



HAL
open science

Extracting meaningful curves from images

Frédéric Cao, Pablo Musé, Frédéric Sur

► **To cite this version:**

Frédéric Cao, Pablo Musé, Frédéric Sur. Extracting meaningful curves from images. [Research Report] RR-5067, INRIA. 2003. inria-00071517

HAL Id: inria-00071517

<https://inria.hal.science/inria-00071517>

Submitted on 23 May 2006

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.

Extracting meaningful curves from images

Frédéric Cao Pablo Musé, and Frédéric Sur

N°5067

Décembre 2003

THÈME 3

 *Rapport
de recherche*

Extracting meaningful curves from images

Frédéric Cao* Pablo Musé[†], and Frédéric Sur[†]

Thème 3 — Interaction homme-machine,
images, données, connaissances
Projet Vista

Rapport de recherche n°5067 — Décembre 2003 — 38 pages

Abstract: Since the beginning, Mathematical Morphology has proposed to extract shapes from images as connected components of level sets. These methods have proved very efficient in shape recognition and shape analysis. In this paper, we present an improved method to select the most meaningful level lines (boundaries of level sets) from an image. This extraction can be based on statistical arguments, leading to a parameter free algorithm. It permits to roughly extract all pieces of level lines of an image, that coincide with pieces of edges. By this method, the number of encoded level lines is reduced by a factor 100, without any loss of shape contents. In contrast to edge detections algorithm or snakes methods, such a level lines selection method delivers accurate shape elements, without user parameter: no smoothing involved and selection parameters can be computed by Helmholtz Principle.

Key-words: Edge detection, Mathematical Morphology, Topographic map, level lines, Helmholtz Principle, Gestalt theory

(Résumé : tsvp)

* IRISA/INRIA, Campus universitaire de Beaulieu, 35042, Rennes Cedex, FRANCE, Frederic.Cao@irisa.fr

[†] École Normale Supérieure de Cachan, 61 avenue du Président Wilson, 94235 Cachan Cedex, France, {sur},{muse}@cmla.ens-cachan.fr

Extraction de courbes significatives dans les images

Résumé : Dès sa naissance, la morphologie mathématique a proposé d'extraire des formes dans les images à partir des composantes connexes des ensembles de niveau. L'efficacité de ces méthodes a été démontrée en reconnaissance et analyse de formes. Dans cet article, nous présentons une méthode de sélection des lignes de niveau (les frontières des ensembles de niveau) les plus significatives dans une image. Cette extraction est basée sur des arguments statistiques, menant à une détection sans paramètre. Elle permet de détecter tous les morceaux de lignes de niveau qui coïncident avec les contours de l'image, en diminuant le nombre de lignes codables d'un facteur 100, sans pour autant réduire le contenu perceptuel des formes. Contrairement aux méthodes de détection de contours ou de contours actifs traditionnelles, cette sélection de lignes de niveau aboutit à des éléments de formes précis, sans paramètre à régler par l'utilisateur. En outre, aucun lissage n'est nécessaire et tous les autres paramètres sont fixés par le principe de Helmholtz.

Mots-clé : Détection de contours, Morphologie Mathématique, carte topographique, lignes de niveau, principe de Helmholtz, théorie de la Gestalt

Contents

1	Introduction	3
2	Meaningful boundaries	7
2.1	Helmholtz Principle	7
2.2	Contrasted boundaries	7
2.3	Maximal boundaries	8
2.4	Discussion on the definition of meaningful contrasted boundaries	9
2.4.1	Interpretation of the number of false alarms	9
2.4.2	Length distribution and meaningfulness	11
2.4.3	Cleaning-up meaningful boundaries	11
2.4.4	Geometrical information reduction	13
2.4.5	Meaningful boundaries vs. Haralick's detector	15
2.4.6	Color edges	15
3	Image quantization and gradient norm distribution	16
4	Local boundary detection	19
4.1	Algorithm	19
4.2	Experiments on locally contrasted boundaries	20
5	Meaningful boundaries or snakes?	25
5.1	Definition of local regularity	26
5.2	Meaningful contrasted and smooth boundary	27
5.3	Comparison with active contours	29
5.4	Experiments on smooth meaningful boundaries	30
6	Meaningful edges	31
6.1	Edges as pieces of level lines	31
6.2	Experiments on edges	32
7	Conclusion	32
A	Appendix: Numerical estimation of the Hausdorff dimension of a curve	37

1 Introduction

Natural images are very complex, and despite the progress of modern computers, we cannot handle the huge amount of information they contain. Thus, the idea of Marr and Hildreth [29] that edges provide a good summary of images is still vivid. Since their seminal works, efforts have been carried on local methods. Marr defined edges as zero-crossings of the laplacian [28], and Haralick [19] proposed a more correct definition which is equivalent to the zero-crossings of $D^2u(Du, Du)$ where

Du and D^2u are respectively the gradient and the second derivative of the image. In his famous paper [2], Canny gives a filter that tries to optimize the edge localization, but which is equivalent to Haralick's. Although they are technically sound, local methods also have an immediate drawback: while edges are usually thought about as curves, these methods detect sets of points with an orientation (*edgels*) that have to be connected afterward. Moreover, they require different thresholds since contrast has no absolute meaning. In addition, they are sensitive to noise (since they use derivatives of the image) and can only be considered through a multiscale process. The choice of these thresholds depends on the observed image, and is not that easy. It is also known that edge is not a completely local concept and that it does not rely entirely on contrast. Indeed, following Gestalt Theory [20, 42], shapes (and thus edges) result from the collaboration of a small set of perceptual laws (called "partial gestalts" by Desolneux, Moisan and Morel [14]), and contrast is only one of them. Among others, we can cite alignments, symmetry, convexity, closedness, good continuation, etc...

Other theories, related to edge detection, explicitly use good continuation, which means in this case regularity of curves. The most famous is certainly the theory of active contours (or snakes) [21], where optimal boundaries result from a compromise between their intrinsic regularity and the extrinsic value of the image contrast along the active contours. The main weaknesses of this theory are the number of parameters and the sensitivity to an initial guess. More recent methods propose to initiate the detection with many contours, most of which will hopefully disappear [9]. But again, there is no measure on the certainty of the remaining detected contours.

The Mathematical Morphology school proposed an alternative to the local approaches above. Following morphologists, the image information is completely contained in a family of binary images that are obtained by thresholding the images at given values [30, 39]. This is equivalent to considering level sets; the level set of u at the value λ is

$$\chi_\lambda(u) = \{x \in \mathbb{R}^2, \quad u(x) \geq \lambda\}. \quad (1)$$

Obviously, if we only consider a coarsely quantized set of different gray levels, information is lost, especially in textures. Nevertheless, it is worth noting how large shapes are already present with as few as 5 or 6 levels. As soon remarked by Serra [39], no information is lost at all, since we can reconstruct an image from the whole family of its level sets, by

$$u(x) = \sup\{\lambda \in \mathbb{R}, \quad x \in \chi_\lambda(u)\}.$$

Thus, the level sets do not only give a convenient way to extract information, they provide a complete representation of images. Alternative complete representations are, for instance, Fourier or wavelets decomposition [27]. But while these last are very adequate for image compression (they are used in the JPEG 2000 standard), they are not very well adapted in shape analysis, since their basic elements have no immediate perceptual interpretation. On the contrary, morphologists soon remarked that boundaries of level sets fits parts of objects boundaries very well. They call level lines the topological boundaries of connected components of level sets, and topographic map of an image, the collection of all its level lines. The topographic map also gives a complete representation of an image and enjoys several important advantages [6]:

- It is invariant with respect to contrast change. It is not invariant to illumination change, since in this case, the image is really different, although it represents the same scene. However, many level lines still are locally the same.
- It is not as local as sets of edges, since level lines are Jordan curves that are either closed or meet the image borders. (This property requires that the image has bounded variations [16]).
- It is a hierarchical representation: since level sets are ordered by the inclusion relation (and so are there connected components), the topographic map may be embedded in a tree structure.
- But most important regarding the main subject of this paper, object contours locally coincide with level lines very well. Basically, level lines are everywhere normal to the gradient as edges. On the other hand, level lines are accurate at occlusions. Whereas, edges detectors usually fail near T-junctions (and additional treatments are necessary), there are several level lines at a junction. The order of the multiple junction coincide with the number of level lines [4]. We shall go back to this in Sect.2.4.5.




Figure 1: Level lines and T-junction. Depending on the grey level configuration between shapes and background, level lines may follow or not (as on the figure) the objects boundary. In any case, junctions appear where two level lines separate. Here, there are two kinds of level lines: the occluded circle and the shape composed of the union of the circle and the square. The square itself may be retrieved by difference.

The level sets representation has recently been used, with success, for image simplification and segmentation. In particular, it was shown that it allowed to define multiscale representation of images [31, 37, 38], while avoiding the main drawbacks of linear scale space theory [23, 43], namely an oversmoothing of contours.

We are convinced that level lines may directly give usable curves for any shape recognition algorithm. The main drawback of the topographic map representation is its lack of compactness. First, since it is complete, it contains all the texture information. The level lines in textures are usually very complicated, and are not always useful for blind shape recognition. (The opposite may be true,

for instance for very accurate image registration). Moreover, because of noise and interpolation, many level lines may follow a single contour. Thus, it is useful, for practical computational reasons, to select only the most meaningful level lines.

Recently, Desolneux et al. proposed a parameterless algorithm to detect contrasted level lines (called *meaningful boundaries*) in grey level images [13]. Their method, which needs no parameter tuning, relies on a perceptual principle called Helmholtz Principle. Experimentally, meaningful boundaries are often very close to minimizers of any reasonable snake energy [15]. This adequation of meaningful boundaries and snakes is a bit paradoxical since, unlike snakes, no local regularity is imposed on meaningful boundaries.

However, the algorithm of Desolneux et al. raises several questions and objections. First, because of image quantization, some edges are missing (lots of them in some low contrasted images). Second, it uses a global information on contrast (the histogram). This yields an over-detection in regions with important contrast and a sub-detection in low contrasted regions (it is the so-called blue sky effect). Third, regularity of edges is not used for the detection.

