] “is that, at their own scale, these things [the transitions in some image features] can all be thought of as spatially localized”. The raw primal sketch was based on convolving the image with the Laplacian of Gaussian kernels of different spatial extensions (variances), detecting the zero crossings and looking for the coincidence of responses for different scales (2 or 3).

Gaussian filtering does nothing else than blurring the image: it deletes the finest structures first but it also progressively ‘blurs’ the shapes, coarsening their *geometrical description*. To cope with this problem, Witkin [54] proposed to consider the family of convolutions results with gaussians of different size as a whole, a multi-scale representation of the signal called the *scale space*. The original idea of Witkin was to detect global contours at large scales and to track them downward along the scale dimension in order to get a fine geometrical description. The scale-space theory then greatly developed to become a fully coherent theory of multi-scale low-level image processing.

Important milestones in this development include Koenderink’s work [26], who formalized the Gaussian scale-space theory in $\mathbb{R}^n$, introducing a fundamental principle of multi-scale analysis, the *causality principle*, which formalizes a very commonsensical idea: details can only disappear by zooming out. To act as a multi-scale analysis, a sequence of descriptions—or models—of a datum must not create information when scale increases. In other words, the scale axis is fundamentally oriented. Koenderink also related Gaussian convolutions with the solutions of the heat equation and thus transposed the scale-space theory from the filtering domain to the domain of continuous evolution modelling by partial differential equations (PDE). [2] then showed that the gaussian scale-space is the unique linear scale-space. Subsequent scale-space models, such as Perona and Malik anisotropic diffusion scheme [41] or the affine morphological scale-space [47, 1], involved non linear PDE. [1] finally derived axiomatically a complete catalogue of all the possible PDE generated scale-spaces, the nature of the analysis depending on the set of the invariances obeyed. Another important family of scale-spaces are generated by morphological operators [23, 46, 40, 7].

1.3 ... to scale-sets

While low-level contour extraction and region delineation are very close problems, region-oriented image analysis, i.e. image segmentation, has followed a rather different development than edge detection. Following Leclerc, the image segmentation problem can be informally defined as the problem “to delineate regions that have, to a certain degree, coherent attributes in the image”. Indeed, “it is an important problem as, on the whole, objects or coherent physical processes in the scene project into regions with coherent image attributes” [30].

Now in general, “coherent” regions can be found at different scales in an image: global structures which are perceived as “coherent” wholes at a large scale of analysis are made up of subparts which appear “coherent” at a lower scale of analysis, and which in turn decompose into more coherent subparts... The existence of nested “coherent” structures in an image is obvious in the case of the texton/texture relationship; however it is much more general. For instance, a region which is “homogenous” to a certain degree, if it is not strictly constant, is then made up of more “homogenous” subparts. Hence, it can be argued that the general organization of the “coherent” regions of an image is a *hierarchical* organization, modelling the existence of nested meaningful structures in an image.

Despite this obvious remark, the image segmentation problem has been traditionally defined as a *partitioning* problem : decompose the image into *disjoint* regions [19], that is produce a single scale, ‘flat’, analysis of the image. A reason for this might be that contrarily to edge detection, region extraction is not always clearly conceived as a low level process [33]. Indeed, as semantic units of a scene get mapped into regional units in an image, there is a possible confusion between the low-level task of primitive structure

extraction from images and the high-level task of object recognition and delineation. Furthermore, even a very classical segmentation algorithm, say a simple region growing method, is able to delineate structures which correspond to certain semantic units, *when its parameters are appropriately tuned*. Hence the temptation to use a low-level segmentation tool as an object recognition tool.

However, automatic parameter tuning is a hard problem which greatly affects the practical operability of most segmentation algorithms. If the images a system has to analyze show a little variability, then it is dubious that a single parameter setting will allow to correctly segment all the images: the target objects will be over-segmented in some images, under-segmented in some other or even over-segmented in some parts of a given image and under-segmented in other parts of the same image. The reason for this is that the parameters of a segmentation algorithm are related to low-level criteria, such as homogeneity or geometrical regularity, which in general are not robust criteria in terms of object *discrimination*, except in very constrained experimental settings (e.g. industrial inspection). Moreover, homogeneity or goodness-of-fit terms are dependent on the image dynamics; hence, if the image contrast changes, the parameters have to be adapted in order for the result to remain satisfactory. This point is of course closely related to the fundamental role of invariance in vision and to the fact that vision algorithms should in general be independent of the viewing conditions (pose, distance, lighting...).

Like Morel and Solimini, we think that "the only basis parameter in computer vision is 'scale'" [38], that is the fineness of analysis, and that in general, the useful scale of analysis for a given task cannot be chosen *a priori*, that is before viewing the image. The task of selecting the useful scale(s) for a given application is inherently a high level recognition task which cannot be robustly solved based on the low level criteria embedded in segmentation algorithms. Hence, a low level segmentation tool should remain scale uncommitted and output a multi-scale description of the image in which higher level processes can navigate freely in order to find their objects of interest. We claim that devising multi-scale low-level analyses is the only way to get robust computer vision systems, and in particular to solve the difficult problems of parameter tuning and of the respect of fundamental invariances.

Those different ideas have led us to develop a multi-scale approach of low level region-oriented image analysis, which we have called the scale-sets theory because, as we shall see, it can be regarded as a region-oriented scale-space theory. The structure of the paper is as follows:

1.4 Outline of the paper

In the spirit of the scale-space theory, the first part of the paper (section 2) derives the general structure of unbiased multi-scale region-oriented image descriptions from a few basic axioms. It then introduces the scale-sets representation, an implicit representation which fully captures the structure of such descriptions. Its general properties are studied and our proposal is compared to related work.

The second part of the paper (section 3) studies the way scale-sets image descriptions can be built according to an energy-minimization principle. In a much classical way now, the mono-scale segmentation problem is put as the problem of finding a piece-wise defined model 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. We show that under broad assumptions, such energies induce monotone sequences of minimal cuts in a hierarchy when λ browses $\mathbb{R}^+$ and we design an efficient algorithm to compute the complete scale-sets representation of these sequences exactly. In a kind of feedback loop, this allows us to design a remarkable procedure to build the hierarchy from which the minimal cuts are extracted. The principle is to progressively regularize the energy by increasing λ and to follow the energy minimum along a monotone sequence of partitions. This strategy leads to a parameter-free region grouping criterion which we implement within a pair-wise region merging scheme. The scale-sets descriptions obtained are linearly scale invariant, which implies that the scale parameter λ is completely removed from the low-level segmentation stage. This property also allows to confer other invariances properties to the solution.

Before concluding, section 4 presents some experimental results. Appendix A recalls some usual definitions on partitions and hierarchies and takes up the global notations of the paper. Appendices B and C provide the proofs applicable to theorems 8 and 10.

2 Unbiased multi-scale segmentations and their scale-sets representation

2.1 Multi-scale segmentations

Let $\mathbf{P}$ be a partitioning algorithm depending on a positive real parameter λ . Given an image I defined over a domain $\mathcal{D}$ and a value of λ , $\mathbf{P}$ outputs a partition P_λ of $\mathcal{D}$. Intuitively, λ behaves as a scale parameter if increasing λ systematically coarsens the partition produced by $\mathbf{P}$, in the sense of the fineness relationship between partitions (see appendix A). Certainly because it is very intuitive, this definition has been adopted as a starting point for multi-scale segmentation by various authors, such as [27][48]. We show here that it can be derived from two fundamental principles of multi-scale analysis, namely causality and absence of boundary delocalization.

Let $P = (P_\lambda)_{\lambda \in \mathbb{R}^+}$ be the complete family of the partitions produced by $\mathbf{P}$ on I when λ browses $\mathbb{R}^+$. P can be thought of as a mapping from $\mathbb{R}^+$ to $\mathbb{P}(\mathcal{D})$ or equivalently as a point of $\mathbb{P}(\mathcal{D})^{\mathbb{R}^+}$. Now let's assume that $\mathcal{D}$ is a part of a topological space and that the algorithm $\mathbf{P}$ only produces partitions into *connected regions*. A partition P_λ is then fully determined by the set of its *boundaries*, which we denote by δP_λ .

The **causality principle** is certainly the most fundamental principle of multi-scale analysis [26]. From this principle, for any couple of scales $\lambda_2 > \lambda_1$, the "structures" found at scale λ_2 should find a "cause" at scale λ_1 . Following Witkin's original idea [54], we apply this principle to the *edges* produced by algorithm $\mathbf{P}$. In this case, the parameter λ behaves as a *scale parameter* if and only if for all $\lambda_2 > \lambda_1$, the boundaries of partition P_{λ_2} are in a one-to-one mapping with a subset of the boundaries of P_{λ_1} (their "cause"). Now, to be topologically meaningful, the mapping which relates δP_{λ_2} to a subset of δP_{λ_1} must be continuous. Hence following definition :

Definition 1 [causal structure] We say that a sequence $(P_\lambda)_{\lambda \in \mathbb{R}^+} \in \mathbb{P}(\mathcal{D})^{\mathbb{R}^+}$ has a **causal structure** when $\forall (\lambda_1, \lambda_2) \in \mathbb{R}^{+2}$ such that $\lambda_2 \geq \lambda_1$ there is a diffeomorphism ϕ of $\mathcal{D}$ such that $\phi(\delta P_{\lambda_2}) \subseteq \delta P_{\lambda_1}$.

In terms of partitions, P has a causal structure when for all $\lambda_2 > \lambda_1$, P_{λ_2} can be morphed to an under-partition of P_{λ_1} , that is can be obtained by first applying a continuous deformation to P_{λ_1} and then deleting some of its boundaries, i.e. merging some of its regions.

In general, the morphism ϕ is dependent on λ_1 and λ_2 . Thus write it down $\phi(\lambda_1, \lambda_2)$. In order to discard sudden displacements of boundaries when scale increases it is then natural to require ϕ 's continuity with respect to λ . As $\forall \lambda \in \mathbb{R}^+$, a trivial solution for $\phi(\lambda, \lambda)$ is $\mathcal{D}$'s identical mapping (Id), one is naturally led to the following definition:

Definition 2 [continuous causal structure] We say that a sequence $(P_\lambda)_{\lambda \in \mathbb{R}^+} \in \mathbb{P}(\mathcal{D})^{\mathbb{R}^+}$ has a **continuous causal structure** if it has a causal structure and the family of diffeomorphisms ϕ verifies

$$\forall \lambda \in \mathbb{R}^+ \quad \lim_{\varepsilon \rightarrow 0} \phi(\lambda, \lambda + \varepsilon) = Id.$$

Finally, if one assumes that the best *localization* of a structure can be achieved at the finest scale of analysis, one should always keep the *geometrical information* gathered at the finest scale and thus discard any boundary deformation when scale increases. This can only be reached by setting $\phi(\lambda_1, \lambda_2) = Id$ for all scales. This additional condition then leads to a structure in which δP_{λ_2} is always a subset of δP_{λ_1} , and thus to a structure in which P_{λ_2} is always an under-partition of P_{λ_1} . Hence the definition:

Definition 3 [unbiased multi-scale segmentation] A sequence $(P_\lambda)_{\lambda \in \mathbb{R}^+} \in \mathbb{P}(\mathcal{D})^{\mathbb{R}^+}$ is called an **unbiased multi-scale segmentation** if and only if the application $\lambda \rightarrow P_\lambda$ is increasing, that is if and only if

$$\forall (\lambda_1, \lambda_2) \in \mathbb{R}^{+2} \quad \lambda_2 \geq \lambda_1 \quad \Rightarrow \quad P_{\lambda_2} \geq P_{\lambda_1}. \quad (1)$$

As a conclusion, if one considers an analysis into *connected regions*, then the intuitive structure of a multi-scale segmentation — a monotone mapping from $\mathbb{R}^+$ to $\mathbb{P}(\mathcal{D})$ — follows from two basic principles applied to the *boundaries* of the partitions: i) the causality principle, which is a *topological/structural* principle and ii) a principle of *geometrical* accuracy. The scale parameter λ then controls a process

which does a topological/structural caricature of information but *no* geometrical caricature. These two conditions obviously meet the intuitive requirements for an "ideal" multi-scale analyzer: an analyzer which delineates *global* structures of an image with a fine, *local* geometry. This property is called "strong causality" in [38]. Note that a partitioning algorithm $\mathbf{P}$ whose results on an image verify the relation 1 can also be regarded as a geometrically unbiased multi-scale edge detector which produces closed contours. We shall get back to the issue of region/contour duality in section 2.3 below.