In this paper, we discuss and answer these three objections, with a significant improvement. Our conclusions are the following: quantization noise can be removed with a very slight smoothing (for instance a gaussian with a standard deviation equal to 0.5). We cannot talk of multiscale edges as in Marr-Hildreth theory, since, in practice, this smoothing is invisible, and this value does not depend on the experiment. We also propose a method considering contrast in a more local way. If we use more local contrast information, we can remove edges in texture. Whether this is a nice thing or not depends on the application: for very accurate registration, texture-edges can be useful, while they must be useless for shape recognition. (For texture recognition, harmonic analysis methods are certainly more efficient.) Last, we introduce a local and stable measure of regularity of a curve and use it for smooth edges detection. As already noticed in [3], regularity is often sufficient to detect some very meaningful edges. Nevertheless, general belief is that both regularity and contrast are useful for edge detection. We experimentally check that contrast and regularity are often very redundant. This redundancy is used to make the detection even more robust, but does not change the results of contrast based detection alone. We are also able to tune automatically the relative weight of regularity and contrast, which is a recurrent question in active contours theory.

The plan is as follows. In Sect. 2, we recall the bases of Helmholtz Principle, the definition of meaningful boundaries of Desolneux, Moisan and Morel. We will justify and discuss this definition, which was not explicitly made in [13]. Errors due to the quantization of the contrast are corrected in Sect. 3. In Sect. 4, we describe a procedure that automatically handles local contrast variations. In Sect. 5, we explain how both contrast and regularity criteria can naturally be mixed in a probabilistic setting by introducing a measure of regularity on random level lines in Sect. 5.1. In Sect. 6, we describe a procedure to cut boundaries into their most meaningful parts, before concluding in Sect. 7.

2 Meaningful boundaries

2.1 Helmholtz Principle

Helmholtz Principle is a perceptual principle asserting that conspicuous structures may be viewed as exceptions to randomness. The unexpected configurations we must be interested in, are given by the perceptual laws of Gestalt Theory [20, 42], as alignments, closedness of sets, parallelism etc... Given a class of geometrical events, we shall say that it is ε -meaningful, if, in average, less than ε of these events are observed in an image of white noise. Thus, features are detected *a contrario*, since we consider them as meaningful if they have a very low probability to occur by chance in an image of noise of the same size. In a well sampled white noise image, two objects that are at a distance lower than the Nyquist distance (that is to say 2 pixels) cannot be considered as independent. In what follows, we use the expression “independent points” to qualify points at Nyquist distance (and thus independent in white noise). It can be shown that Helmholtz Principle allows to detect some elementary and important geometrical structures (see [15] and references therein). All the parameters can be reduced to the choice of ε , which is the expected number of detections in a white noise image (which we call “false detections”). In practice, ε can always be taken equal to 1 since the number of objects we want to detect is quite large. Moreover, meaningful perceptual objects correspond to very small values of ε (about 10^{-10}). As a consequence, the detection is in practice parameter free. We refer the reader to [13, 14] for further details.

2.2 Contrasted boundaries

In order to illustrate Helmholtz principle, we recall here the definition of meaningful boundaries given in [13]. It will be also useful since we will discuss this definition in the next sections. Let u be a grey level image. We search the level lines of u along which contrast assumes unexpectedly high values. We compare these values with the empirical histogram of contrast. By translation invariance, we assume that the contrast values are identically distributed following the law of the random variable X defined by

$$\forall \mu > 0, \quad P(X > \mu) = \frac{\#\{x, |Du| > \mu\}}{\#\{|Du| > 0\}}. \quad (2)$$

We denote by $H_c(\mu)$ this empirical probability. In [13], Desolneux, Moisan and Morel proposed the following definition.

Definition 1 ([13]) *Let N_l be the number of level lines of u . A level line C is an ε -meaningful boundary if*

$$NFA(C) \equiv N_l H_c(\min_{x \in C} |Du(x)|)^{1/2} < \varepsilon, \quad (3)$$

where l is the length of C . This number is called number of false alarms (NFA) of C .

In (3), the NFA is the product of the number of level lines and the probability that a random curve Γ containing $\frac{1}{2}$ independent samples has its contrast larger than $\min_{x \in C} |Du(x)|$ everywhere, when we assume that the contrast values on curve samples are mutually independent. If the NFA is very small, this means that this assumption is certainly not valid, leading to an *a contrario* detection. Notice also that meaningful boundaries are invariant with respect to affine contrast changes. In what follows, this detection model will be referred to as the MB model (meaningful boundaries model).

2.3 Maximal boundaries

Since level lines are nested, meaningful boundaries can also be embedded in a tree structure. As remarked by Desolneux, Moisan and Morel [12], meaningful boundaries usually appear in parallel groups. This is due to the fact that images are interpolated and that edges are thick (at least two or three pixels), when images are suitably sampled with respect to Shannon-Nyquist theory. These boundaries are redundant, and in applications, it may be useful to eliminate some of them, for time and memory saving. The previous authors use the notion of *maximal monotone section* in a level line tree, as introduced by Monasse [32]. They consider parts of branches of the tree of meaningful boundaries, such that any node has only one son and the grey level is monotone in this part. Such a branch part is called a monotone section. It is maximal if it is not strictly contained in another monotone section.

Definition 2 ([13]) *We say that a meaningful boundary is maximal meaningful if it has a minimal NFA in a maximal monotone section.*

Figure 2 illustrates that the loss of information of maximal meaningful boundaries is negligible compared to the gain of information compactness.




Figure 2: Maximal meaningful boundaries. 1. Original image, 83,759 level lines 2. All meaningful boundaries: 11,505 detections. 3. Maximal meaningful boundaries. Only 883 boundaries remain, while the visual loss is very weak.

Since meaningful boundaries inherit the tree structure of the topographic map, they can be used to reconstruct an image, thus defining an image operator, see Fig. 3. It is a connected operator as defined by Salembier and Serra [38] (but it is not a filter by reconstruction). It is neither a contrast invariant operator, since it explicitly uses the gradient value (it only commutes with affine global contrast change), nor an idempotent operator.

As remarked by Salembier [37], an operator pruning the topographic maps preserves edges very well. Contrary to local operators as for instance grain filter [41], the meaningful boundary reconstruction does not simply remove leaves of the tree but also inner nodes corresponding to possibly large (but low contrasted) level lines.



Figure 3: Original image on the left (99,829 level lines). Right: reconstruction from the 429 maximal meaningful boundaries. The gray level may not be really significant since, on edges, the maximal meaningful level line has an intermediate level between both sides of the edge. It would be more perceptually adequate to set the gray level to the brighter or darker meaningful level line. Nevertheless, for contrast independent shape recognition purposes, we do not use the grey level value, but only the geometry of level lines. The most important is that we preserved the main shapes, while removing the textures.

2.4 Discussion on the definition of meaningful contrasted boundaries

2.4.1 Interpretation of the number of false alarms

In this section, we give a precise interpretation of Def. 1, which was not explicit in [13]. Let us first recall the following classical lemma.

Lemma 1 *Let X be a random variable and $H(x) = P(X \geq x)$. Then for all $t \in [0, 1]$,*

$$P(H(X) < t) \leq t.$$

Assume that X is a random variable described by the inverse repartition function $H(\mu) = P(X \geq \mu)$. Assume that u is a random image such that the values $|Du|$ are independent with the same law as X . Let now E be a set of random curves (C_i) in u such that $\#E$ (the cardinality of E) is independent from each C_i . For each i , we note $\mu = \min_{x \in C} |Du(x)|$ and L_i the (random) length of

C_i . We assume that L_i is independent from the pixels crossed by C .

We say that C_i is ε -meaningful if

$$NFA(C_i) = \#E \cdot H(\mu_i)^{L_i} < \varepsilon.$$

Proposition 1 *The expected number of ε -meaningful curves in a random set E of random curves is smaller than ε .*

Proof. Let us also denote by X_i the binary random variable equal to 1 if C_i is meaningful and to 0 else. Let also $N = \#E$.

$$\mathbb{E} \left(\sum_{i=1}^N X_i \right) = \mathbb{E} \left(\mathbb{E} \left(\sum_{i=1}^N X_i | N \right) \right).$$

We have assumed that N is independent from the curves. Thus, conditionally to $N = n$, the law of $\sum_{i=1}^N X_i$ is the law of $\sum_{i=1}^n Y_i$, where Y_i is a binary variable equal to 1 is $nH(\mu_i)^{L_i} < \varepsilon$ and 0 else. By linearity of expectation,

$$\mathbb{E} \left(\sum_{i=1}^N X_i | N = n \right) = \mathbb{E} \left(\sum_{i=1}^n Y_i \right) = \sum_{i=1}^n \mathbb{E}(Y_i).$$

Since Y_i is a Bernoulli variable, $\mathbb{E}(Y_i) = P(Y_i = 1) = P(nH(\mu_i)^{L_i} < \varepsilon) = \sum_{l=0}^{\infty} P(nH(\mu_i)^{L_i} | L_i = l)P(L_i = l)$. Again, we have assumed that L_i is independent of the gradient distribution in the image. Thus conditionally to $L_i = l$, the law of $nH(\mu_i)^{L_i}$ is the law of $nH(\mu_i)^l$. Let us finally denote by $(\alpha_1, \dots, \alpha_l)$ the l (independent) values of $|Du|$ along C_i . We have

$$\begin{aligned} P(nH(\mu_i)^l < \varepsilon) &= P \left(H \left(\min_{1 \leq k \leq l} \alpha_k \right) < \left(\frac{\varepsilon}{n} \right)^{1/l} \right) \\ &= P \left(\max_{1 \leq k \leq l} H(\alpha_k) < \left(\frac{\varepsilon}{n} \right)^{1/l} \right) \text{ since } H \text{ is nonincreasing} \\ &= \prod_{k=1}^l P \left(H(\alpha_k) < \left(\frac{\varepsilon}{n} \right)^{1/l} \right) \text{ by independence} \\ &\leq \frac{\varepsilon}{n} \text{ from Lemma 1.} \end{aligned}$$

This term does not depend upon l , thus

$$\sum_{l=0}^{\infty} P(nH(\mu_i)^{L_i} < \varepsilon | L_i = l)P(L_i = l) \leq \frac{\varepsilon}{n} \sum_{l=0}^{\infty} P(L_i = l) = \frac{\varepsilon}{n}.$$

Hence,

$$\mathbb{E} \left(\sum_{i=1}^N X_i | N = n \right) \leq \varepsilon.$$

This finally implies $\mathbb{E} \left(\sum_{i=1}^N X_i \right) \leq \varepsilon$, which exactly means that the expected number of meaningful curves is less than ε . \square