2.2 The scale-sets representation

As was argued above, a low-level segmentation algorithm should not focus on a specific scale but should output a multi-scale description of an image. In a region-oriented approach, the result should be an unbiased multi-scale segmentation, in the sense of definition 3. However, this objective poses a representational problem: a multi-scale segmentation P is an increasing mapping from $\mathbb{R}^+$ to $\mathbb{P}(\mathcal{D})$, and as this mapping starts from the *continuous* real line, it cannot be directly represented on a computer.

One can of course approximate the multi-scale structure by *sampling* the scale axis, i.e. by setting a series $(\lambda_1, \lambda_2 \dots \lambda_k)$ of increasing scales and collecting the associated partitions, ending with a pyramid of segmentations $(P_{\lambda_1} < \dots < P_{\lambda_k})$. However, the sampling points are necessarily arbitrary: they must be chosen *before* viewing the image. Such a strategy is thus *scale-committed*. Instead, we would like to get a complete representation of P which would allow to browse the multi-scale structure of an image freely, e.g. which would allow to set any value of λ *a posteriori*, and retrieve the corresponding partition P_λ . This can be achieved by what we call the scale-sets representation of P , as we now explain.

Rather than focusing on the partitions P_λ , consider the *regions* which compose them. Formally, consider $H = \bigcup_{\lambda \in \mathbb{R}^+} P_\lambda$. Obviously, as P is a sequence of monotone partitions, the regions which belong to H are either disjoint or nested. Thus, if P_0 is the absolute over-partition and if for a sufficiently large scale L , P_L is the absolute under-partition $\{\mathcal{D}\}$, then H is a hierarchy on $\mathcal{D}$ (see appendix A). For a digital image defined on a domain $\mathcal{D}$ of N pixels, H is finite and contains at most $2N - 1$ elements, a bound which is reached when H is a binary hierarchy.

Also consider for each region x of H the set $\Lambda(x)$ of the scales for which x belongs to P_λ :

$$\Lambda(x) \triangleq \{\lambda \mid x \in P_\lambda\}.$$

One can easily verify that the relation 1 implies that $\Lambda(x)$ is an *interval*, which can be written

$$\Lambda(x) = [\lambda^+(x), \lambda^-(x)[.$$

Definition 4 We call $\Lambda(x)$ the **interval of persistence** of set x , $\lambda^+(x)$ its **scale of appearance** and $\lambda^-(x)$ its **scale of disappearance** in P .

$\lambda^+(x)$ and $\lambda^-(x)$ are the equivalent of the so-called inner and outer scales of x in scale-space theory [31]. One then immediately verifies that a set $x \in H$ disappears when its father in H appears, that is

$$\forall x \in H \quad \lambda^-(x) = \lambda^+(\mathcal{F}(x)). \quad (2)$$

Hence, the couple $S = (H, \lambda^+)$, where H provides the sets which compose the partitions of P and λ^+ gives access to the range of scales over which they "live", captures the complete structure of P . We thus propose following definition :

Definition 5 [scale-sets representation] If $P = (P_\lambda)_{\lambda \in \mathbb{R}^+}$ is an unbiased multi-scale segmentation then the structure $\mathcal{S}(P) \triangleq (H, \lambda^+)$ where

$$\begin{aligned} H &\triangleq \bigcup_{\lambda \in \mathbb{R}^+} P_\lambda \\ \lambda^+ &\triangleq \begin{cases} H \mapsto \mathbb{R}^+ \\ x \mapsto \min\{\lambda \in \mathbb{R}^+ \mid x \in P_\lambda\} \end{cases} \end{aligned}$$

is called the **scale-sets representation** of P .

From equation 2, λ^+ is an *increasing function* with respect to the set-inclusion partial order:

$$x \subset y \quad \Rightarrow \quad \lambda^+(x) < \lambda^+(y).$$

In the domain of hierarchical classification, such a function is called an *index* on H and the couple (H, λ^+) is called an *indexed hierarchy* or a *stratified hierarchy*. In our context, H represents a family of regions of the image's domain and λ^+ maps them onto a 1D "scale" axis, hence the scale-sets terminology. The whole structure completely represents a volume of partitions, yet in an implicit manner. Any partition P_λ can then be immediately retrieved by *sectioning* the scale-sets structure:

Definition 6 [sections of a scale-sets] Let $\mathcal{S} = (H, \lambda^+)$ be a scale-sets on $\mathcal{D}$. Let B_λ be the family of boolean predicates on H :

$$B_\lambda(x) \text{ is true} \quad \Leftrightarrow \quad \lambda^+(x) \leq \lambda$$

then the coarsest cut of H whose elements verify B_λ is called $\mathcal{S}$ ' **section** of scale λ , or $\mathcal{S}$ ' **λ -section**. We denote it by $S_\lambda(\mathcal{S})$.

$S_\lambda(\mathcal{S})$ is the partition made up of the largest regions of H which have a scale of appearance lower than λ . Obviously, thanks to the increasingness of λ^+ , this partition exists and is unique for all $\lambda \geq 0$. Furthermore, if P is an unbiased multi-scale segmentation, then

$$\forall \lambda \in \mathbb{R}^+ \quad P_\lambda = S_\lambda(\mathcal{S}(P)).$$

Hence the family $(P_\lambda)_{\lambda \in \mathbb{R}^+}$ is the family of the sections of $\mathcal{S}(P)$.

In practice, a scale-sets $\mathcal{S}$ can be stored in a computer as a real-valued tree. It can be graphically represented by a dendrogram in which the y axis represents the scale axis. Each *horizontal cut* in this diagram then corresponds to a section of $\mathcal{S}$ (see figure 1).

2.3 Discussion

2.3.1 Related work

As regards the image segmentation problem, the idea of obtaining partitions as cuts in sets hierarchies goes back to the famous split and merge algorithm by Horowitz and Pavlidis [21]. However, the authors considered regular hierarchies (quad-trees) and only looked for a single cut, using a homogeneity predicate-based formulation of the partitioning problem.

Later, different authors have proposed to return a stack of monotone partitions - fine to coarse - as a segmentation result, also called a *pyramid of segmentations*, either based on a structural / graph-based approach [37, 24], on an energy minimization-based approach [27, 3, 13] or on a morphological approach [48, 46].

Other authors consider returning a tree of regions, or hierarchy, built with a classical region merging algorithm [43] or by a recursive divisive approach based on binary Markov Random Fields [42].

In all these works, different levels of details are proposed; however either the levels are not related to the values of a scale parameter, or the scale axis is sampled.

As we have seen, in order to obtain a complete representation of a sequence of monotone partitions with respect to a real parameter, one needs to consider an *implicit* representation of the partitions, as the family of the sections of a stratified hierarchy. A monotone mapping $\lambda \rightarrow P_\lambda$ is fully characterized by a set of critical events: namely the appearance of new regions at some specific scales, which are unions of some regions existing at lower scales. By nature, these events are discrete and in the case of an image defined on a discrete domain, they are in finite number. A complete representation can thus be obtained by making *explicit* these critical events and the scales at which they occur. It corresponds to a reverse point of view from the direct — pyramidal — point of view. The pyramidal approach amounts to setting some specific scales and asking: which sets are present at that scales? From the scale-sets point of view, the key-point is: at which scale does a specific set appear¹? Answering this question will be a key-point in the energy minimization-based approach to build scale-sets descriptions presented below.

¹This reversal of point of view can be made mathematically precise: if $\mathbf{P}$ is viewed as a multivalued function of λ returning *sets*, then $\Lambda(x)$ can be written $\Lambda(x) = \mathbf{P}^{-1}(x)$.

Figure 1: A scale-sets image description and some of its sections. This scale-sets was obtained by the scale climbing algorithm from a watershed over-segmentation of the image (see below).

2.3.2 Ultrametric distances and the region/contour duality

It is well known that the datum of a stratified hierarchy on $\mathcal{D}$ is equivalent to that of an **ultrametric distance** δ on $\mathcal{D}$, that is of a distance which verifies

$$\forall (x, y, z) \in \mathcal{D}^3 \quad \delta(x, z) \leq \max\{\delta(x, y), \delta(y, z)\}.$$

The closed balls of an ultrametric space $(\mathcal{D}, \delta)$ constitute a hierarchy H and the diameter of the balls is an index on H . Conversely, a couple (H, λ^+) induces an ultrametrics $\lambda_\star^-$ on $\mathcal{D}$: $\lambda_\star^-(x, y)$ is the smallest scale for which x and y are grouped, i.e. belong to a same region. Formally

$$\lambda_\star^-(x, y) \triangleq \lambda^+(\{x\} \vee \{y\})$$

where $\vee$ denotes the supremum operator in the sup-semi-lattice $(H, \subseteq)$. The elements of (H, λ^+) section of scale λ are the maximal balls of $(\mathcal{D}, \lambda_\star^-)$ which have a diameter lower than λ .

Now, if $\mathcal{D}$ is a domain of a topological space and if the regions of H are *connected* then the ultrametrics is obviously fully determined by its values for the couples of *neighbouring* points in $\mathcal{D}$. $\lambda_\star^-(x, y)$ then

represents the scale at which the *boundary* between x and y *disappears*. Thus, a multi-scale segmentation *into connected regions* can be equivalently represented by a **contour disappearance map** (see figure 2). This contour map has this particular property that thresholding it always leads to closed contours, hence to a partition of the image’s domain.

Figure 2: The contour disappearance map associated with the scale-sets of figure 1 (logarithmic scale).

For practical uses, this contour-oriented representation can be further simplified. Let $G = (\mathcal{D}, \mathcal{N})$ denote the region adjacency graph of P_0 , i.e. of the base of H . Value its edges by $\lambda_\star^-$ and consider a *minimum spanning tree (MST)* T of $(G, \lambda_\star^-)$. Removing the k edges of T which have a lower scale than λ then splits the tree into $k + 1$ connected components which correspond to the regions of the section of scale λ of $(H, \lambda^+)^2$. The scale-sets representation, by an indexed inclusion tree, and the indexed spanning tree representation are two dual representations of an unbiased multi-scale segmentation, in the sense of the duality between connected region-based and closed contour-based descriptions.

Fernand Meyer’s morphological multi-scale segmentation approach is based on the spanning tree representation [34, 35, 36]. Meyer proposes to start from a watershed transform of a gradient of the image, build the region adjacency graph of the catchment basins and value its edges by a measure of dissimilarity. A multi-scale segmentation is then given by the MST of this graph. Meyer proposes different measures of dissimilarity which amount to simulating different synchronous flooding processes of the gradient image. The scale parameter only depends on the contrast, the surface or the volume (contrast $\times$ surface) of the regions.

In [18] we proposed another way to obtain multi-scale segmentations from a dissimilarity-based grouping approach. We defined a *cocoon* of a valued graph as a connected set of nodes whose maximal internal dissimilarity is lower than the minimal dissimilarity with the exterior nodes. We proved that the set of the cocoons of a graph is a hierarchy and released an associated ultrametrics. Cocoons hierarchies are related to complete-linkage clustering while the MST of a graph can be computed by single-linkage clustering [29, 18].

However, dissimilarity-based approaches do not allow to introduce *geometrical criteria* in the segmentation process. In contrast, the multi-scale analyses which we propose here are based on optimizing two-term-based energies, one term being a geometrical regularization term. We now turn to the energy-minimization part of our approach.

²This comes from the fact that the MST of a graph (G, d) , where d is a dissimilarity function, is a graph-based representation of the maximal ultrametrics δ bounded above by d . When d is already an ultrametrics then $\delta = d$.

3 Scale-sets generation from a variational principle

The first part of the paper has discussed the general structure of unbiased multi-scale segmentations and the way they can be represented exactly, considering a scale-sets representation. We shall now show that one can build scale-sets whose sections approximate the solutions of a general energy minimization problem which embeds most optimization-based formulations of the segmentation problem (variational, Markovian, minimal encoding). The energies involved have two global terms whose relative weight is controlled by a real parameter which, as we shall see, generally behaves as a scale parameter.

3.1 Scale as an inherent parameter of approximate modelling problems

In a very classical way now, we put the single-scale segmentation problem as an optimal piece-wise image modelling problem [39, 15]. Let $\mathcal{I}$ be the set of the images one is interested in and $\mathcal{M}$ be a set of possible models of these images. Given an image $I \in \mathcal{I}$, the objective is to pick in $\mathcal{M}$ the "best model" of I . By a "comparison principle" [27], this can always be formalized as the problem of finding $M \in \mathcal{M}$ which minimizes a certain energy $E : \mathcal{I} \times \mathcal{M} \mapsto \mathbb{R}^+$. Such an approach contains three aspects: the definition of $\mathcal{M}$, the definition of E , and the optimization itself.

In deterministic approaches of segmentation, $\mathcal{M}$ is a set of piece-wise defined images: piece-wise constant or smooth [39], polynomial [30, 25]. . . In probabilistic approaches, $\mathcal{M}$ is a set of piece-wise defined stochastic processes whose realizations are images: piece-wise Gaussian processes. . . In what follows, we focus on deterministic models; however a strictly parallel reasoning can be followed for probabilistic models.

In a deterministic framework, a model of an image is also an image. Hence, one can at once choose a distance D between models and images and define the "best model" M of I as the one which minimizes $D(I, M)$. However, the solution for this type of problem is in general not satisfactory: most of the time the model set contains the image set, so that there is always a trivial solution $M = I$ which achieves $D(I, M) = 0$. For instance, if one looks for a piece-wise constant model of an image, then there is always an exact solution which consists in tessellating the image into as many regions as pixels. In practice, one looks for regions which correspond to coherent phenomena which are larger than a pixel. Hence, one does not look for the absolutely flat zones of an image but look for regions over which the image is *approximately* constant. A *precision of approximation* then needs to be introduced.

If I is an image and ε is a positive number, let us say that a model $M \in \mathcal{M}$ such that $D(I, M) \leq \varepsilon$ is an ε -approximation of I . For a given ε , there are a priori numerous ε -approximations of I . The way to choose one of them must then be principled, which can again be done through an energy-minimization approach: assume that we are capable of defining an energy $C : \mathcal{M} \mapsto \mathbb{R}^+$ which captures our preferences a priori for the different models; the problem of the ε -approximation of an image I can then be expressed as a constrained optimization problem which we denote by $\mathcal{P}_\varepsilon(I)$:

$$\text{Find } M \in \mathcal{M} \text{ which minimizes } C(M) \text{ subject to } D(I, M) \leq \varepsilon. \quad (3)$$

The energy C can be of different kinds. It can translate objective knowledge on the data or be based on relatively subjective criteria and aim to favour a specific interpretation. However, this constrained problem takes a particular meaning when C can be interpreted as a measure of the *complexity* of a model (objective or subjective). In this case, the problem $\mathcal{P}_\varepsilon(I)$ formalizes the idea that *between two models that equally fit the data, the simpler of the two should be preferred*. It thus can be viewed as an expression of Occam's Razor principle. This principle has a dual formulation : *between two equally complex models, the closer to the data of the two should be preferred*, which leads to a dual minimization problem $\mathcal{P}_\gamma^*(I)$:

$$\text{Find } M \in \mathcal{M} \text{ which minimizes } D(I, M) \text{ subject to } C(M) \leq \gamma. \quad (4)$$

These two problems are well known in information theory where they formalize lossy compression issues [49]. In this context, the energy C represents the communication rate or the number of bits needed to encode the model and D measures the distortion of the data caused by the compression. Given a certain quota of losses, the problem is to obtain the best compression ratio or dually, given a certain compression ratio, the problem is to minimize the losses. Classically, the minimization of a constrained

problem is obtained by minimizing its associated Lagrangian. The Lagrangians associated to $\mathcal{P}_\varepsilon(I)$ and $\mathcal{P}_\gamma^*(I)$ respectively read:

$$E_\mu(I, M) = C(M) + \mu D(I, M) \quad (5)$$

$$E_\lambda^*(I, M) = \lambda C(M) + D(I, M) \quad (6)$$

where μ and λ are the Lagrange multipliers.

Under large assumptions, one proves that there is a bijection between the solutions of $\mathcal{P}_\varepsilon(I)$ (resp. $\mathcal{P}_\gamma^*(I)$) and the minima of $E_\mu(I, M)$ (resp. $E_\lambda^*(I, M)$). Every (single) value of ε (resp. γ) corresponds to a particular value of μ (resp. λ) for which the minimum of the Lagrangian is identical to the solution of the problem under constraint.

With the lagrangian expressions, one finds again the usual energies used in image analysis. $E_\lambda^*(I, M)$ corresponds for instance to the general expression of a regularized ill-posed variational problem. D is then usually called a "goodness-of-fit" term and C a "regularization" term. For segmentation purposes, the archetype of this type of formulation is given by Mumford and Shah's functional [39]. Within the framework of Bayesian inference, $E_\lambda^*(I, M)$ can also be interpreted as the Gibbs potential associated to a conditional probability of the model given the data, or posterior probability [15]. In this case, $\exp(-D(I, M))$ corresponds to the probability of the data knowing the model, or likelihood, and $\exp(-\lambda C(M))$ corresponds to the prior probability of the model. Minimizing $E_\lambda^*(I, M)$ is equivalent to maximizing $\exp(-E_\lambda^*(I, M))$ which is proportional to the posterior probability of M .

Now, viewing an energy $E_\mu(I, M)$ as the lagrangian of a constrained minimization problem allows to interpret its minimization as the resolution of a problem of optimal ε -approximation of an image and thus puts in light the role of the parameter μ . μ controls the precision ε of the approximation found and thus can be thought of as a *scale parameter* along the image's values dimension: the distance between the optimal model and the image monotonously decreases as μ increases. Besides, the minima of $E_\mu(I, M)$ and $E_\lambda^*(I, M)$ are identical when $\lambda = 1/\mu$. Consequently, the families of solutions for the two dual constrained optimization problems coincide. Every precision ε corresponds to a bound of "complexity" $\gamma(\varepsilon)$. Looking for an ε -approximation of minimal complexity is equivalent to looking for the closest approximation whose complexity is lower than $\gamma(\varepsilon)$. The precision and the complexity of an optimal model thus vary altogether and the Lagrangian parameter is the bridge between them.

3.2 Statement of the problem and strategy adopted

We have described above two different notions of scale associated with piece-wise image description problems: from a structural/ensemblistic point of view, a "scale parameter" allows to browse a sequence of nested spatial partitions and thus controls the "structural" fineness of a description. From a modelling point of view, a "scale parameter" controls both the approximation's fidelity and the approximation's complexity. One would naturally like to see both notions of scale coincide, i.e. to see the modelling/optimization multi-scale framework match the structural/ensemblistic multi-scale framework.

However, in general, partitions minimizing an energy of the form 6 are not monotone, mainly because boundary displacements occur when λ increases. An example involving Mumford-Shah's functional for piecewise constant image approximation is provided in [16]. Furthermore, it is well known that even the simplest energy-minimization problems on discrete partitions are NP-hard (see e.g. [8] for a proof of it involving Potts energy).

As the minima of usual two term-based energies on partitions are not nested, we propose to enforce this structure and put the multi-scale segmentation problem as the one to **find a scale-sets description of an image whose sections are as close as possible to the partitions minimizing a two term-based energy of the form 6**. As we shall see, imposing a hierarchical structure to the solution will lead us to a fast approximation scheme of the solutions.

The way we address this problem is then driven by the structure of the scale-sets representation. Two distinct components are needed to fully determine a multi-scale segmentation: a hierarchy of regions H and a mapping λ^+ of these regions onto a one-dimensional axis, the 'scale' axis. H provides the sets but only induces a *partial order* on them through the inclusion relation. λ^+ then defines a *total order* on the structure, which implicitly defines a family of monotone cuts of H . We shall study the construction of those two components independently and in *reverse order*:

1) In section 3.3 we assume that we already know the structural part of the solution, i.e. the hierarchy H , and we prove that for an important class of two term-based energies, the family of H minimal cuts *is* a multi-scale analysis. Hence, for this class of energies, knowing H determines λ^+ . Furthermore, we show that λ^+ can be efficiently computed by a dynamic programming-based algorithm.

2) In section 3.4 we get back to the construction of the hierarchy itself. The results of section 3.3 lead us to a natural principle to build a scale-sets whose sections track the energy's minima, which is called the **scale climbing** principle. The method is parameter-free and leads to scale-sets which verify important properties.

3.3 Multi-scale minimal cuts of a hierarchy

This section develops the first part of the strategy described above. From there on, we shall restrict the type of energies considered to what we call **affine separable energies (ASE)**, the energies which can be written for a partition P

$$E_\lambda(P) = \sum_{R \in P} \lambda C(R) + D(R) \quad (7)$$

Most energies usually considered for image segmentation can be separated. Examples of it are the Mumford-Shah functional [39], Potts' priors for Markov Random Fields [15], and most description length criteria such as the one proposed by Leclerc [30]. From a probabilistic point of view, the separability of E_λ amounts to an independence assumption between the processes that generated the different regions of an image.

The datum of an ASE E_λ on the partitions of a domain $\mathcal{D}$ is equivalent to the datum of a couple (C, D) of energies on the *parts* of $\mathcal{D}$. We can thus also write $E_\lambda = (C, D)$.

3.3.1 Hierarchical minimization of separable energies

Forget for a while the scale parameter and simply consider a separable energy $E(P) = \sum_{R \in P} E(R)$. If H is a hierarchy then H 's cut which minimizes E can be easily computed by dynamic programming (DP):

$\forall x \in H$, let $C^*(H(x))$ be the minimal cut of H 's partial hierarchy rooted at x . Let $E^*(H(x))$ be the energy of this cut. As the union of two cuts of two disjoint hierarchies is a cut of their union, and as the energy E is separable, then $\forall x \in H$ the following Bellman's dynamic programming equations holds :

$$E^*(H(x)) = \min \left\{ E(x), \sum_{s \in \mathcal{S}(x)} E^*(H(s)) \right\} \quad (8)$$

$$C^*(H(x)) = \begin{cases} \{x\} & \text{if } E(x) \leq \sum_{s \in \mathcal{S}(x)} E^*(H(s)) \\ \bigcup_{s \in \mathcal{S}(x)} C^*(H(s)) & \text{otherwise} \end{cases} \quad (9)$$

One can thus optimize E starting from the leaves of H and applying equations 8 and 9 successively for all the nodes of H in hierarchical order. $C^*(H)$ and $E^*(H)$ then provide the minimal cut of H and its energy. The overall procedure has a linear complexity with respect to H 's size.