2.4.2 Length distribution and meaningfulness

Let us consider an ε -meaningful boundary C containing L_C points, and $\mu = \min_{x \in C} |Du(x)|$. The computation of the NFA of C involves the probability that, on a curve contained in noise, the gradient is everywhere greater than μ , knowing that the length of the curve is $L = L_C$. The contrario model gives the law of $\min |Du|$ on a curve, conditionally to its length. Precisely, the independence assumption yields

$$P(\min_{x \in C} |Du(x)| \geq \mu | L = l) = H_c(\mu)^l.$$

The MB model implicitly considers that curves are equally probable, independently from their length. Since all the curves we consider are level lines, this is obviously not true in white noise. Indeed, there are much more short level lines than long ones. Thus, the fact that we observe a very long level line is, in itself, a reason to detect it. This is already suggested in the definition of the NFAs, since as soon as $H_c(\mu) < 1$, a curve becomes meaningful if it is long enough. Experiments confirm this fact: more than 50% of the level lines with a length greater than 1000 are meaningful. The minimal value of the gradient along very long meaningful boundaries is generally very small: less than 1. Are these detections false detections? As far as we observed in real images, the fact that a level line is very long is never casual. The cause of this rare event may be that the curve is an edge or that a smooth gradient generated by a light source. Level lines due to light sources are not edges properly speaking: they do not follow objects boundaries, contrast is weak, and they are not as smooth as most contours (they have no shape in the usual meaning). But, in terms of image analysis, they obviously have an interpretation. From this view point, the proportion of long meaningful level lines is certainly still underestimated. Indeed, if we define meaningful boundaries as large deviations from curves generated by noise, the probability of interest should rather be given by the joint law of the length and the minimal gradient, that is to say

$$P(L \geq l, \min_{x \in C} |Du(x)| \geq \mu).$$

If we could estimate this probability, we would be able to define NFAs which would take values much smaller than the current ones. In particular all very long curves would certainly be meaningful. The interpretation is that, in a white noise image of the same size as the real image, we do not observe level lines of such a length. However, the fact that these lines have a perceptual interpretation does not imply that they correspond to the usual concept of shape (e.g. illumination gradient). In addition, the joint law of the minimal gradient and the length of the curves cannot be accurately estimated.

2.4.3 Cleaning-up meaningful boundaries

Proposition 1 asserts that if a curve is a meaningful boundary, then it cannot be *entirely* generated in white noise. On the other hand, can we guarantee that no part of a meaningful boundary is contained in noise? Or, for a given meaningful boundary, can we give an upper bound of the size of the part of the boundary that is likely to be contained in noise (i.e. a non-edge region)? To answer this question, we use the *a posteriori* length distribution

$$P(L \geq l | \min_{x \in C} |Du(x)| \geq \mu). \quad (4)$$

Contrary to the probability appearing in Def. 1, this one penalizes long curves not only through the gradient value. To compute it, we need the a priori distribution $P(L \geq l)$ that a level line in noise has a length greater than l . This law can be correctly estimated for $l \leq 1000$ (to give an order of magnitude), see Fig. 4. For higher values, there are too few level lines. By using Bayes' rule, we derive

$$P(L \geq l \mid \min_{x \in C} |Du(x)| \geq \mu) = \frac{\sum_{k=l}^{\infty} P(\min_{x \in C} |Du(x)| \geq \mu \mid L = k) P(L = k)}{\sum_{k=1}^{\infty} P(\min_{x \in C} |Du(x)| \geq \mu \mid L = k) P(L = k)}.$$

(The denominator is nothing but $P(|Du| > \mu)$). By the a contrario assumption (independence of




Figure 4: Log10 of the inverse repartition function of length of level lines in a white noise image. The average length is about 3.5, meaning that most level sets enclose a single pixel.

the gradient along curves), we can still write

$$p_{\mu}(l) \equiv P(L \geq l \mid \min_{x \in C} |Du(x)| \geq \mu) = \frac{\sum_{k=l}^{\infty} H_c(\mu)^k P(L = k)}{\sum_{k=1}^{\infty} H_c(\mu)^k P(L = k)}. \quad (5)$$

Let us now consider an image u with N_{ll} level lines. We also denote by N_l the number of all possible subcurves of these level lines. (N_l is the sum of the squared length of the lines if they are closed).

Assume that Γ is a piece of length L , contained in a non-edge part, described by the noise model. We want to estimate the probability that L is greater than $l > 0$, knowing that $|Du| \geq \mu$. This is exactly $p_{\mu}(l)$, the probability defined in (5). Then, as in Prop. 1, the number $N_l \cdot p_{\mu}(l)$ is an upper bound of the expected number of *pieces* of lines of length greater than l with gradient larger than μ (see Sect. 6 of this paper and [13]). For a fixed μ , let be l such that $N_l \cdot p_{\mu}(l) \leq \varepsilon$. Then, we know that, in average, we cannot observe more than ε pieces of level line with a length larger than l and a gradient every greater than μ . We make the assumption that a point with a gradient less than μ is

located in noise. Let us remove any piece of length l containing such a point. Then all remaining points belongs to a piece of curve with length greater than l with gradient larger than μ , which cannot be due to chance.

This yields a clean-up algorithm for boundary detection.

1. Detect meaningful boundaries.
2. For a fixed $\mu > 0$, let $\mathcal{L}(\mu) = \inf\{l, N_l \cdot p_\mu(l) < \varepsilon\}$.
3. For any meaningful boundary, remove all subcurve of length $\mathcal{L}(\mu)$ containing a point where $|Du| \leq \mu$.

This introduces a parameter, μ . When μ gets larger, $\mathcal{L}(\mu)$ decreases, so that the clean-up remove more numerous but smaller pieces of curves. The choice of μ can be determined by applicative considerations. Detected edges may be used for different purposes, for instance shape recognition or image matching. Letting $|Du|$ less than 1, means that we cannot locate edges with an accuracy better than one pixel. Thus choosing $\mu = 1$ for all images is not restrictive. We also check that for μ about 1, we obtain values of $\mathcal{L}(\mu)$ less than a few hundreds, which is compatible with the empirical estimation of the *a priori* length distribution.




Figure 5: Meaningful boundary clean-up. On the left the original image. In the middle, the meaningful boundaries with local histograms, see Sect. 4. Boundaries are found in the sky. They are detected since the gradient in the sky is regular because of the smoothly changing illumination. The gradient value is about 0.2. Even though they are not smooth at small scale (they cannot be well located, due to the too small gradient), they are nearly parallel at large scales, which can be explained, a posteriori. Now, these boundaries may not be very useful for shape recognition purposes, because of their bad localization. On the right, the result after the clean-up procedure with a gradient threshold equal to 1.

2.4.4 Geometrical information reduction

Caselles et al. claim that pieces of level lines are the basic objects of image analysis [4]. Then, suitably encoded (that is to say in a stable and invariant way) pieces of level lines could directly be used to feed a shape recognition algorithm. There is no theoretical obstruction to encoding all level lines. Lisani et al. [26] describe an encoding method for shape matching. Intrinsic frame for pieces

of level line are defined from inflexion points or bitangent points. These normalized curves are taken as shape elements (we call them *codes*); their number is basically proportional to the number of inflexion points. Thus, long and oscillating curves are very costly, and at the time being, the method is not applicable in reasonable time if all the topographic map is encoded. The MB model is used to compute a raw primal sketch [28], that contains most information of shape contours. Experiments in [13] and this article show that the MB model give sufficient information: it is sufficiently complete and we now discuss the geometrical compression it provides.

A natural geometrical measure for image complexity is the total variation. If u is an image in L^1 , its total variation is

$$TV(u) = \sup \left\{ \int u \operatorname{div} g, \quad g \in C^1(\mathbb{R}^2, \mathbb{R}^2), \quad \sup |g| \leq 1 \right\}. \quad (6)$$

When $Du \in L^1$, $TV(u) = \int |Du|$. The total variation also has the classical geometrical interpretation, given by the coarea formula [16],

$$TV(u) = \int_{\mathbb{R}} \mathcal{H}^1(\partial\chi_\lambda(u)) \, d\lambda, \quad (7)$$

where \mathcal{H}^1 is the one-dimensional Hausdorff measure. When u only assumes a finite number of values, the total variation is nothing but the sum of the perimeters of the level sets. Obviously, approximation of $\int |Du|$ by mere finite differences is much faster than summing all level lines perimeter. Both estimates give very close results (relative errors about 0.3%). The ratio between the total variation and the sum of meaningful boundaries perimeter gives an amount of the geometrical compression. This ratio is experimentally about 15. Similar results are obtained with the total curvature

$$TC(u) = \int_{\lambda} \int_{\partial\chi_\lambda(u)} |\kappa(s)| \, d\mathcal{H}^1(s) = \int |Du| \left| \operatorname{div} \frac{Du}{|Du|} \right|. \quad (8)$$

(The second equality is a consequence of the coarea formula.)

If we now encode shape elements by Lisani's algorithm, the gain is often much better than that, especially for textured or low contrasted images. We think the reason is the following. Recent studies proved [18, 24], that under scale invariance hypotheses, images have not a bounded total variation. This blow-up is mainly due to small objects (noise and microtexture). The corresponding level lines are usually too short to be meaningful boundaries, and the observed total variation decrease is due to their elimination. On the other hand, these lines are too simple to contain an inflexion point, a flat point or a bitangent. Thus, they are not encoded either, even if we keep all level lines. On the opposite, in noise there are only very few long level lines. These lines have a very complicated fractal geometry, and numerically contain many inflexion points. The distance between two inflexions points in white noise is approximately 3. (This, in turn, is also the mean length of level lines in white noise.) As a result, a single long line in white noise may produce several thousands of shape elements, which is the empirical approximate number of codes obtained for all meaningful boundaries for images of size 500×500 ! Such lines may also appear in real images, in low contrasted areas

with an illumination gradient as in Fig. 5. Their contribution is negligible for the total variation (the gradient is low), while it is disproportioned for shape coding. Thus, for practical shape recognition, we must ensure as in Sect. 2.4.3 above, that no pieces of level lines due to noise are detected.

2.4.5 Meaningful boundaries vs. Haralick's detector

In this section, we comment the main differences between the meaningful boundary model and the classical edge detector introduced by Haralick. The meaningful boundaries are based on the topographic map of grey level images, which gives a complete topological representation of grey level images. Caselles, Coll and Morel [4, 6] detail all the properties of this representation. A first advantage of this representation is its stability: even with an important amount of noise, many level lines do not change much. A second advantage is its invariance with respect to global contrast change. But the main property is the structure of this representation: it is a set of nested curves that are either closed or meet the image boundary. As a consequence, level lines have two of the main properties usually expected in edge detection or image segmentation: they are curves (and not sets of points), and are embedded in a hierarchical structure [17, 33, 40]. Moreover, away from critical points, level lines coincide with isophotes. As a consequence, for almost any level, the gradient is almost everywhere normal to level lines, which makes level lines good candidates for edges.