This procedure has been used for classification purposes in the CART algorithm [9] and for wavelet bases construction in the Best-basis algorithm [12]. In these applications, the hierarchies involved are quad-trees. Salembier and Garrido also used this minimal cut algorithm to build optimal image thumbnails in a rate/distorsion sense [43]. The cut is extracted from a hierarchy of regions obtained by classical region-merging technique which has no particular affinity with the minimization problem. We also employed this method in an image segmentation algorithm based on minimizing a Minimum Description Length (MDL) criterion in a hierarchy of cocoons [18, 17].

The dynamic programming procedure allows to extract a single optimal partition from a hierarchy. Given an affine separable energy E_λ , we shall now study the behavior of H 's minimal cut with respect to λ .

3.3.2 Multi-scale minimal cuts

Let H be a hierarchy and E_λ be an ASE. For each $\lambda \in \mathbb{R}^+$, we call λ -cut the cut of H which minimizes E_λ . We denote this cut by $C_\lambda^*(H)$. A region $x \in H$ which belongs to $C_\lambda^*(H)$ is called λ -optimal.

Examining the structure of the dynamic programming equation 8, one can easily verify that

Proposition 7 *A set $x \in H$ is λ -optimal if and only if the following two properties hold:*

- i) x is partially optimal, i.e. $\forall Y \in \mathbb{C}(H(x)) \quad E_\lambda(\{x\}) \leq E_\lambda(Y)$.*
- ii) x is maximal in H for the property i), that is no $y \in H$ such that $x \subset y$ is also partially optimal.*

In what follows $P_\lambda^*(H)$ will denote the set of all *partially* λ -optimal nodes of H , that is the set of the nodes which verify the property i) of proposition 7. Our approach then rests on the following theorem:

Theorem 8 [Multi-scale minimal cuts]

Let $E_\lambda = (C, D)$ be an affine separable energy. If either term C or term D is monotone with respect to the fineness partial order relation on the partitions of $\mathcal{D}$, then for any hierarchy H on $\mathcal{D}$, the sequence $(C_\lambda^(H))_{\lambda \in \mathbb{R}}$ of the minimal cuts of H for E_λ is a multi-scale segmentation.*

More precisely, if C is decreasing in $\mathbb{P}(\mathcal{D})$, that is if

$$\forall (P, Q) \in \mathbb{P}^2(\mathcal{D}) \quad Q \geq P \Rightarrow C(Q) \leq C(P) \quad (10)$$

or if D is increasing in $\mathbb{P}(\mathcal{D})$ then $(C_\lambda^(H))_{\lambda \in \mathbb{R}}$ is an increasing sequence of partitions. If C is increasing or D is decreasing then $(C_\lambda^*(H))_{\lambda \in \mathbb{R}}$ is a decreasing sequence.*

The proof is provided in appendix B.

In what follows, we will focus on energies which lead to increasing sequences of solutions, fine to coarse when λ increases, hence to ASes built either on a decreasing regularizer C or on an increasing goodness-of-fit term D . Such energies shall be called **multi-scale ASes (MASEs)**.

Note that the decreasingness of a separable energy in the partition lattice is equivalent to the *subadditivity* of the corresponding energy on the set lattice. More precisely, one easily shows that a separable energy $C(P) = \sum_{R \in P} C(R)$ verifies 10 if and only if

$$\forall (R, S) \in \mathcal{P}^2(\mathcal{D}) \quad R \cap S = \emptyset \Rightarrow C(R \cup S) \leq C(R) + C(S). \quad (11)$$

Usual regularizing energies are decreasing. Quantifying the complexity of a segmentation by its number of regions gives a decreasing regularizer. Also, summing up any positive quantity (length, curvature, ...) along the boundaries of the partition leads to a decreasing regularizer: deleting a boundary by merging some adjacent regions systematically reduces the energy. Indeed, the decreasingness condition matches Occam's idea that entities should not be multiplied without necessity. Starting from this idea, the regions of a solution should only be multiplied if it increases the model's goodness-of-fit. Hence among two ordered solutions one should a priori prefer the coarser of the two.

On the contrary, usual goodness-of-fit energies are increasing. Assume one models the image within a region by a certain parametric model (e.g. constant, polynomial) and measures the distance in the L_p norm. Indeed, the overall energy (the sum of the residuals of an L_p regression) is always larger when one fits a single model to the union of two regions rather than two separate models to each region. In other words, in the context of fixed order parametric modelling in the L_p norm, increasing the number of model pieces always improves the model's fidelity.

As a conclusion, the energies which naturally arise in image segmentation lead to multi-scale solutions in a hierarchical framework. We shall now proceed with the consequences of the theorem 8.

3.3.3 Appearance, disappearance and persistence of regions

Indeed, if $(C_\lambda^*(H))_{\lambda \in \mathbb{R}}$ is a multi-scale segmentation, it can be represented by a scale-sets structure. Moreover, this scale-sets can be computed exactly and efficiently by generalizing the above dynamic programming procedure. We shall first explain the relationship between partial optimality and maximality in H and the scales of appearance and disappearance of H 's sets.

For each region x of the hierarchy, let

$$\Lambda^*(x) \triangleq \{\lambda | x \in C_\lambda^*(H)\}$$

be the set of the scales for which x belongs to the λ -minimal cut of H . Also define

$$\Lambda_{\uparrow}^*(x) \triangleq \{\lambda | x \in P_{\lambda}^*(H)\}$$

which is the set of the scales for which x is partially λ -optimal.

When E_{λ} is a MASE, if a given region x is partially optimal for a scale λ then so is it for any larger scale $\lambda' > \lambda$. Thus, $\forall x \in H$, the set $\Lambda_{\uparrow}^*(x)$ is an interval of the form $[a, +\infty[$. Now, from the maximality condition ii) in proposition 7, x is globally optimal at scale λ if and only if no upper node on its branch to the top of the hierarchy is also partially λ -optimal. We thus define

$$\Lambda_{\downarrow}^*(x) \triangleq \bigcup_{y \in H, x \subset y} \Lambda_{\uparrow}^*(y)$$

which represents the set of the scales for which x *cannot* be maximal and thus cannot be globally optimal. As $\Lambda_{\downarrow}^*(x)$ is a union of intervals of the form $[a, +\infty[$, it is also an interval of this form.

Finally, the two conditions of global optimality of a node provided by proposition 7 can be summarized by

$$\Lambda^*(x) = \Lambda_{\uparrow}^*(x) \setminus \Lambda_{\downarrow}^*(x).$$

We thus conclude that

Proposition 9 *If H is a hierarchy and E_{λ} is a MASE then $\forall x \in H$ the set $\Lambda^*(x)$ is an interval of the form $\Lambda^*(x) = [\lambda^+(x), \lambda^-(x)[$, where*

$$\lambda^+(x) = \inf \Lambda_{\uparrow}^*(x) \tag{12}$$

$$\lambda^-(x) = \inf \Lambda_{\downarrow}^*(x) = \min_{y \in H, x \subset y} \lambda^+(y) \tag{13}$$

The rightmost part of equation 13 shows that, knowing $\lambda^+(x)$ for each element x of H , λ^- can be easily computed by a top-down traversal of H . We shall now explain how λ^+ can be computed efficiently by generalizing the dynamic programming method described above.

3.3.4 Functional dynamic programming and persistent hierarchies

In the single-scale DP method, each node $x \in H$ is attributed by a single energy value. We now consider attributing each node $x \in H$ by a *function* of λ which represents its energetic behaviour with respect to scale.

$\forall x \in H$, $E_{\lambda}(x)$ represents energy of x at scale λ . Now rewrite E_{λ} as a function of λ :

$$E_x : \lambda \rightarrow E_x(\lambda) = \lambda C(x) + D(x).$$

and call it the **self-energy** of x . E_x is an *affine* function of λ which has a positive slope and y -intercept. For each partial hierarchy $H(x)$, also rewrite the energy of its minimal cuts as a function of λ :

$$E_x^* : \lambda \rightarrow E_x^*(\lambda) = E_{\lambda}^*(H(x)).$$

and call it the **partial energy** of x .

Clearly, for any base node of H (whose partial hierarchy reduces to a single set) these two energy functions coincide, that is $\forall b \in \underline{H} : E_b^* = E_b$.

Now, as addition and infimum operations on functions are constructed from the point-wise operations, for any node $x \in H$, dynamic programming equation 8 can be rewritten as a functional equation:

$$E_x^* = \inf \{E_x, \sum_{s \in S(x)} E_s^*\} \tag{14}$$

We then have the

Proposition 10 [Functional dynamic programming] Let H be a hierarchy and $E_\lambda = (C, D)$ a multi-scale affine separable energy, then $\forall x \in H$

i) E_x^* is a piece-wise affine, non decreasing, continuous and concave function.

ii) $\forall \lambda \in \mathbb{R}^+$

$$E_x^*(\lambda) = \begin{cases} \sum_{s \in \mathcal{S}(x)} E_s^*(\lambda) & \text{if } \lambda < \lambda^+(x) \\ E_x(\lambda) & \text{otherwise} \end{cases} \quad (15)$$

iii) If C is strictly decreasing or D is strictly increasing then $\lambda^+(x)$ is real and is the unique solution of

$$E_x(\lambda) = \sum_{s \in \mathcal{S}(x)} E_s^*(\lambda).$$

The proof is provided in appendix C.

Figure 3: Computation of partial energy function E_c^* and of scale of appearance $\lambda^+(c)$ of a node c knowing E^* for its sons. (a) Case of a first level node with two sons. (b) General case.

The function E_x^* which provides the energies of the optimal cuts of $H(x)$ with respect to λ is thus a rather simple piece-wise affine function (see figure 3(b)). It can be explicitly stored for each node of H , typically by the list of the endpoints of its affine pieces. Then, for each node x , computing E_x^* only requires two operations on piece-wise affine functions, a sum and an intersection with an affine function, which both can be computed exactly (see figure 3). If the driving energy (C or D) is strictly monotone, then by the proposition 10iii), the intersection between $\sum_{s \in \mathcal{S}(x)} E_s^*$ and E_x is unique and provides $\lambda^+(x)$. Note that when the driving energy is (even casually) additive, the intersection may exist and be unique, but it

also may not exist, in which case $\lambda^+(x) = +\infty$, or span all over $\mathbb{R}$, in which case $\lambda^+(x) = -\infty$. Knowing $\lambda^+(x)$, the form of E_x^* is then provided using equation 15. This functional dynamic programming (FDP) method thus allows to compute λ^+ and E^* for each region of H by a single bottom-up traversal of H . Afterwards, $\lambda^-(x)$ can be computed by a top-down traversal of H using equation 13.

After implementing two traversals of the hierarchy, bottom-up then top-down, one gets three quantities on each node: its partial energy function and its scales of appearance and disappearance. In particular, the partial energy $E_{\hat{H}}^*$ of the top of the hierarchy gives the energy of each λ -minimal cut $C_\lambda^*(H)$ of H . We simply denote it by E^* . Let also $C^*(\lambda) = C(C_\lambda^*(H))$ denote the complexity of $C_\lambda^*(H)$, that is in a compression context the rate function, and $D^*(\lambda) = D(C_\lambda^*(H))$ its distance to the image, or distortion function. Knowing E^* , these functions are simply given by :

$$\begin{cases} C^*(\lambda) &= \frac{\partial E^*}{\partial \lambda}(\lambda) \\ D^*(\lambda) &= E^*(\lambda) - \lambda \frac{\partial E^*}{\partial \lambda}(\lambda). \end{cases}$$

$C^*(\lambda)$ decreases and $D^*(\lambda)$ increases, hence the operational rate/distorsion curve $C^*(D^*)$ decreases (see figure 4). This result, which holds in a particular operational context, matches the general rate/distorsion monotonous relationship discovered by Shannon [49].

Figure 4: Typical energy curves associated to the minimal cuts of a hierarchy.

After FDP, some nodes might have an empty range of persistence, whenever $\lambda^+(x) \leq \lambda^-(x)$, which means that they never enter an optimal cut of H . These non persistent nodes do not belong to the scale-sets which represent H 's minimal cuts for E_λ and may thus be deleted from the initial hierarchy. Hence,

Definition 11 [*Persistent hierarchy and scale-sets associated to (H, E_λ)*]

Let H be a hierarchy and E_λ be a MASE, then the hierarchy $H^* \subset H$ defined by

$$H^* \triangleq H \setminus \{x \in H \mid \lambda^+(x) \leq \lambda^-(x)\}$$

is called the *persistent hierarchy* associated with H and E_λ . (H^*, λ^+) is the scale-sets representation of the minimal cuts of H for E_λ .

Figure 5 illustrates the different steps of the computation of this scale-sets and figure 6 shows the effect of the scale-sets representation on the family of minimal cuts of a hierarchy.

Figure 5: Multi-scale optimization in a hierarchy. (a) An initial binary hierarchy H drawn with a vertical axis corresponding to the levels of the regions in the hierarchy. After a bottom-up traversal of H , each region is attributed a scale of appearance in an optimal solution for the multi-scale energy considered. (b) Represents hierarchy H with a vertical axis corresponding to the scale axis. In this representation, the non persistent nodes, which never appear in an optimal solution, are upper than some of their ancestors. A top-down traversal of H then allows to compute the scales of disappearance of the regions, which reveal these non persistent nodes. (c)-(d) The persistent hierarchy H^* obtained after removing the non persistent regions. Remark that deleting some intermediate stages leads to n-ary hierarchies. (d) Illustrates that the scale of appearance is an increasing criterion in H^* .

3.3.5 Discussion

Negative scales ? As one might have noticed, nothing above constraints the scale parameter to be *positive*. In particular, all the base regions a priori have a scale of appearance of $-\infty$, and the top of the hierarchy has a scale of disappearance of $+\infty$. These two cases are side effects due to the relativity of the optimality criterion: the appearance depends on the energy of the sub-structures of a region and the

Figure 6: The scale-sets representation maps the family of the minimal cuts of a hierarchy into a family of horizontal cuts, or sections. (a) a 30×30 image. (b) Persistent hierarchy indexed by the level of the nodes. (c) Minimal cuts of representation (b). (d) Partition obtained by sectioning the scale-sets at half the scale of appearance of the top of the hierarchy. (e) Persistent hierarchy indexed by the scale of appearance. C is the section corresponding to the partition in (d). (f) Minimal cuts of representation (e). The hierarchy was obtained by the scale-climbing algorithm using the Mumford-Shah functional (see below). The resulting scale-sets is clearly bimodal, revealing the rather binary structure of the image.

disappearance depends on the energy of super-structures which are both undefined at the limits of the perceptual domain. The infinite ranges of persistence of both the base and the top are thus meaningless.

Apart from the extrema, it can be noted that, if energy D also decreases (this may only happen casually), some regions which do not belong to the base might also get a negative scale of appearance. The idea of obtaining negative scales is somewhat unnatural: what does a negative regularization mean? If no regularization at all is needed to prefer a region described as a whole rather than described by its parts, it means that the “goodness-of-fit” energy D itself embeds a regularizing term. A negative scale of appearance can then be interpreted as an “anti-regularization” factor needed to compensate the internal regularization of D , in order to ... split the region! At extreme settings, if D is also perfectly decreasing then the top of H appears at a negative scale: the optimal cut immediately jumps to the top and the whole image is merged for all positive scales.

Conversely, it can be shown that if D increases then all the scales of appearance are positive (this appears as a consequence of proposition 14 below). Thus, “true” opponent energy schemes —involving

a decreasing prior and an increasing goodness-of-fit term— lead to intuitive solutions, belonging to the positive scales domain. Other reasons might be invoked to restrict the range of solutions to positive scales. For example, in minimum encoding frameworks negative scales correspond to negative code lengths, which are meaningless. In practice, we only used "true" or nearly "true" opponent energies and we restricted the final solution to the positive scales domain.

Multi-dimensional extension One can easily check that the optimization technique can be extended to multi-dimensional regularizers by considering a vectorial "scale-parameter" $\lambda \in \mathbb{R}^n$ and also a vector of regularizing energies $C : \mathcal{P}(X) \rightarrow \mathbb{R}^n$. If $\langle \cdot, \cdot \rangle$ denotes the inner product of $\mathbb{R}^n$, one can then consider an energy of the form $E_\lambda(R) = \langle \lambda, C(R) \rangle + D(R)$. If all the components of C are sub-additive then all the framework is still valid. The partial energy functions are then piece-wise affine hyperplanes of $\mathbb{R}^{n+1}$.

3.4 The scale climbing principle

The previous section has shown how *given a hierarchy of regions*, the full sequence of its optimal cuts for a multi-scale energy E_λ could be computed and represented exactly. The methodology is general-purpose: it can be used after any process which builds a hierarchy. One could for example use it on a quad-tree or on a hierarchy obtained by recursively splitting the image using Shi and Malik's normalized cut criterion [50]. However, the results obtained so far put forward a natural way to build a hierarchy whose minimal cuts seek the global minima of the energy considered.

As we have seen, given a hierarchy H and a multi-scale energy E_λ , H 's λ -cut climbs the hierarchy when λ increases. It jumps to a higher region x each time λ reaches a specific scale namely $\lambda^+(x)$. Increasing λ amounts to regularizing the solution by strengthening the constraint on the model's complexity or — by duality — by softening the constraint on the model's goodness-of-fit. Thus consider starting from the finest solution — the one which involves the slightest regularization ($\lambda = 0^+$) — and simulating a continuous increase of λ . For most classes of models, the 0^+ -regularization solution is straightforward, e.g. for a piece-wise constant model it is made up of the flat zones of the image. Call this initial partition C_0 . Now check among all the under-partitions of C_0 *which one would be reached first* during a continuous increase of λ . Assume it is reached at λ_1 and call it C_{λ_1} . Collect it in the scale-sets as its λ_1 -section and repeat the process: increase λ until a different partition is reached, say at λ_2 , collect it in the scale-sets, and so on. As we shall now explain, such a continuous simplification process can be handled exactly within our hierarchical framework.

3.4.1 "Pure" scale climbing

We first define what we call the "pure" scale climbing strategy to build H . Assume you want to build H from scratch by progressively adding sets to the solution. Let H_0 be the minimal hierarchy on $\mathcal{D}$, that is the one which does not contain any other set than the singletons and $\mathcal{D}$ itself. Let H_k be the hierarchy obtained after k sets have been added to H_0 . Let R_k be the set of the regions that can be added to H_k without breaking the hierarchical structure of the solution: $R_k = \{R \in \mathcal{P}(\mathcal{D}) | R \notin H_k \text{ and } H_k \cup \{R\} \text{ is a hierarchy}\}$. Indeed, at each step k , the scale of appearance of any $R \in R_k$ in a minimal cut of the hierarchy $H_k \cup \{R\}$ is well defined. The "pure" scale climbing algorithm is then

Definition 12 *Pure scale climbing*

*Starting from the empty hierarchy H_0 , recursively add to H_k the region R of R_k which would get the lowest scale of appearance in $H_k \cup \{R\}$. Stop when $R_k = \emptyset$. The final hierarchy is called the **scale climbing hierarchy (SCH)** associated with E_λ .*

As shows the next proposition, the pure scale climbing algorithm follows a sequence of *upward global minima* of E_λ when the scale increases:

Theorem 13 *Let H_k be the hierarchy obtained after k steps of the scale climbing strategy for the multi-scale energy E_λ and let H^* be the hierarchy obtained after the scale climbing procedure is completed. Let C_k^u be the coarsest cut of H_k which is not the trivial cut $\{\mathcal{D}\}$. Let $R^* = \operatorname{argmin}_{R \in R_k} \lambda^+(R)$ be the next region to be added to H_k by scale climbing and $S^* = \operatorname{argmin}_{R \in R_k \setminus R^*} \lambda^+(R)$ be the region of R_k which*

has the second lowest scale of appearance after R^* . Then, for all $\lambda \in [\lambda^+(R^*), \lambda^+(S^*)]$, H^* 's λ -minimal cut minimizes E_λ in the set of all the under-partitions of C_k^u .

Sketch of the proof: Obviously, the set R_k of the regions which can be added to the solution gets smaller and smaller, namely whatever region is added at each step: $R_{k+1} \subset R_k$. Furthermore, one shows that the scales of appearance of the sets which remain in R_{k+1} after a completion step always *increase*, that is $\forall R \in R_{k+1} : \lambda^+(R)$ in $H_{k+1} \cup \{R\}$ is always greater than $\lambda^+(R)$ in $H_k \cup \{R\}$ (see [16]). Hence, adding R^* to H_k ensures that it will persist in the final scale-sets at least over the range $[\lambda^+(R^*), \lambda^+(S^*)]$. Moreover, as adding sets of R_k to H_k allows to build any under-partition of C_k^u , the partition built is optimal among all the under-partitions of C_k^u . ■

Hence, the pure scale climbing strategy tracks a global minimum among all the under-partitions of the current solution along the scale dimension.

The point is then that adding the sets in the order of their scale of appearance always leads to adding some *supersets* of the sets previously added. Indeed, as the scales of appearance of the completion hypotheses always increase after a completion step, no reversal of order can occur. As a consequence, all sets of a pure scale climbing hierarchy are persistent. This bottom-up direction of construction matches the direction of the dynamic programming method to compute λ^+ , thus if the dynamic programming computation has already been made on the partial solution H_k , then the value of $\lambda^+(R)$ for any $R \in R_k$ can be obtained by a single dynamic programming step. The whole framework thus points out a privileged way to build the hierarchy, namely *the causal way*.

Please note that, whereas scale climbing is a discrete procedure, it properly simulates a *continuous* increase of λ . The point is that the procedure explicitly computes the scales at which a simplification of model would occur. It thus can be viewed as a *graduated non-convexity* [6] or a *continuation* minimization method [30] constrained to produce ordered partitions. Indeed, the "scale" parameter λ provides a natural embedding of the energy. The following proposition gives a rate/distortion interpretation of this continuation :

Proposition 14 *The scale climbing strategy amounts to choosing at each completion step the region $R \in R_k$ which minimizes*

$$-\frac{\Delta D}{\Delta C} = -\frac{D(\{R\}) - D(P)}{C(\{R\}) - C(P)}$$

where P is the coarsest minimal cut of the partial hierarchy $H_{k+1}(R)$ which is not $\{R\}$.

proof : $\lambda^+(R)$ represents the scale for which the energy of $\{R\}$ becomes equal to the energy of the coarsest minimal cut P of $H_{k+1}(R)$ which is not $\{R\}$. Hence:

$$\begin{aligned} \lambda^+(R) &= \lambda \text{ such that } \lambda C(\{R\}) + D(\{R\}) = \lambda C(P) + D(P) \\ &= -\frac{D(\{R\}) - D(P)}{C(\{R\}) - C(P)}. \end{aligned}$$

■

The expression $\frac{\Delta D}{\Delta C}$ is the discrete form of the derivative of the distortion D with respect to the rate C . The scale climbing strategy can thus be interpreted as a steepest descent strategy along the operational rate/distortion curve of the problem.