Following Haralick [19], edges are the maxima of the gradient norm in the direction of the gradient, such that the gradient is larger than a given threshold. Thus, for a grey-level image u , they are the zero-crossings of $D^2u(Du, Du)$. It is well-known that this quantity is numerically sensitive to noise and that a smoothing is necessary. Thus in practice, u is first convolved with a gaussian with standard deviation σ (we denote by g_σ this gaussian and $u_\sigma = g_\sigma * u$) and the points where $D^2u_\sigma(Du_\sigma, Du_\sigma)$ changes sign and $|Du_\sigma| > \mu$ are edges points. Although there have been some attempts to automatically determine the scale parameter σ [25], edge detection widely remains multiscale as predicted by Marr [28]. On the contrary, meaningful boundaries do not need smoothing for suitable sampled images (see next section) and can be run with no image dependent threshold. Haralick's detector provides with a set of points or a few pixels long curves. The way they should be connected is far from obvious and may lead to a very high computational complexity; this is one more problem structurally handled by level lines. Last but not least, Haralick's operator is inefficient for corners and junctions. Indeed, at those points, the gradient direction is very badly estimated and edges may be severely cut. At a junction, level lines bifurcates. Figure 6 shows the meaningful boundaries and Canny's filter near two junctions. junction [5].

2.4.6 Color edges

Let $u = (u_R, u_G, u_B) : \mathbb{R}^2 \rightarrow \mathbb{R}^3$ be a color image given by its three components in the RGB space. Since topographic maps are based on the total order relation on gray levels, they cannot be immediately generalized for multi-channel images. In [7], Caselles, Coll and Morel have proposed to define topographic maps based on the luminance and have shown that it contains most of the image information. The luminance is defined by

$$L(u) = \frac{u_R + u_G + u_B}{\sqrt{3}}.$$




Figure 6: Junction and level lines. On the left, the original image. Middle, Haralick's detector implemented with Canny's filter on the area designated on the left image. Note how the contour is broken at the junction, due to the bad estimate of the gradient direction, and the high number of edge pieces. Right: detailed view of meaningful boundaries on the region. There are two level-lines, each corresponding to an edge part.

From the topographic map of $L(u)$, we can generalize the MB model by using a color gradient. At each point let us denote by n the direction of the gradient, that is $n = \frac{Du}{|Du|}$. We then denote by $\frac{\partial u}{\partial n}$ the gradient of u in the direction of n , that is to say the vector

$$\frac{\partial u}{\partial n} = \left(\frac{\partial u_R}{\partial n}, \frac{\partial u_G}{\partial n}, \frac{\partial u_B}{\partial n} \right) = (Du_R \cdot n, Du_G \cdot n, Du_B \cdot n). \quad (9)$$

We then apply the MB algorithm by replacing the norm of the single-channel image gradient by $|\frac{\partial u}{\partial n}|$. As can be expected from [7], the results are almost the same as for grey-level images. Of course, the results are slightly dependent on the choice of luminance and of the gradient norm. It is well known that different combinations of colors may lead to the same luminance. It is experimentally quite seldom, and we did not observe any real gain or loss to use color.

3 Image quantization and gradient norm distribution

A common belief about edges is that the value of the gradient should be very large. The subjective contours (see Kanisza [20], for instance) are seen as some exceptions to this rule. On the contrary, we believe that a lot of edges are nearly subjective (this is nearly always the case in paintings after the Renaissance). In any case, numerical evidence proves that many edges have a contrast that is not larger than 4 of 5, which is quite comparable to texture and noise. This, of course, makes gradient thresholds quite difficult to tune in classical edge detection. This also has an unexpected consequence for meaningful boundaries: the gradient distribution does not really need to be accurate for large values but for *small* values. This means that grey levels have to be suitably quantized in smooth areas, while it does not have a real importance around very contrasted edges. Image quantization is performed after image sampling, and creates high frequencies variations. As a consequence, coarsely

quantized images are not well sampled, in the sense that their spectrum does not vanish close to the critical frequency. This also implies a coarse quantization of the norm of the gradient, especially for small values. This was known to be a problem for the gradient orientation [11], but not for the norm. Let us use the simple gradient estimate on a 2×2 mask, defined by


$$\frac{\partial u}{\partial x} \simeq u_x(i, j) = \frac{1}{2}(u(i+1, j) + u(i+1, j+1) - u(i, j) - u(i, j+1)), \quad (10)$$

$$\frac{\partial u}{\partial y} \simeq u_y(i, j) = \frac{1}{2}(u(i, j+1) + u(i+1, j+1) - u(i, j) - u(i+1, j)). \quad (11)$$

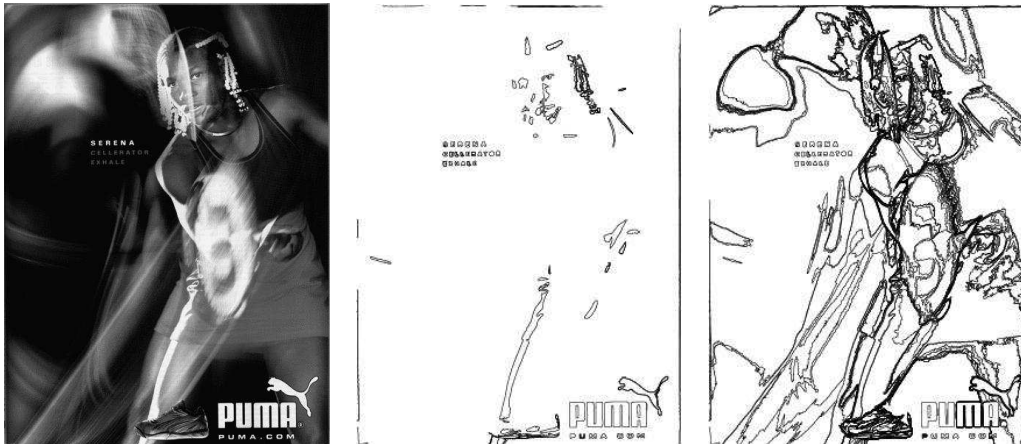
In this case, $u_x(i, j) = u_y(i, j) = 0$ if and only if $u(i+1, j+1) = u(i, j)$ and $u(i+1, j) = u(i, j+1)$. Thus u does not need to be constant on the four points (i, j) , $(i+1, j)$, $(i, j+1)$ and $(i+1, j+1)$. It may even assume very different values. For instance, this numerical gradient is everywhere equal to zero on a perfect chessboard. However, such an image is not well sampled and the gradient computation is not correct. In [11], for alignment detection purposes, the proposed dequantization was a $(\frac{1}{2}, \frac{1}{2})$ translation of the image by using Shannon's interpolation. It is well known that this creates Gibbs effect along edges. Since Gibbs effect leaves the gradient orientation modulo π unchanged, this is not a problem for alignment detection, while it is not suitable for the norm of the gradient.

In order to dequantize the norm of the gradient, we simply propose to apply a slight gaussian smoothing to the image. Remark that this is not the same as to use a different mask to compute the gradient, since level lines are also extracted on the smoothed image. Numerically, a standard deviation of 0.5 is nearly the smallest possible value. We emphasize that this convolution only aims at removing quantization noise, but not to remove small scales elements of the image. In practice, there is hardly any visual difference before and after smoothing. On Fig. 7, we display the meaningful boundaries without and with smoothing. The convolution dramatically changes the gradient distribution close to 0. It also removes many critical and saddle points due to quantization. While the image seems nearly unchanged, the detected lines change incredibly. Of course, one can argue that smoothing creates some structures in the image, and this is why we obtain more detections. This argument is not true. Indeed, the Helmholtz principle asserts that there is less than ε detections in a *well-sampled* image of noise. Smoothing the image is equivalent to decrease Nyquist's rate. Thus, we have to apply the same algorithm as above after a suitable sampling of level lines. For a very slight smoothing of standard deviation 0.5, we did not observe any change.

Remark 1 In order to deal with unstable values of the gradient, we could think of relaxing the definition of a meaningful boundary. As for alignments [12], we could require that a large proportion of points of the boundary have a large gradient. That is to say, the tested events would be of the kind "at least k points among l have a gradient larger than μ_i ". Under the a *contrario* assumption of independence, the probability of such an event is given by the tail of the binomial law. This definition gives many more detections, and avoids all the accidental misdetections due to bad gradient computation. On the other hand, a lot of level lines pass from one object to another. In this case, the cleaning procedure becomes necessary, since this a *contrario* model accepts that a long part of a boundary may not coincide with an edge.



(a) Left: inverse repartition function of the gradient of the image below with no gradient norm dequantization (detailed view near 0), that is to say $P(|Du| > \mu)$ as a function of μ . Right: the same image after convolution with a gaussian with standard deviation equal to 0.5. For $|Du| = 4$ the probability drops from 0.92 to 0.56.



(b) On the left, the original image is much quantized since it has a very low contrast. This leads to bad gradient estimation and a lot of missing detections (middle). Gradient dequantization leads to more correct detections (right).

Figure 7: Influence of image quantization on meaningful boundaries.

4 Local boundary detection

In the model above, the values of the gradient are random variables whose distribution is empirically estimated. It is simply the histogram of the gradient in the image. One can argue that this distribution is too global. This also yields what we call the “blue sky effect”. Consider an image containing two parts: a contrasted or textured one (e.g. ground) and a smooth one (e.g. sky). Then, we can observe an over-detection in the ground, and an under-detection in the sky. Indeed, the sky only contributes with small values in the histogram. Thus we tend to detect anything which is more contrasted than the sky, and nearly anything is detected in the ground. On the contrary, the contrasted ground makes the detection more difficult for regions with a small contrast. This is not in agreement with human vision, since we locally adapt our perception of contrast. Objects are masked in contrasted regions, while our accuracy is improved in low contrasted regions (up to some physiological thresholds). In this section, we address this local adaptivity to contrast. It does not use new concepts and is an adaptation of the meaningful boundary model. We first describe the algorithm, then show experiments.

4.1 Algorithm

Assume that we have detected a closed boundary. Then it divides the image into two connected components: the interior and the exterior of the curve. Then, we can compute the empirical contrast distribution in the interior on the one hand and in the exterior on the other hand. We then independently detect new meaningful boundaries in each connected component. We then apply this procedure recursively. Since the size of the level line tree is finite, it is clear that we end the detection in a finite number of steps.