Moreover, the scale climbing algorithm produces hierarchies which are robust to linear transforms of the energies :

Theorem 15 [Linear invariance of SCH]

If (H^, λ^+) is the SCH associated with the energy $E_\lambda = (C, D)$, then the SCH associated with $E'_\lambda = (\mu C, \nu D)$ is $(H^*, \frac{\nu}{\mu} \lambda^+)$.*

Sketch of the proof: The proof can be made by induction on the dynamic programming step of the computation of λ^+ , on the basis that if $f_{(a,b)}(x) = ax + b$ is an affine function and $I_x(f, g)$ denotes the intersection abscissa of two affine functions f and g then

$$I_x(f_{(\mu a, \nu b)}, f_{(\mu a', \nu b')}) = \frac{\nu}{\mu} I_x(f_{(a,b)}, f_{(a', b')}).$$

This relation remains valid for sum and infimum operations on piece-wise affine functions. Hence, as all scales of appearance are multiplied by ν/μ , the scale climbing order is unchanged by some linear transforms of the energies. ■

This invariance property has very important consequences.

The *structure* of a SCH is invariant by any linear transform on $\mathbf{C}$ or $\mathbf{D}$. Adjusting the relative weight of the initial energies just stretches the scale axis of the final scale-sets proportionally. The energies used can thus be stated up to a multiplicative term and the right balance between complexity and goodness-of-fit of a model can always be set *a posteriori*. This allows us to achieve our initial goal, which was to extricate the scale parameter from the low level segmentation stage. It becomes a real "potentiometer" which allows to explore the scale dimension.

This linear invariance also allows for the transfer of invariance properties from the energies to the final description. For instance, if the goodness-of-fit term is only scaled when the image values are scaled then so is the SCH. For example, one can easily verify that the solutions obtained with the Mumford and Shah's piece-wise constant model are invariant to some linear transforms of the image values and some homotheties of the image.

3.4.2 Binary scale climbing

Of course, the pure scale climbing strategy is computationally intractable: if $|\mathcal{D}| = N$ then the initial set R_0 of possible completions has almost 2^N elements. In order to get a practical algorithm, we thus restrict the search space to a local search space. As before, we denote by C_k^u the coarsest cut of H_k which is not the whole domain. Instead of considering all the connected supersets of the regions of C_k^u , we only consider the supersets which can be obtained by merging pairs of adjacent regions of C_k^u . The algorithm thus obtained works as follows:

Definition 16 [*Binary scale climbing*] Let E_λ be a multi-scale energy and P_0 a fine segmentation of $\mathcal{D}$. Set $P \leftarrow P_0$ and $H \leftarrow P_0$. While $|P| \neq 1$ merge the pair $\{R, S\}$ of adjacent regions of P whose scale of appearance $\lambda^+(R \cup S)$ is minimal and add $R \cup S$ to H . The final hierarchy H is called the **Binary Scale Climbing Hierarchy (BSCH)** associated with E_λ and P_0 .

The construction of a BSCH is parameter-free, given an over-segmentation and a couple of antagonist energies on the regions of an image. It is based on a classical priority queue region merging algorithm which can be efficiently implemented using a region adjacency graph and a heap structure to manage the queue of merging hypotheses [28, 20]. The bottom-up dynamic programming optimization is realized simultaneously to the construction of the hierarchy. Each time a pair of regions is merged, a DP step allows to compute the scale of appearance of the new region with its neighbors. A top-down propagation in H finally provides both the scales of disappearance of the regions and the persistent hierarchy. Note that the initial binary hierarchy then becomes n-ary. We show in [16] that the whole procedure can be implemented with a worst case complexity in $\mathcal{O}(N^2 \log N)$, where N is the size of the initial over-segmentation. For usual images, building a BSCH takes almost linear time.

While the search space is reduced, the major properties of the "pure" scale climbing strategy remain valid. The global minimum is only tracked locally, among the immediate under-partitions of the current solution. However, the local step is still chosen by steepest descent along the rate/distorsion curve and the final structure is also linearly invariant. Furthermore, the persistent hierarchies obtained are almost binary hierarchies (see below), which indicates that most of the time the "pure" scale climbing completion step would certainly have been chosen among the immediate supersets of the current upper sets, like in binary scale climbing.

4 Experimental results

We shall now consider some experimental results.

We first consider Mumford and Shah's classical piece-wise constant approximation [39]. The discrete

version of this model assigns following energies to a region R :

$$\begin{aligned} C(R) &= |\delta R| \\ D(R) &= \sum_{x \in R} \|I(x) - \overline{I_R}\|^2. \end{aligned} \quad (16)$$

C is the length of the boundary of R , and D is the sum of the quadratic deviations from the mean value of the image within the region. This second term is valid for any number of channels in the image.

Remark that a natural scale arises after the computation of a scale-sets, namely the scale of appearance of the top of the hierarchy, $\lambda^+(\mathcal{D})$ which represents the minimum regularization strength needed to get the image modeled by a single region. This scale depends on the image content, contrast and size (recall that with the Mumford-Shah functional, the scales of appearance are linearly covariant with respect to the image dynamics and size). Hence sectioning a scale-sets at some scales defined *relatively* to $\lambda^+(\mathcal{D})$ gives an adaptive analysis.

Figure 7 shows sequences of sections at scales $\lambda^+(\mathcal{D})/2^k$, $k = 1 \dots 8$, for the scale-sets obtained on some images. We call such a sequence the *dichotomy* of a scale-sets. Figure 7 pictures the models obtained, that is the piece-wise constant descriptions. Grouping started from the image pixels using Mumford-Shah's energy; hence the algorithm is absolutely parameter-free. With our implementation, the average computation time is 42 seconds on a 1,2GHz computer for a 256×256 image. Of course, starting from an initial over-segmentation of the image rather than at the pixel level greatly reduces computation time. With the same settings than above, starting from a watershed over-segmentation, which is a secure over-segmentation of an image [5], reduces the computation time to 2 to 3 seconds (the average size of the initial catchment basins ranging from 15 to 20 pixels).

Figure 8 pictures the scale-sets corresponding to the images of figure 7. The scale axis is represented with a logarithmic scale, which is the natural representation for intensive units (it allows to map ratios into differences). As $\log(x) < 0$ when $x < 1$, an origin of the representation has to be chosen in order to get a representation belonging to $\mathbb{R}^+$. Again, a mapping relative to $\lambda^+(\mathcal{D})$ is natural:

$$\lambda \rightarrow \begin{cases} \log\left(\frac{\lambda p}{\lambda^+(\mathcal{D})}\right) & \text{if } \lambda \geq \lambda^+(\mathcal{D})/p \\ 0 & \text{otherwise} \end{cases} \quad (17)$$

where p is a parameter called the "precision" of the representation. In the examples, $p = 2^8$, hence matching the dichotomies in figure 7. The same mapping was used for the contour-oriented representation (ultrametrics) of figure 9 (null log-scales are mapped to white and the log-scale of appearance of the whole domain is mapped to black).

Figure 7 clearly illustrates that meaningful objects appear at different scales in an image. The MR image of a head in figure 7(a) is a good example of it. Large scale values allow to separate the whole head from its background. Decreasing the scale parameter then progressively adds details to the description: the brain and the skull are first described as wholes and then internal details appear. Note that the scale-sets 8(a) is clearly bimodal, revealing the object/background structure of the image. The right hand side of the hierarchy corresponds to the head while the left hand side corresponds to the background. The partial hierarchy which corresponds to the background is characteristic of an unstructured region: no large persistent region, no particular stage during the grouping. Note that the contours corresponding to a scale-sets are based on global information and not on local intensity. Hence, the poorly contrasted contour of the top of the head in figure 7(a) gets the same "saliency" than the lower contour of the head which is very contrasted: both are equally meaningful because they participate in the boundary of the same globally coherent zone.

Figures 10 and 11 show results obtained with a more sophisticated geometrical model than the simple length-based energy of Mumford and Shah.

Let $L = (m_1, m_2, \dots, m_k)$ be a sequence of k 2D points defining a closed 2D polygonal line. If $v_i = m_{i+1} - m_i$ denotes the displacement vector associated with the segment $[i, i+1]$, we say that the quantity

$$S(L) = \sum_{i=1}^k |\widehat{v_i, v_{i+1}}|$$

Figure 7: Dichotomies of some scale-sets obtained by scale climbing from the image pixels using Mumford-Shah's piece-wise constant model. k is the order of the dichotomy (see text). The bottom row represents the original images.

where $\widehat{a, b}$ denotes the angle made up by vectors a and b , is the **concavity energy** of polyline L .

The name "concavity" comes from the fact that this energy has a value of 2π for all convex polygons and increases as the total curvature of the concavities of a polygon increases. To measure the concavity of a discrete region, we first perform a slight polygonal approximation of its boundary with a classical algorithm in order to suppress the pixelic artefacts.

Figure 8: scale-sets corresponding to the images of figure 7. The scale axis is logarithmic with a precision p of 256 (see equation 17).

Figure 9: Contour disparition maps (ultrametrics) corresponding to the images in figure 7. Logarithmic scale axis of precision $p = 256$.

In examples 10 and 11, the initial over-partition was obtained by a watershed algorithm based on the magnitude of a Canny-Deriche edge detector tuned to $\alpha = 2$ [11]. As before, we considered piece-wise constant image modelling and employed the L_2 norm as a fidelity criterion (equ. 16). The difference with the Mumford-Shah functional is that the regularizing energy is a concavity energy.

We observed on several experiments that this energy produces more meaningful regions than Mumford-Shah's energy (the scale-sets of figure 1 was also computed with this energy). A point which might appear surprising for a piece-wise constant model is that this energy allows to find the right delineation for very textured units at large scales, such as the village in figure 10. The point is that energy minimization

criteria are relative criteria: even if a very textured region is absolutely not constant, if it is different enough from its environment then for a low precision, the best piece-wise constant model is a model which separates the region from its background.

Figure 10: An aerial image.

Figure 11: Group photograph.

5 Conclusion and future work

A general methodology for multi-scale region-oriented image analysis has been presented. It can be viewed as a synthesis between structural —hierarchical grouping— approaches and energy minimization based approaches or as a region-oriented counterpart to the scale-space theory. The main result of the paper is that usual piece-wise modelling problems expressed in a variational way intrinsically contain a free parameter which behaves as a scale parameter. We then developed a methodology which provides an approximation of the full structure of the solutions with respect to scale, which we called a scale-sets description in reference to the scale-space theory. The scale-climbing algorithm is absolutely parameter-free given a multi-scale energy and a base segmentation which can be the pixel-wise tessellation of the image. The fundamental property of the structure obtained is its linear scale invariance which ensures that scale tuning is totally deferred to a subsequent stage.

The final hierarchical representations allow the dynamical browsing of the complete set of the multi-scale solutions. The useful scale of description for a specific application can thus be tuned *a posteriori*, either interactively or automatically (e.g. by actively searching for specific objects). Also, partitions made up of objects living at different scales may be composed, which correspond to non horizontal cuts in the hierarchy.

While we only present segmentation results on natural images, the multi-scale modelling framework is not dedicated to a specific kind of image, nor is it to the partitioning of image data. In [51], the scale-climbing algorithm method is used to perform multi-scale piece-wise plane descriptions of range images obtained from a classical stereovision algorithm.

Different applications of the proposed framework have already been developed by different authors. [10] proposed a stereovision algorithm based on matching the regions of two scale-sets descriptions of a stereo pair. The robustness of the algorithm comes from the fact that regions of different levels of the hierarchies can be matched, hence avoiding the classical problem of differences in level of segmentation in region-oriented stereo-matching. We observed that, due to its robustness, the scale-climbing algorithm often segments in much the same way the matching regions of a stereo-pair, however not necessarily at the same scale. [52] also used the contour-based representation of a scale-sets (the ultrametrics) to register a cadaster graph onto an aerial image segmented with the scale-climbing algorithm.

Salembier and Garrido and co-workers have developed a number of image processing techniques based on hierarchies of regions. They have shown how such structures could be used for a large number of processing goals such as object detection and recognition, visual browsing and region-oriented image compression, filtering, image data-base indexing and similarity based retrieval. All these techniques can naturally rely on scale climbing hierarchies. We refer the reader to the authors' papers [45, 14, 44, 43] for detailed descriptions. Part of our future work will be centered on the development of such applications using SCHs.

[35, 32, 55] have proposed interactive segmentation systems based on the multi-scale approach of Meyer [34]. Again, these interactive tools could be based on SCHs.

In the scale-space theory, the evolution diagram (when scale increases) of the features for which the analysis is causal (e.g. the image maxima) is known as the scale-space *fingerprint* of an image [22]. The scale-sets itself, that is the stratified hierarchy structure, is the fingerprint associated with a multi-scale segmentation. Like a scale-space fingerprint, it could for example be used in content-based image retrieval problems.

Obviously, the persistence of the regions of a SCH is a very meaningful information. The range of scales over which a region lives is a measure of its stability, in the sense of Leclerc [30]. It represents the efforts needed to merge it with its neighborhood and thus quantifies the region's saliency. Starting from a point in the image and climbing the associated branch of the hierarchy, a chain of nested regions is browsed. When the "growing" region reaches a maximal coherent part of an image, a peak of persistence can be observed, which generally means that the region delineated corresponds to a meaningful object of the scene. The issue of exploiting the persistences (and also energetical variations related to merging operations) in order to extract salient objects of an image will be part of our future work.

In conclusion, we would like to get back to some methodological aspects of the approach.

First and back with David Marr, we'd like to stress the "duality" between the "representation and the processing of information" [33, first sentences, p.3]. The present paper clearly shows that the investigation of how the solutions to a problem can be represented can lead to important hints on the way to solve it. Also, our methodology relies on addressing an intractable problem into a subspace of the space of its solutions (the cuts of a hierarchy rather than the whole partition lattice), in which it can be exactly solved. As we have seen, finding exact resolution methods in structured subspaces can then, in a kind of feedback loop, lead to privileged ways to build the subspaces and thus to fine approximation methods. Finally, we would like to go back over the functional dynamic programming principle. It amounts to handling the explicit form of a parametric energy towards its parameters and considering the inverse problem of finding the values of the parameters for which an element of the search space is optimal. When such an approach is successful, it turns a parameter into a real "potentiometer" which allows flexible algorithm tuning. It also allows to quantify the stability of a solution, a point which is certainly

as important as obtaining one. This “inverse parametric” approach is certainly worth being investigated for other computer vision or pattern recognition problems.

Appendices

A Definitions and notations

This section recalls some general definitions and puts the global notations of the paper.

- Given a set $\mathcal{D}$, a subset $S \subset \mathcal{P}(\mathcal{D})$, where $\mathcal{P}(\mathcal{D})$ denotes the set of the parts of $\mathcal{D}$, is called a set system on $\mathcal{D}$.
- A set system made up of non overlapping sets covering $\mathcal{D}$ is a **partition** of $\mathcal{D}$. We denote by $\mathbb{P}(\mathcal{D})$ the set of the partitions of $\mathcal{D}$. Recall that a partition P is *finer* than a partition Q , $P \leq Q$, if and only if $\forall p \in P \exists q \in Q : p \subseteq q$. The inclusion relation, $\subseteq$, is a partial order on $\mathcal{P}(\mathcal{D})$ and the fineness relation, $\leq$, is a partial order on $\mathbb{P}(\mathcal{D})$. $\subseteq$ and $\leq$ induce complete lattice structures on $\mathcal{P}(\mathcal{D})$ and $\mathbb{P}(\mathcal{D})$.
- A set system H of $\mathcal{D}$ is a **hierarchy** on $\mathcal{D}$ if and only if
 - $\emptyset \notin H$
 - $\mathcal{D} \in H$
 - $\forall x \in \mathcal{D} : \{x\} \in H$
 - $\forall (X, Y) \in H^2 : X \cap Y \in \{\emptyset, X, Y\}$
(X and Y are either disjoint or nested)

The set $\mathcal{D}$ is called the top of H and we denote it by $\hat{H}$. The subset $\{\{x\}, x \in \mathcal{D}\}$ of H is called its base—or bottom—and we denote it by $\underline{H}$. For image segmentation problems, $\mathcal{D}$ is the domain of the image under scope, and H 's base represents the finest partition considered: either the absolute over-partition of $\mathcal{D}$, into individual pixels, or a coarser one, such as its tessellation into flat zones [48] or into the catchment basins of an image gradient [4].

- A hierarchy H can be represented as a rooted tree $T = (H, S)$ (see figure 12). The nodes of T represent the regions of H (the root is the top of H), and the edges S of T represent the covering relation between the elements of H , i.e. $(x, y) \in S \Leftrightarrow x \subset y$ and $\nexists z \in H | x \subset z \subset y$. For any edge $(x, y) \in S$, y is called x 's father and x is one of y 's sons. We denote the father of a node x by $\mathcal{F}(x)$ and the set of the sons of a set y by $\mathcal{S}(y)$.

Figure 12: A hierarchy H represented as a tree. $H(x)$ is the partial hierarchy rooted at x . C_1 and C_2 are two cuts, resp. made up of circled and diamond-shaped nodes.

- $\forall x \in H$, the subset of H defined by $H(x) \triangleq \{y \in H, y \subseteq x\}$ is a hierarchy on x which we call H 's **partial hierarchy** induced by x . The associated tree is the complete subtree rooted at x (see figure 12).
- A **cut** C of a hierarchy H is a subset of H which intersects any path from the base to the top of H exactly once. Equivalently, the cuts of a hierarchy H on a set $\mathcal{D}$ are the partitions of $\mathcal{D}$ that can be obtained by picking sets in H . A cut C can be graphically represented by a curve which divides the tree into two sets of nodes: the ones which remain connected to the top and the ones which remain connected to the bottom (see figure 12). The elements of the cut itself are the upper nodes remaining connected to the bottom. We denote by $\mathbb{C}(H)$ the set of all the possible cuts of a hierarchy H . Please note that:
 - Two cuts C_1 and C_2 of a hierarchy are in general *unordered* (C_1 is neither coarser nor finer than C_2). For example, the two cuts in figure 12 are unordered, which can be checked because their representative curves intersect.

ii) The number of possible cuts of a hierarchy is in general very large. One easily shows that if H is a balanced binary tree built on a set of size N then $|\mathbb{C}(H)| > \sqrt{2}^N$ (we provide a precise approximation of this cardinal in [16]). It is interesting to compare this bound to the number of partitions of a planar domain in *connected regions*, which is lower than 4^N [53].

B Proof of theorem 8

We focus on the case of a decreasing energy $\mathbb{C}$. The increasing case is symmetrical and the case of a monotony condition on $\mathbb{D}$ is obtained by the fact that the minimum of $\lambda\mathbb{C} + \mathbb{D}$ is the same as the minimum of $\mathbb{C} + \frac{1}{\lambda}\mathbb{D}$. The inversion in λ explains the change of direction of monotony in the solutions.

Let $\lambda_1 \in \mathbb{R}$. Expanding the partial λ_1 -optimality property (proposition 7-i) for all x in $P_{\lambda_1}^*(H)$ then gives:

$$\forall x \in P_{\lambda_1}^*(H) \quad \forall Y \in \mathbb{C}(H(x)) :$$

$$\begin{aligned} \lambda_1 \mathbb{C}(x) + \mathbb{D}(x) &\leq \sum_{y \in Y} \lambda_1 \mathbb{C}(y) + \mathbb{D}(y) \\ \Leftrightarrow \lambda_1 \left(\mathbb{C}(x) - \sum_{y \in Y} \mathbb{C}(y) \right) &\leq \sum_{y \in Y} \mathbb{D}(y) - \mathbb{D}(x) \end{aligned}$$

As $\mathbb{C}$ is decreasing, and by definition Y is a partition of x , we get

$$\Delta \mathbb{C} = \mathbb{C}(x) - \sum_{y \in Y} \mathbb{C}(y) = \mathbb{C}\left(\bigcup_{y \in Y} y\right) - \sum_{y \in Y} \mathbb{C}(y).$$

Hence, if $\mathbb{C}$ is a sub-additive function on sets then $\Delta \mathbb{C} \leq 0$, which implies that $\forall \lambda_2 \in \mathbb{R}$

$$\lambda_2 \geq \lambda_1 \Rightarrow \lambda_2 \Delta \mathbb{C} \leq \lambda_1 \Delta \mathbb{C} \leq \sum_{y \in Y} \mathbb{D}(y) - \mathbb{D}(x)$$

which expresses the partial λ_2 -optimality of x . We thus obtain a “partial causality” property:

$$\forall (\lambda_1, \lambda_2) \in \mathbb{R}^2 \quad \lambda_2 \geq \lambda_1 \Rightarrow (\forall x \in H \quad x \in P_{\lambda_1}^*(H) \Rightarrow x \in P_{\lambda_2}^*(H)) \quad (18)$$

Let’s now examine the conditions of global optimality of a node x . The condition i) of proposition 7 can be rewritten

$$\forall \lambda_1 \in \mathbb{R} \quad (\forall x \in H \quad x \in C_{\lambda_1}^*(H) \Rightarrow x \in P_{\lambda_1}^*(H)).$$

So, using the “partial causality” equation 18 we obtain $\forall (\lambda_1, \lambda_2) \in \mathbb{R}^2$

$$\lambda_2 \geq \lambda_1 \Rightarrow (\forall x \in H \quad x \in C_{\lambda_1}^*(H) \Rightarrow x \in P_{\lambda_2}^*(H)).$$

Now, if $x \in P_{\lambda_2}^*(H)$ then no partially λ_2 -optimal node y such that $y \subset x$ may be maximal (verify proposition 7ii)). As y is λ_2 -optimal if and only if it is a maximal partially λ_2 -optimal node, we conclude that $\forall (\lambda_1, \lambda_2) \in \mathbb{R}^2$

$$\begin{aligned} \lambda_2 \geq \lambda_1 &\Rightarrow \forall x \in C_{\lambda_1}^*(H) \quad \forall y \in C_{\lambda_2}^*(H) \quad x \cap y \neq \emptyset \Rightarrow x \subseteq y \\ \Leftrightarrow \lambda_2 \geq \lambda_1 &\Rightarrow C_{\lambda_2}^*(H) \geq C_{\lambda_1}^*(H). \end{aligned}$$

C Proof of proposition 10

i) The proof is made by induction on the dynamic programming step in H .