The situation is actually a bit more complicated. First, this method depends on the order we use to describe the image boundaries. We simply choose to start with the most meaningful boundaries. Second, boundaries are not always closed. In this case, their endpoints belong to the image border. They still cut the image into two connected components. Unfortunately, there is no clear notion of interior and exterior. An algorithmic choice is made, but is purely algorithmic and arbitrary from a perceptual point of view [32]. Thus, we cannot rely on this choice of interior, which conflicts with closed boundaries. However, we can first apply the detection to open boundaries, then to the closed ones. (Open boundaries contain all the closed ones, since level lines are nested.) More precisely, we proceed as follows.

Let us call R_0 the root boundary, that is the (non-meaningful) boundary containing all the image. If C is a boundary, we denote by $\text{Int } C$ its interior.

1. Set $R \leftarrow R_0$. (Local root.)
2. set \mathcal{M} , the set of already stored in R meaningful boundaries. Initially, \mathcal{M} is empty.
3. Let $R' \leftarrow R \setminus \cup_{C \in \mathcal{M}} \text{Int } C$.
4. compute the histogram of $|Du|$ in R' .

5. Use this histogram and detect the maximal meaningful boundaries included in R' . Let \mathcal{N} be the maximal meaningful boundaries defined by $C \in \mathcal{N}$ if and only if

$$\begin{cases} \text{Int}(C') \subsetneq \text{Int}(C) \Rightarrow NFA(C) < NFA(C') \\ \text{Int}(C) \subset \text{Int}(C') \Rightarrow NFA(C) \leq NFA(C'). \end{cases} \quad (12)$$

Otherwise said, the boundaries in \mathcal{N} have an optimal NFA. Note that this is stronger than the maximality defined in Sect. 2.3 since we go across monotone sections. We call the boundaries in \mathcal{N} the total maximal boundaries. The subtree with root equal to R that remains by keeping only the boundaries in \mathcal{N} has only two levels: the local root R , and \mathcal{N} . Since the interior of open boundaries is arbitrary, we do not mix the detection of open and closed boundaries. In practice, this means that if we detect an open meaningful boundary C , we apply the definition of total maximal boundary (12) only to open boundaries containing C or contained in C .

6. if $\mathcal{N} \neq \emptyset$, then we have detected new boundaries in the complementary of the already detected ones. Then,
- (a) set $\mathcal{M} = \mathcal{M} \cup \mathcal{N}$. By construction, all the closed boundaries in \mathcal{M} have disjoint interior.
 - (b) return to step 3.
7. If $\mathcal{N} = \emptyset$, there are no new boundaries in the local root and in the complementary of the currently detected boundaries. We then continue the research at lower levels of the tree. For any boundary $C \in \mathcal{M}$,
- (a) store C .
 - (b) set $R \leftarrow C$, and $\mathcal{M} \leftarrow \emptyset$.
 - (c) return to step 3.

Remark 2 Each boundary may be tested more than once. Thus, the number of false alarms has to be multiplied by the maximal number of visits of a boundary, which is bounded by above by the level lines tree depth. In fact, each detected boundary often lies in the middle of the local root, and this divides the tree depth by 2. Thus in practice, the maximal number of visits of a boundary is like the logarithm of the initial tree depth. In practice, it is always much smaller than 100.

4.2 Experiments on locally contrasted boundaries

In Fig .9, we show the difference between the detection with a global contrast histogram and the updated local histogram. To give an idea of magnitude of the number of false alarms, the boundary delimiting sky and the foreground has a NFA equal to 10^{-357} . This means, that, in average, we need to look 10^{357} images to find such a contrasted curve by chance! The smaller boundaries around the opening on the top of the tower have NFAs about 10^{-10} .

Very interestingly, using local contrast removes boundaries in texture. This is logical since the local




Figure 8: Example of local research of meaningful boundary. (a) the initial boundaries. They are oriented such that the tangent and the interior normal form a direct frame. We compute the NFA of each boundary. In solid line, we draw the total meaningful ones. Two are open, one is closed. Remark that the interior are disjoint, because of total maximality. While we detect some open curves, we ignore the closed ones. (b) While we detect new meaningful boundaries, we compute the contrast histogram in the complementary of the interior of the open detected boundaries and resume research in this part of the image. In R' , the exterior of the detected open boundaries, we detect a total maximal boundary. Remark that this boundary may have been already detected but rejected because of open boundaries. We assume here that no new open boundaries are meaningful. Thus, we keep this closed boundary. (c) We resume the research (with recomputed histogram) in the exterior (white part) of the detected boundaries, until we cannot find new ones. When this is over, we then compute the local contrast histogram in each region R_1 , R_2 , R_3 and look for boundaries inside them. (d) A boundary R_4 has been detected in R_1 . Compute the local histogram in $R_1 \setminus R_4$ and detect boundaries. (e) Finally, we scan for boundaries in R_4 with new local contrast histogram.




Figure 9: Influence of local contrast. From left to right: original image, maximal meaningful boundaries, local maximal meaningful boundaries. Regularity is not taken into account. There are 280,000 boundaries in the initial image, 652 in the second one and 193 in the last one. Texture is removed since local contrast (for instance) on the church tower is much more demanding than the global histogram. As the texture is uniform, no level line is a large deviation to the empirical local contrast, yielding no detection. This is very good for shape analysis where we often want to distinguish texture from real shapes.

contrast in textured regions (as on the tower) assumes larger values than in the rest of the image. Thus, this decreases the NFA of boundaries and most of them simply disappear in textured regions. This is a masking phenomenon.

Let us explain why this is useful for shape recognition. In general, a shape recognition algorithm can be divided in four steps:

1. shapes extraction
2. (invariant) encoding
3. comparison: compute some distance between encoded shapes
4. decision: accept or reject pairs of matching shapes

Present and future applications need to compare images in huge databases, where we have no a priori that two images, or two shapes should match. Since every procedure in the above methodology is very costly, it is interesting to limit the number of encoded shapes and to try to keep the “most meaningful”.

At the time being, there is no general model of shapes [44]. Nevertheless, for shape recognition algorithm, we can give empirical observations of what a “good shape” is. First, it should be stable in terms of extraction. This is generally expressed in terms of contrast and regularity, and the method we describe in this paper gives quantitative arguments. (Regularity is the object of the next section.) For encoding, a good shape should not be too simple, especially if we are interested in an invariant recognition. For instance, most convex shapes are very alike in affine invariant shape recognition. Assume that we have chosen an affine invariant distance between shapes. If we want to be sure that two convex shapes match, the distance between them has to be very small. Indeed, two convex shapes can casually be close to each other, while the probability that it occurs for more complex shapes is very small (this means that recognition is relative to the database and to the query [35]). On the other hand, a shape should not be too complex, since complexity usually makes the encoding longer and more difficult. Because of occlusions, we usually try to match pieces of shapes. Very complex shapes will be divided in numerous pieces, making computations longer.

Now, it is well known that texture are strongly damaged by compression. Thus level lines in texture may not be reliable when two images come from different sources (with different quality, compression rate etc...). Moreover, they are very complex, and yield many encoded pieces of curves. If these curves match for two different images, then those images are certainly exactly the same. Now, the computational cost may be too high for some applications, where we may want to detect a particular shape (a logo for instance) in a database. Thus it may be useful to automatically remove contrasted regions corresponding to texture. This is what the local contrast detection makes in practice.

The argument above is reversed for stereo images registration. In this case, we have the strong a priori that the images are nearly the same, and the goal is to register them as best as possible. In this application, textures can also give some useful information. (See Fig 10).

The elimination of textures is the first aspect of the “blue sky” effect. On the contrary, local contrast should make curves in low contrasted areas more detectable. This is also what we empirically observe: we detect illumination gradient (See Sect. 2.4.2, Fig. 5 and 11). They can be due to the

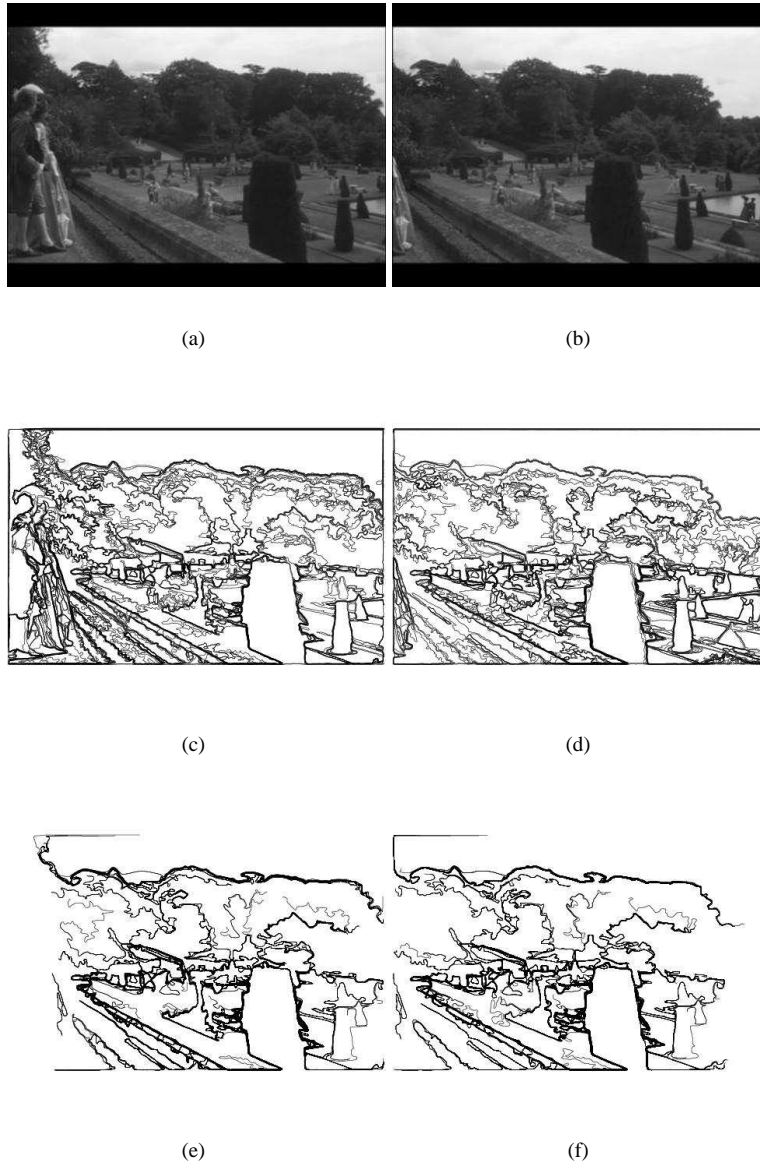


Figure 10: Image registration. (a) and (b) are two images from a movie during a rightward traveling. (c) and (d) are the meaningful boundaries in the previous images. (e) and (f) are the pieces of level lines of (c) and (d) that match with a number a false alarms less than 10^{-7} . We use the algorithm developed by Lisani, Musé and Sur [26, 34, 35], which uses an a contrario definition of shape matching.

vicinity of the light source, or to the variation of the orientation of the surface of a three dimensional object with respect to the light source. Such lines do not correspond to the usual notion of shapes (objects). Nevertheless, it is logical to detect them as remarkable structures.




Figure 11: Illumination, local contrast and regularity. Left: original image. Middle: meaningful contrasted boundary. Right: meaningful contrasted and smooth boundary with local contrast. With contrast only, a single boundary appears on the right with the contrast due to illumination. If contrast is localized, then more boundaries are detected. If we also add a regularity constraint (see Sect. 5.1 below), there are still more detections. These boundaries are very different from texture since they are nearly convex and parallel. They are eliminated by the cleaning procedure described in Sect. 2.4.3.

5 Meaningful boundaries or snakes?

In [15], Desolneux, Moisan and Morel compared the MB model with variational snake theory. This may seem a bit weird since the MB model only uses contrast observations along a curve, while snakes are also required to be smooth. In fact, the explanation for natural images is that contrasted boundaries often locally coincide with objects. Thus, they are also incidentally smooth. Whereas smoothness seems to be optional for the detection, it may give a better localization of the contour. In this section, we propose a method to incorporate smoothness to the detection. There are only few additional detections, while the position of the maximal meaningful boundaries may change a little bit. The NFA also significantly decreases. The small number of new detections and the fact that each partial detector can detect most image edges prove *a contrario* that contrast and regularity are not independent in natural images.

An *a contrario* model of regularity has been proposed in [3]. It assumes that the variation of the orientation of the tangent between two samples is a random value uniformly distributed in $(-\pi, \pi)$. Thus, the implicit *a contrario* model is random walks with isotropic and independent increments. This model is not really adapted for the following reason. All the curves we detect are level lines, thus boundaries of compact sets. As a consequence, they do not self-intersect. While the local influence is not clearly visible, this implies that long level lines are much more regular than random walks. This logically leads to an overdetection of long level lines because the independence assumption is strongly violated at very long range. The solution we propose is to stick to Helmholtz principle: “no

detection in white noise”. Thus we have to learn the regularity of level lines in white noise, and use this as the a priori distribution.

5.1 Definition of local regularity

Let $l_0 > 0$ be a fixed positive value. Let C be a rectifiable planar curve, parameterized by its length. Let $x = C(s_0) \in C$. With no loss of generality, we assume that $s_0 = 0$.

Definition 3 We call regularity of C at x (at scale l_0) the quantity

$$R_{l_0}(x) = \frac{\max(|x - C(-l_0)|, |x - C(l_0)|)}{l_0}. \quad (13)$$




Figure 12: Regularity definition. The regularity at x is obtained by comparing the radius of the circle with l_0 . The radius is equal to l_0 if and only if the curve is a straight line. If the curve has a large curvature, the radius will be small compared to l_0 .

Of course, this definition really makes sense if the length of C is larger than $2l_0$. This definition of regularity (see Fig. 12) is related to the Hausdorff dimension of C around x . First, $R_{l_0}(x) \leq 1$, with equality if and only if either $C((-l_0, 0))$ or $C((0, l_0))$ is a line segment. On the contrary, if $R_{l_0}(x)$ is small, then the curve is highly curved around x .

We can also interpret $R_{l_0}(x)$ as a function of the local curvature. Indeed, if C is a circle with large enough radius ρ , then

$$R_{l_0}(x) = \text{sinc} \left(\frac{2l_0}{\rho} \right), \quad \text{where } \text{sinc } x = \frac{\sin x}{x}. \quad (14)$$

This approximation is valid when l_0 is small compared to ρ . In this case, the regularity is a non-increasing function of the curvature.

This definition is not purely local, but it is also less sensitive to noise compared to differential measures as the curvature. Let

$$\mathcal{H}_{l_0}(r) = P(x \in C, C \text{ is a white noise level line and } R_{l_0}(x) > r). \quad (15)$$

This distribution only depends on l_0 and can be empirically estimated. Of course, we learn it on level lines whose length is much larger than l_0 in order to avoid quantization effects.

Remark 3 As expected, the distribution \mathcal{H}_{l_0} is very different in white noise and natural images. In natural images, the histogram of R_{l_0} has a peak at 1, corresponding to real objects boundaries (which often contain alignments). In some textured images, as paintings, most edges are not real but subjective and this is clearly visible on the histogram of R_{l_0} . See Fig. 13. The distribution also clearly depends on l_0 . When l_0 grows, the histogram mode moves to lower values. However, we obtain the same qualitative behavior as above. In Appendix A, we use these distributions to compute the Hausdorff dimension of white noise level lines. We then quantitatively check that they are much more smooth than (self-intersecting) isotropic random walks.

Remark 4 In [18], Gousseau and Morel prove that natural images are not smooth. In particular, their total variation seems to be unbounded. This is not a contradiction with the experiment we displayed. Indeed, images are not smooth because they contain many small objects due to texture. Our regularity computation is not completely local, so we remove little objects before computing the histograms.

Again, the choice of l_0 is a natural question. Of course l_0 should be larger than Nyquist distance. It should not be too large either. In experiments we have chosen $l_0 = 10$. But, since NFAs are additive, we may also choose several reasonable values of l_0 (say $l_0 = 5, 10, 20$) and multiply the NFAs by the number of l_0 . In practice, changing l_0 influences the number of samples and best NFAs are attained for small l_0 .

5.2 Meaningful contrasted and smooth boundary

Now that we have a background model of regularity, we use it to detect regular curves a contrario. It is natural to assume, in the background model, that contrast and regularity are independent. Thus

$$P(C \text{ is contrasted and smooth}) = P(C \text{ is smooth}) \times P(C \text{ is contrasted}).$$

Definition 4 Let C be a level line. Let

$$\nu = \min\{|Du(x)| > \mu_i, x \in C\}, \quad (16)$$

$$\rho = \min\{|R_l(x)|, x \in C\}, \quad (17)$$

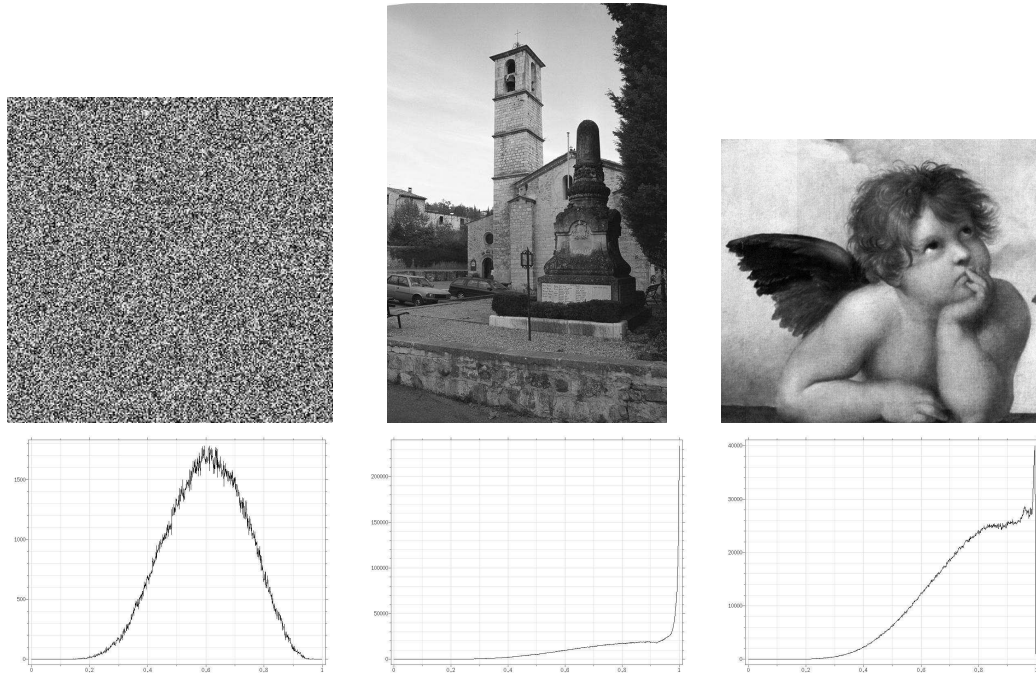


Figure 13: Regularity histograms. Upper row: a white noise image, a scanned photograph and a scanned photograph of a painting. Bottom row: the three regularity histograms for $l_0 = 10$. Since its histogram vanishes near 1, white noise does not contain any alignments or smooth curves, as foreseen. Nearly all natural images (containing true edges) have a regularity histogram like the second one. The third image contains mostly subjective edges, as it is composed of painted strokes. As a consequence, the regularity histogram is much less concentrated around 1 as for “natural” images. If we now unzoom the three images (with an adequate smoothing before downsampling), then the first histogram remains unchanged (scale invariance), while the other two have regularity histograms like the second one. Indeed, after unzooming, most textures and small scale features disappear, and small gaps get filled.

be respectively the minimal quantized contrast and regularity along C . Let

$$NFA_{cs}(C) = N_{ll} H_c(\nu)^{1/2} (\mathcal{H}_{l_0}(\rho))^{1/2l_0}. \quad (18)$$

We say that C is a ε -meaningful smooth boundary if $NFA_{cs}(C) < \varepsilon$.

The number of false alarms is the product of number of level lines and the probability that the contrast and the regularity are simultaneously larger than the observed values along a curve with prescribed length taken in the background model. The probability is computed in the a contrario model where contrast and regularity are independent and local observations are mutually independent.

As above, this research can be recursively performed by computing local histograms of the gradient. In experiments, detection results are qualitatively equivalent with or with no regularity. On the other hand, NFA may decrease a lot for smooth boundaries. Even though the detection is not changed in one single image, it is still interesting to decrease the NFA as much as possible. Indeed, we may want to detect boundaries not in a single image but in a database (for instance in shape recognition applications). We can evaluate the size of a “universal database” (containing all the images ever seen by human kind) to 10^{15} . Thus, curves with a NFA lower than 10^{-15} in a single image can also be considered as universally meaningful.