Indeed, $\forall x \in H$, its self-energy $E_x = \lambda\mathbb{C}(x) + \mathbb{D}(x)$ is affine with $\mathbb{C}(x) \geq 0$ and is thus a piece-wise affine, non decreasing, continuous and concave function (NDCC). In particular, as $\forall b \in \underline{H} \quad E_b^* = E_b$, the partial energies of the base nodes are NDCC. Now take $x \in H$ and assume that $\forall s \in \mathcal{S}(x)$, E_s^* is NDCC. As finite sums and infima of NDCC functions are also NDCC, then according to equation 14, E_x^* is also

NDCC.

ii) Let $E_{S(x)}^* = \sum_{s \in S(x)} E_s^*$. By the dynamic programming principle the range $\Lambda_{\dagger}^*(x)$ of partial λ -optimality of a node x reads

$$\Lambda_{\dagger}^*(x) = \{\lambda | E_x(\lambda) \leq E_{S(x)}^*(\lambda)\}.$$

According to proposition 9, this range is given by

$$\Lambda_{\dagger}^*(x) = [\lambda^+(x), +\infty[.$$

Thus,

$$E_x(\lambda) \leq E_{S(x)}^*(\lambda) \Leftrightarrow \lambda \geq \lambda^+(x)$$

which means that E_x^* has the piece-wise form given by equation 15.

iii) As all the functions involved are continuous

$$\lambda^+(x) = \min\{\lambda | E_x(\lambda) = E_{S(x)}^*(\lambda)\}.$$

Thus consider $\Delta = E_x - E_{S(x)}^*$ and look for its zero crossings. $\forall \lambda \in \mathbb{R}$, $E_{S(x)}^*(\lambda)$ represents the energy of a cut $P(\lambda)$ of $H(x)$ which is a strict over-partition of $\{x\}$, hence

$$E_{S(x)}^*(\lambda) = \lambda \sum_{R \in P(\lambda)} C(R) + \sum_{R \in P(\lambda)} D(R)$$

and $\forall \lambda \in \mathbb{R}$

$$\frac{\partial E_{S(x)}^*}{\partial \lambda}(\lambda) = \sum_{R \in P(\lambda)} C(R).$$

We thus get

$$\frac{\partial \Delta}{\partial \lambda}(\lambda) = \frac{\partial E_x}{\partial \lambda}(\lambda) - \frac{\partial E_{S(x)}^*}{\partial \lambda}(\lambda) = C(x) - \sum_{R \in P(\lambda)} C(R).$$

So, if C is strictly subadditive, then $\forall \lambda \in \mathbb{R} : \frac{\partial \Delta}{\partial \lambda}(\lambda) < 0$. Δ is thus strictly decreasing and crosses zero for a unique $\lambda \in \mathbb{R}$ which is $\lambda^+(x)$.

References

- [1] Luis Alvarez, Frédéric Guichard, Pierre-Louis Lions, and Jean-Michel Morel. Axioms and Fundamental Equations of Image Processing. *Arch. Rational Mech. Anal.*, 123:199–257, 1993.
- [2] J. Badaud, Witkin A. P., Baudin M., and Duda R. O. Uniqueness of the gaussian kernel for scale-space filtering. *IEEE trans. Pattern Anal. Mach. Intell.*, 8:26–33, 1986.
- [3] C. Ballester, V. Caselles, and M. González. Affine Invariant Segmentation by Variational Method. In *Proc. of 9th RFIA, Paris*, pages 379–390, January 1994.
- [4] S. Beucher and C. Lantuejoul. Use of watersheds in contour detection. In *Proceedings of Int. Workshop on Image Processing, Real-time Edge and Motion Detection/Estimation, Rennes (France)*, september 1979.
- [5] S. Beucher and Fernand Meyer. The morphological approach to segmentation : the watershed transformation. In Dougherty, editor, *Mathematical Morphology in Image Processing*, chapter 12, pages 433–481. Marcel Dekker, 1993.
- [6] Andrew Blake and A. Zisserman. *Visual reconstruction*. MIT Press, Cambridge, MA, 1987.
- [7] Joseph H. Bosworth and Acton Scott T. Morphological scale-space in image processing. *Digital Signal Processing*, 13(3), 2003.

- [8] Yuri Boykov, Olga Veksler, and Ramin Zabih. Fast approximate energy minimization via graph cuts. *IEEE trans. on Pattern Analysis and Machine Intelligence*, 2001.
- [9] L Breiman, J H Friedman, R A Olshen, and C J Stone. *Classification and regression trees*. Wadsworth and Brooks, 1984.
- [10] Nesrine Chehata, Franck Jung, Marc Pierrot-Deseilligny, and George Stamon. A region based approach for 3D-roof reconstruction form HR satellite stereo pairs. In *DICTA 03, Vol II, Sydney, Australia*, pages 889–898, December 2003.
- [11] Rachid Deriche. Using Canny’s criteria to derive a recursively implemented optimal edge detector. *International Journal of Computer Vision*, 1(2):167–187, 1987.
- [12] David Donoho. CART and best-ortho-basis selection : a connection. *Annals of statistics*, 25:1870–1911, 1997.
- [13] F. Fuchs and H. Le Men. Detecting planar patches in urban scenes. In *Proceedings of SPIE, Visual Information Processing VIII, Orlando, USA*, volume 3716, pages 167–176, avril 1999.
- [14] Luis Garrido and Philippe Salembier. Representing and retrieving regions using binary partition trees. In *IEEE Int. Conference on Image Processing ICIP’99, Kobe, Japan*, pages 25–28, October 1999.
- [15] Stewart Geman and D. Geman. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6(6):721–741, 1984.
- [16] Laurent Guigues. *Modèles multi-échelles pour la segmentation d’images (Multi-scale models for image segmentation)*. PhD thesis, Univ. de Cergy-Pontoise, France, 2003.
- [17] Laurent Guigues, Hervé Le Men, and Jean-Pierre Cocquerez. Segmentation d’image par minimisation d’un critère mdl dans une pyramide de segmentations. In *actes du XVIII^{ème} colloque GRETSI sur le Traitement des Images et du Signal, Toulouse, France*, 2001.
- [18] Laurent Guigues, Hervé Le Men, and Jean-Pierre Cocquerez. The hierarchy of the cocoons of a graph and its application to image segmentation. *Pattern Recognition Letters*, 24(3):1024–1066, May 2003.
- [19] Haralick and Shapiro. Glossary of computer vision terms. *Pattern Recognition*, 24(1):69–93, 1991.
- [20] K. Haris, S. N. Estradiadis, N. Maglaveras, and A. K. Katsaggelos. Hybrid image segmentation using watersheds and fast region merging. *IEEE Trans. on Image Processing*, 7(12):1684–1699, 1998.
- [21] Steven L. Horowitz and Theodosios Pavlidis. Picture Segmentation by a Tree Traversal Algorithm. *J. of the Assoc. for Comp. Mach.*, 23(2):368–388, April 1976.
- [22] Paul T. Jackay, W. W. Boles, and M. Deriche. Morphological fingerprints and their use in object recognition in range images. In *Proc. of IEEE Conf. on Acoustics, Speech and Signal Processing, Adelaide, Australia*, pages V5–V8, 1994.
- [23] Paul T. Jackay and M. Deriche. Scale-space properties of the multiscale morphological dilation-erosion. *IEEE trans. Pattern Anal. Mach. Intell.*, 9:532–550, 1987.
- [24] Jean-Michel Jolion and Annick Montanvert. The adaptive pyramid, a framework for 2d image analysis. *CVGIP: Image Understanding*, 55(3):339–348, May 1992.
- [25] Tapas Kanungo, Byron Dom, Wayne Niblack, David Steele, and Jacob Sheinvald. MDL-Based Multi-Band Image Segmentation Using a Fast Region Merging Scheme. Technical Report RJ 9960 (87919), IBM Research Division, May 1995.
- [26] J. J. Koenderink. The structure of images. *Biol. Cybern.*, 50:363–370, 1984.

- [27] 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.
- [28] T. Kurita. An Efficient Agglomerative Clustering Algorithm Using a Heap. *Pattern Recognition*, 24(3):205–209, 1991.
- [29] Lance and Williams. A general theory of classificatory sorting strategies. *The Computer Journal*, 9:373–380, 1967.
- [30] Yvan G. Leclerc. Constructing Simple Stable Descriptions for Image Partitioning. *International Journal of Computer Vision*, 3(1):73–102, 1989.
- [31] Tony Lindeberg. *Scale-space Theory in Computer Vision*. Kluwer Academic, 1994.
- [32] B Marcotegui, Francisca Zanoguera, P Correia, R Rosa, Ferrand Marqués, R Mech, and M Wollborn. A video object generation tool allowing friendly user interaction. In *Proc. of ICIP'99, Kobe (Japan)*, October 99.
- [33] David Marr. *Vision : a computational investigation into the human representation and processing of visual information*. W.H. Freeman and Company, NY 29-61, 1982.
- [34] Fernand Meyer. Graph Based Morphological Segmentation. In *Proc. 2nd IAPR TC-15 workshop on Graph based Representations, Haendorf, Austria*, pages 51–60, 1999.
- [35] Fernand Meyer. Morphological multiscale and interactive segmentation. In *Proc. of IEEE-EURASIP Workshop on Nonlinear Signal and Image Processing, Antalya, Turkey*, June 1999.
- [36] Fernand Meyer. An overview of morphological segmentation. *Int. Journal of Pattern Recognition and Artificial Intelligence*, 15(7):1089–1118, 2001.
- [37] Annick Montanvert, Peter Meer, and Azriel Rosenfeld. Hierarchical image analysis using irregular tessellations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(4):307–316, April 1991.
- [38] Jean-Michel Morel and Sergio Solimini. *Variational methods in image segmentation*. Progress in Nonlinear Differential Equations and Their Applications, vol. 14, Birkhäuser, Boston, 1995.
- [39] David Mumford and Jiambo Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Comm. on Pure and Applied Math.*, 17(4):577–685, 1989.
- [40] K.-R. Park and C.-N. Lee. Scale-space using mathematical morphology. *IEEE trans. Pattern Anal. Mach. Intell.*, 8:1121–1126, 1996.
- [41] P. Perona and Jitendra Malik. Scale space and edge detection using anisotropic diffusion. *IEEE trans. on PAMI*, 12(7):629–639, July 1990.
- [42] G. Poggi and A. R. P. Ragozini. Image segmentation by a tree-structured markov random field. *IEEE Signal Processing Letters*, 6:155–157, July 1999.
- [43] Philippe Salembier and Luis Garrido. Binary Partition Tree as an Efficient Representation for Image Processing, Segmentation, and Information Retrieval. *IEEE Trans. on Image Processing*, 9(4):561–576, April 2000.
- [44] Philippe Salembier and Luis Garrido. Visual segments tree creation for MPEG-7 description schemes. In *IEEE Int. Conference on Multimedia and Exposition ICME'2000, NY city, NY USA*, 2000.
- [45] Philippe Salembier, Albert Oliveras, and Luis Garrido. Anti-extensive Connected Operators for Image and Sequence Processing. *IEEE Trans. on Image Processing*, 7(4):555–570, April 1998.
- [46] Philippe Salembier and Jean Serra. Flat Zones Filtering, Connected Operators, and Filters by Reconstruction. *IEEE trans. on Image Processing*, 4(8):1153–1160, August 1995.

- [47] Guillermo Sapiro and A. Tannenbaum. Affine Invariant Scale Space. *Int. J. of Computer Vision*, 11:24–44, 1993.
- [48] Jean Serra and Philippe Salembier. Connected operators and pyramids. In *SPIE conf. on Image Algebra and Mathematical Morphology, San Diego, CA*, volume 2030, pages 65–76, 1993.
- [49] Claude Shannon. Coding theorems for a discrete source with a fidelity criterion. *IRE Natl. Conv. Record*, Part 4:142–163, 1959.
- [50] Jiambo Shi and Jitendra Malik. Normalized Cuts and Image Segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905, August 2000.
- [51] Franck Taillardier, Laurent Guigues, and Rachid Deriche. A framework for constrained multi-scale range image segmentation. In *Proc. of IEEE Int. Conf. on Image Processing (ICIP'03), Barcelona, Spain*, September 2003.
- [52] Roger Trias-Sanz and Marc Pierrot-Deseilligny. An edge-based method for registering a graph onto an image with application to cadastre registration. In *Proc. of the 2004 International Conference on Pattern Recognition (ICPR 2004), Cambridge, UK*, 2004.
- [53] Jia-Ping Wang. Stochastic Relaxation on Partitions With Connected Components and its Application to Image Segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(6):619–636, June 1998.
- [54] Andrew W. Witkin. Scale space filtering. In *Proc. of 8th Int. Joint Conf. on Artificial Intell., Karlsruhe, West Germany*, pages 1019–1021, 1983.
- [55] Francisca Zanoguera, B Marcotegui, and Fernand Meyer. A toolbox for interactive segmentation based on nested partitions. In *proc. of ICIP'99, Kobe (Japan)*, October 1999.