5.3 Comparison with active contours

Active contours is one of the most popular techniques of boundary detection. The first works of Kass, Witkin and Terzopoulos [21] have been improved and generalized by many authors. Recent models are more intrinsic, can be expressed implicitly (which ease the possible topological changes of the active contours) and can use image statistics [8, 36]. In this section, we do not focus on any particular active contour model, but try to compare a generic model with meaningful boundaries. Such a comparison has already been made by Desolneux, Moisan and Morel [15] for meaningful boundaries. Even though these boundaries are only contrast-based, they show that they are very close to active contours in general and particularly to the model of Kimmel and Bruckstein [22]. Since in this paper we have also introduced a regularity criterion, comparison is even more adequate. Let us briefly give a generic active contour model: it is a curve that fits shape contours (hence contrast should be large along the contour) and which is also as smooth as possible. The problem usually assumes a variational formulation. An optimal curve minimizes an energy of the type

$$E(C) = \int_C g(|Du(C(s))|) + \lambda h(\text{curv}(C(s))) ds, \quad (19)$$

where Du is the gradient of a given grey-level image, g is a nonincreasing function, $\text{curv}(C(s))$ is the curvature of C at point $C(s)$, h is a nondecreasing function and s is the arc-length. The optimal curve is a trade-off between the external energy depending on the image gradient, and the internal energy depending on the curve itself only. Such a model can accurately give the position of the contour. However, it has several drawbacks:

- The model assumes that there is a contour: It cannot be used as a detection algorithm. This also explains why active contours are also introduced in Bayesian models, where the real question is: knowing that one object is present, what is the best candidate?

- The initialization is crucial.
- The optimal balance parameter λ (which, for homogeneity reasons, can also be viewed as a scale parameter) is unknown and depends on the image. It has a strong influence on the result.

If we now only consider the homogeneity of the different energy terms, we have to minimize a potential of the form $Lg(|Du|) + \lambda Lh(\text{curv } C)$, L being the length of the curve. Let us now consider the meaningful smooth boundary model. A meaningful curve has a small probability to occur in the a contrario model. Our regularity measure is a non increasing function of the curvature (see (14)). Thus, for a meaningful curve, the quantity

$$(H_c(|Du|)^{L/2} \mathcal{H}_{l_0}(R_{l_0}(C)))^{L/2l_0}$$

is small. Let us now take the logarithm of this expression. We obtain an expression of the type

$$L(E_{ext}(|Du|) + E_{int}(\text{curv } C)),$$

where E_{ext} is a non increasing function of $|Du|$, and E_{int} is a non decreasing function of the curvature. The model is qualitatively alike a snake model. Nevertheless, there are three major differences:

1. There is a quantitative criterion to decide if the curve has to be detected. Contrary to snakes algorithm, meaningful boundaries detection is *not* a minimization algorithm. It is well known in active contours model that the value of the energy of the minimizer has no interpretation. All that we can say is that a candidate is better than another one. Our model gives a meaning to the energy-like term. Thus, there is no need for a minimization since we can give thresholds under which a candidate has to be detected.
2. Meaningful boundaries are level lines. Thus, no initialization by hand is needed.
3. We do not have to fix the weight functions g and h as well as the scale parameter λ .

5.4 Experiments on smooth meaningful boundaries

In general, adding a regularity criterion does not qualitatively change the result. This is conform with the observation of Desolneux et al. in [15]. Remark also that the criterion we give does not eliminate irregular level lines. Indeed,

$$NFA_{cs}(C) \leq N_{ll} H_c(\nu)^{l/2},$$

(with the same notations as in Def. 4) since $\mathcal{H}_r(\rho) \leq 1$. We can only detect more lines. Of course, smooth boundaries NFA decrease a lot (about 10^{-15}), and this can modify maximal meaningful boundaries. As it was already observed in [3], contrast and regularity are often very redundant, and this explains why the same curves are detected.

Fig. 14, (INRIA desk) is very geometrical and shows the redundancy between contrast and regularity. We can also define NFA for smooth boundaries, with no care of contrast, as

$$NFA_{reg}(C) = N_{ll} (\mathcal{H}_{l_0}(\rho))^{l/2l_0}. \quad (20)$$

We retrieve most edges in the desk image with this definition.




Figure 14: Regularity detectability. The original is the left most. In middle, we display the 204 detected contrasted smooth boundaries as defined in Def. 4. On the right, the 96 smooth boundaries, with no contrast information, defined in (20). All the main boundaries are already present. Of course, contrast may be the main cause of small NFA, since regularity acts at larger scales. For instance, the window panes have NFA about 10^{-150} with contrast and 10^{-15} with regularity only (which still make them detectable in any image database). The desk on the bottom right has a NFA equal to 10^{-60} with contrast and 10^{-20} with regularity, which is already very small.

6 Meaningful edges

6.1 Edges as pieces of level lines

A very long curve may become meaningful even though it is not very contrasted and smooth (See Sect. 2.4.2). This phenomenon did not appear in Desolneux et al. work [13], because of quantization effects. Indeed, the gradient computation for a quantized image implies that there many critical points (zero gradient). If a level line is very long and does not fit an edge, then the probability that it does not cross a critical point is close to zero. Thus (too) many long curves were eliminated. Because of occlusions, meaningful boundaries may not coincide with a whole object boundary. But each of its parts is the part of an object boundary. Thus, it is natural to consider that the basic objects for image analysis are pieces of level lines. This was also proposed in [4, 13]. The drawback is that pieces of level lines are not embedded in a tree structure and we can non longer apply the notion of (monotone) maximality (Sect. 2.3). The following definition of edges incorporates smoothness information to the edges of [13].

Consider a set of digital curves C_1, \dots, C_N , with length (in independent points) equal to l_1, \dots, l_N . Then the number of all possible sampled connected subcurves is bounded by $N_l = \sum_{i=1}^N l_i^2$.

Definition 5 Let Γ be a (connected) subcurve of one of the C_i . Let ν and ρ the minimal quantized contrast and regularity in Γ (See (16) and (17)). Let l the length of Γ . We define the number of false alarms of Γ by

$$NFA_{ed}(\Gamma) = N_l H_c(\nu)^{l/2} (\mathcal{H}_{l_0}(\rho))^{l/2l_0}. \quad (21)$$

We say that Γ is an ε -meaningful edge if $NFA(\Gamma) < \varepsilon$.

We say that Γ is a maximal ε -meaningful edge if and only if

1. for all $\Gamma', \Gamma' \subset \Gamma \Rightarrow NFA(\Gamma') \geq NFA(\Gamma)$.
2. for all $\Gamma', \Gamma \subset \Gamma' \Rightarrow NFA(\Gamma') > NFA(\Gamma)$.

This detection destroys the tree structure of the topographic map and we loose the notion of monotone maximality introduced in Sect. 2.3. We can apply the edge detection in two different ways:

- apply it to all the image level lines
- apply it only to meaningful boundaries.

In the first case, we automatically get all the meaningful pieces of all level lines. We may get thick sets of edges concentrated along true edges. The tree structure can no longer be used. (However, it is possible to use another notion of maximality that make those wide strokes thinner. [1, 10].) In the second case, we already know that the a contrario model is false on the data curves (else they would not be meaningful boundaries). Thus, it is not possible to apply the same detection thresholds on this boundaries. However, it is possible to perform a variational research: we can select the most meaningful pieces in each meaningful boundary. Practically, the meaningful edge algorithm needs no parameter at all, and leads to equivalent result to the cleaning-up method of Sect. 2.4.3. However, edges maximality may force to prefer too short edges if contrast is very high on a small portion of line, which is not the case for the cleaning procedure that removes low-contrasted parts.

6.2 Experiments on edges

We first display an image with smooth edges alone, that is to say we ignore contrast and only use the regularity criterion: on the Valbonne church image (Fig. 15), only the main structures are recovered. They correspond to very strong edges of man made objects. On the contrary, textures are ignored. See for instance the edge of the tree on the right-hand side of the image. Finally we display an experiment showing edges within meaningful boundaries (Fig 16). We first compute the maximal meaningful boundaries, then only keep their most meaningful parts. This may remove pieces of a boundary that are detectable when compared to noise, but which are much less meaningful than other parts of the boundary.

7 Conclusion

The soundness of Marr and Hildreth's edge detection doctrine [29] has been widely proved in the literature. The results in the previous section show that, moreover, the output of edge detection can directly be sets of curves (and not points to be connected as in the classical Canny detector [2]). Even more important, this can be done with no parameters. We also justify, a posteriori, all the algorithms (as active contours and segmentation) that combines both contrast and regularity. However, we can add that for large scale objects ("large" means in practice more than 20 pixels), regularity and




Figure 15: Valbonne church. The original image is the same as in Fig. 9. Meaningful edges with only regularity criterion ($l_0 = 5$). Only the very regular edges are detected. The other ones are mainly due to texture and appear with contrasts. This experiments show that regularity alone gives comparable results with classical edge detections. It also suffers from the same drawbacks. At junctions (and corners) regularity edges are broken. They do not provide good shape elements by themselves.




Figure 16: Edges. Left: original image (Composition 8 by Kandinsky). Middle: local maximal meaningful contrasted and smooth boundaries. Right: edges within the boundaries. Some boundaries that do not correspond to shapes are detected in the middle. It is due to the fact that the background of the painting is not uniform at all and that shapes are surrounded by a halo, whose level lines are detected.

contrast usually express the same information, that is to say, they allow to detect the same things. This redundancy is compatible with the usual notion of shapes, and we prove that we can detect both properties with no a priori model but a qualitative a contrario model. Regularity information makes detection even surer but are usually not necessary. Besides, incorporating several detection criteria is more heavy, and a compromise between accuracy and speed as to be found. In our shape recognition experiments, we usually prefer to ignore regularity. Another conclusion of our experiments is that no detector allows to detect all the shapes in an image and *only them*. This “failure” is completely coherent with Gestalt Theory: shape (and contour) is not a local notion. First, many contours are subjective and configurations of the type of Kanizsa’s triangle often appear at lower degree. For such contours, all local methods are doomed to fail. Moreover, some curves may appear as contours and good shapes out of their context. If we crop a small part of an image, a boundary may appear contrasted, smooth and fairly complex, as a good shape. However, these small shapes may be masked by much larger ones which impose the natural scale of the image. Thus, the notion of edge is only relative, to a context, or to a precise application. We think that the different models we exposed permit to cover most applications. For instance, in image registration, we may want to use any information for a better accuracy, *e.g.* small variation in texture. Digital elevation models generation uses texture for registration. In this case, we have the very strong a priori knowledge that registration makes sense. If we now want to compare two images in a large database, it is certainly better to keep the most meaningful curves that correspond to shapes as best as possible, and to eliminate all other lines. For practical shape matching by shape elements comparison [26, 35, 34], the MB model with local contrast and cleaning-up automatically eliminates most edges due to texture or small illumination gradient. It gives the most complete and compact shape elements of natural images.

Acknowledgements. We thank Jean-Michel Morel and Agnès Desolneux for all discussions and advice.

References

- [1] A. Almansa. *Échantillonnage, interpolation et détection. Applications en imagerie satellitaire*. PhD thesis, ENS Cachan, 2002.
- [2] J. Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(6):679–698, 1986.
- [3] F. Cao. Good continuation in digital images. In *Proceeding of ICCV 03, Nice*, volume 1, pages 440–447, 2003.
- [4] V. Caselles, B. Coll, and J.M. Morel. A Kanizsa program. In *Progress in Nonlinear Differential Equations and their Applications*, volume 25, pages 35–55, 1996.
- [5] V. Caselles, T. Coll, and J.M. Morel. Junction detection and filtering. In Cucker Felipe, editor, *Foundations of computational mathematics*, pages 23–42. Springer, 1997.
- [6] V. Caselles, T. Coll, and J.M. Morel. Topographic maps and local contrast changes in natural images. *International Journal of Computer Vision*, 33(1):5–27, 1999.

-
- [7] V. Caselles, T. Coll, and J.M. Morel. Geometry and color in natural images. *JMIV*, 16(2):89–105, 2002.
- [8] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. *International Journal of Computer Vision*, 22(1):61–79, 1997.
- [9] T. Chan and L. Vese. Active contours without edges. *IEEE Transactions on Image Processing*, 10(2):266–277, 2001.
- [10] E. d’Angelo. Recherches de formes géométriques significatives. application au traitement d’otolithes marins. Technical report, IRISA, 2003.
- [11] A. Desolneux, S. Ladjal, L. Moisan, and J.M. Morel. Dequantizing image orientation. *IEEE Transactions on Image Processing*, 11(10):1129–1140, 2002.
- [12] A. Desolneux, L. Moisan, and J.M. Morel. Meaningful alignments. *International Journal of Computer Vision*, 40(1):7–23, 2000.
- [13] A. Desolneux, L. Moisan, and J.M. Morel. Edge detection by Helmholtz principle. *Journal of Mathematical Imaging and Vision*, 14(3):271–284, 2001.
- [14] A. Desolneux, L. Moisan, and J.M. Morel. A grouping principle and four applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(4):508–513, 2003.
- [15] A. Desolneux, L. Moisan, and J.M. Morel. Variational snake theory. In S. Osher and N. Paragios, editors, *Geometric Level Set Methods in Imaging, Vision, and Graphics*. Springer Verlag, 2003.
- [16] L.C. Evans and R. Gariepy. *Measure Theory and Fine Properties of Functions*. CRC Press Ann Harbor, 1992.
- [17] P. Felzenszwalb and D. Huttenlocher. Image segmentation using local variation. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition*, pages 98–104, 1998.
- [18] Y. Gousseau and J.M. Morel. Are natural images of bounded variation? *SIAM J. of Math. Anal.*, 33(3):634–648, 2001.
- [19] R. Haralick. Digital step edges from zero crossing of second directional derivatives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6:58–68, 1984.
- [20] G. Kanizsa. *La Grammaire du Voir*. Diderot, 1996. Original title: *Grammatica del vedere*. French translation from Italian.
- [21] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1:321–331, 1987.
- [22] R. Kimmel and A.M. Bruckstein. Regularized laplacian zero crossings as optimal edge integrators. In *Image and Vision Computing, IVCNZ01*, New Zealand, 2001.

- [23] J.J. Koenderink. The structure of images. *Biol. Cybern.*, 50:363–370, 1984.
- [24] A.B. Lee, D. Mumford, and J. Huang. Occlusion models for natural images: A statistical study of a scale-invariant dead leaves model. *International Journal of Computer Vision*, 41(1-2):35–59, 2001.
- [25] T. Lindeberg. Feature detection with automatic scale selection. *International Journal of Computer Vision*, 30(2):79–116, 1994.
- [26] J.L. Lisani, L. Moisan, P. Monasse, and J.M. Morel. On the theory of planar shape. *SIAM Multiscale Modeling and Simulation*, 1(1):1–24, 2003.
- [27] S. Mallat. *A Wavelet Tour in Signal Processing*. Academic Press, 2nd edition, 1999.
- [28] D. Marr. *Vision*. N.York, W.H. and Co, 1982.
- [29] D. Marr and E. Hildreth. Theory of edge detection. *Proceeding of Royal Society of London*, 207:187–207, 1980.
- [30] G. Matheron. *Random Sets and Integral Geometry*. John Wiley N.Y., 1975.
- [31] F. Meyer and P. Maragos. Nonlinear scale-space representation with morphological levelings. *J. of Visual Comm. and Image Representation*, 11:245–265, 2000.
- [32] P. Monasse. *Morphological Representation of Digital Images and Application to Registration*. PhD thesis, Université Paris IX Dauphine, 2000.
- [33] D. Mumford and J. Shah. Optimal approximation by piecewise smooth functions and associated variational problems. *Communication on Pure and Applied Mathematics*, XLII(4), 1989.
- [34] P. Musé, F. Sur, F. Cao, and Y. Gousseau. Unsupervised thresholds for shape matching. In *IEEE Int. Conf. on Image Processing, ICIP*, 2003.
- [35] P. Musé, F. Sur, and J.M. Morel. Sur les seuils de reconnaissance de formes. *Traitement du signal*, 19(5/6), 2003.
- [36] N. Paragios and R. Deriche. Geodesic active regions and level set methods for supervised texture segmentation. *International Journal of Computer Vision*, 46(3):223–247, 2002.
- [37] P. Salembier and L. Garrido. Binary partition tree as an efficient representation for image processing, segmentation, and information retrieval. *IEEE Transactions on Image Processing*, 9(4):561–576, 2000.
- [38] P. Salembier and J. Serra. Flat zones filtering, connected operators, and filters by reconstruction. *IEEE Transactions on Image Processing*, 4(8):1153–1160, 1995.
- [39] J. Serra. *Image Analysis and Mathematical Morphology*. Academic Press, 1982.

- [40] J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905, 2000.
- [41] L. Vincent. Grayscale area openings and closings, their efficient implementation and applications. In J. Serra and P. Salembier, editors, *Proceedings of the 1st Workshop on Mathematical Morphology and its Applications to Signal Processing*, pages 22–27, Barcelona, Spain, 1993.
- [42] M. Wertheimer. Untersuchungen zur Lehre der Gestalt, II. *Psychologische Forschung*, 4:301–350, 1923.
- [43] A.P. Witkin. Scale space filtering. In *Proc. of IJCAI, Karlsruhe*, pages 1019–1021, 1983.
- [44] S.C. Zhu. Embedding Gestalt laws in markow random fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(11):1170–1187, 1999.

A Appendix: Numerical estimation of the Hausdorff dimension of a curve

In order to compute the Hausdorff dimension of identically distributed random curves from the histogram of regularity, we proceed as follows. Let C be a curve.

Definition 6 *The Hausdorff measure of dimension α is defined by*

$$\lim_{\delta \rightarrow 0} \inf_{(B_i)_{\delta\text{-covering}}} \sum_i |B_i|^\alpha,$$

where the B_i form a covering of C and $|B_i|$ is the diameter of B_i . The family (B_i) is a δ -covering of C if $C \subset \cup_i B_i$ and for all i , $|B_i| < \delta$.

The problem to estimate this quantity is that it makes no sense to let $\delta \rightarrow 0$ for digital curves. Indeed, even for white noise, the precision is bounded from below by Nyquist distance. We assume that the curve is self-similar. This allows to examine it at larger and larger scales, instead of letting δ go to 0. Let us cut a curve with length $L = 2Nl$ in N chunks of length $2l$. We measure the regularity $R_l(i)$ at the middle point x_i of each piece. The balls with radius $R_l l$ nearly form a covering of C . It is not a covering because the endpoint of the curve chunk may not be the most remote point from the center (see (13)). Nevertheless, we approximate the measure of C by

$$\mathcal{H}^\alpha(C) \simeq \sum_{i=1}^N (2l R_l)^\alpha \simeq 2^{\alpha-1} L l^{\alpha-1} \overline{R_l}^\alpha,$$

where $\overline{R_l}$ is the mean regularity along C . Let us now consider the curve λC with $\lambda > 1$. We can make the same procedure as above with chunks whose length is equal to $2\lambda l$. Thus we evaluate the measure of λC by

$$\mathcal{H}^\alpha(\lambda C) \simeq 2^{\alpha-1} \lambda L (\lambda l)^{\alpha-1} \overline{R_{\lambda l}}^\alpha.$$

But, if we now use pieces of curves of length $2l$, we also obtain

$$\mathcal{H}^\alpha(\lambda C) \simeq 2^{\alpha-1} \lambda L l^{\alpha-1} \overline{R}_l^\alpha.$$

Thus

$$\lambda^\alpha \overline{R}_{\lambda l}^\alpha = \lambda \overline{R}_l^\alpha,$$

yielding

$$\log(\overline{R}_{\lambda l}) = \left(\frac{1}{\alpha} - 1\right) \log \lambda + \log \overline{R}_l. \quad (22)$$

We can evaluate α by examining the histograms of R_l as a function of l .

For random walks with independent increments, we find $\alpha = 2.02$, whereas the true dimension is 2. For level lines in white noise, we find $\alpha = 1.78$. As expected, the level lines of a white noise image are more regular than random walks.



Unité de recherche INRIA Lorraine, Technopôle de Nancy-Brabois, Campus scientifique,
615 rue du Jardin Botanique, BP 101, 54600 VILLERS LÈS NANCY
Unité de recherche INRIA Rennes, Irista, Campus universitaire de Beaulieu, 35042 RENNES Cedex
Unité de recherche INRIA Rhône-Alpes, 655, avenue de l'Europe, 38330 MONTBONNOT ST MARTIN
Unité de recherche INRIA Rocquencourt, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex
Unité de recherche INRIA Sophia-Antipolis, 2004 route des Lucioles, BP 93, 06902 SOPHIA-ANTIPOLIS Cedex

Éditeur
INRIA, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399