Contrast Invariant Detection of Good Continuations, Corners and Terminators

Frédéric Cao

INRIA

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Contrast Invariant Detection of Good Continuations, Corners and Terminators

Frédéric Cao

N° 4542

Septembre 2002

THÈME 3

ISSN 0249-6399 ISRN INRIA/RR--4542--FR+ENG

Contrast Invariant Detection of Good Continuations, Corners and Terminators

Frédéric Cao*

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

Rapport de recherche n° 4542 — Septembre 2002 — 39 pages

Abstract: We propose a statistical criterion of digital curves regularity. It is not defined from an *a priori* model but *a contrario* to some statistics on random walks. This allows to automatically define detection thresholds in terms of a false detection rate. We apply this algorithm to the level lines of gray level images and experimentally check the statement of the Gestalt Theory following which regularity makes curve conspicuous without any contrast information, and that edges are really often good continuations. We also define good continuations breakings, which are good candidates for T-junctions, corners and Julesz's terminators. We also show that detection is *not* improved by shape scale space.

Key-words: Computer Vision, Gestalt Theory, Helmholtz principle, false detection rate, good continuation, corners, T-junctions, terminators, level lines.

* e-mail: fcao@irisa.fr

Unité de recherche INRIA Rennes
IRISA, Campus universitaire de Beaulieu, 35042 Rennes Cedex (France)
Téléphone : +33 2 99 84 71 00 — Télécopie : +33 2 99 84 71 71

Détection invariante par changement de contraste de bonnes continuations, coins et terminators

Résumé : Dans cet article, nous proposons un critère de régularité statistique pour un ensemble de courbes échantillonnées. Ce critère n'est pas défini à partir d'un modèle de régularité, mais *a contrario* du simple modèle de marche aléatoire à accroissement isotropes et indépendants. Les seuils de détection sont fixés par les données et s'expriment en termes de nombre de fausses alarmes. L'algorithme proposé est appliqué aux lignes de niveau d'images en niveau de gris. On montre que conformément à la théorie de la Gestalt, les bords des objets dans les images naturelles sont très souvent des courbes de régularité improbable et sont détectables sans aucune information de contraste, et l'algorithme proposé permet de retrouver la plupart des contours indépendamment du contraste. On définit aussi des ruptures de bonnes continuations, qui sont de bons candidats pour les coins, les jonctions en T et les terminators de Julesz. Nous prouvons également que lisser les courbes ne favorise pas la détection.

Mots-clés : Vision par ordinateur, théorie de la Gestalt, principe de Helmholtz, nombre de fausses alarmes, bonne continuation, coins, jonctions en T, terminators, lignes de niveau.

1 Introduction

The aim of this paper is to define a new method to detect regular curves in digital images. This is an important topic in computer vision, since regularity is often a property satisfied by shapes contours. This is not a coincidence, but on the contrary, a curve often has to be smooth to be viewed as a contour. Indeed, smoothness is an important element that contributes to what Gestaltists [19, 35] have called "good shapes" formation or shape "Prägnanz". Among other elements entering into the composition of good shapes, we can cite, without being exhaustive, convexity, closedness, symmetry and contrast. A structure satisfying one or several of these properties is very likely to be viewed as a shape and its complementary is then considered as the background. This is particularly true for man made objects which are very often closed, symmetric, contrasted, convex and have smooth or even straight boundaries. In this case, the properties above collaborate and reinforce dramatically the shape formation. But more interestingly they may also conflict. For instance, in Fig. 1, quoted from Kanizsa [18], convexity (even approximate) conflicts with symmetry and it is not always obvious to decide what is the shape and what is the background. In [10], the elementary geometrical properties

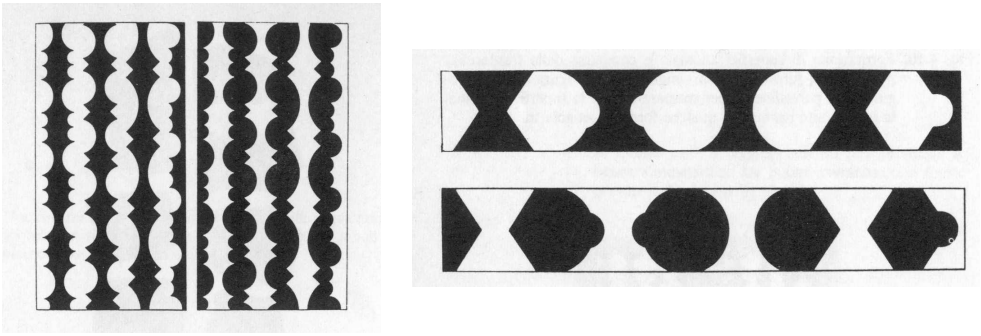

Figure 1: Convexity is often more salient than symmetry. We usually see the most convex parts as the shape, and tend to consider what is left as the background although it is more symmetric (After Kanizsa [18])

cited above have been called *partial gestalts*. They are the most simple and widely independent geometrical properties that yield the formation of perceptual shapes. One of the major statements of Gestaltists was that these partial gestalts defined *grouping* laws (that is how elementary objects assemble to form shapes) or *masking* laws (how they conflict and inhibit). A simple but impressive illustration of the relevance of masking is Fig. 2. The line of the left-hand side is present on the right-hand side. It is not very difficult to find it, but it generally takes some time. Moreover, if we are not aware of the left-hand figure, we simply do not see it on the right. Following gestaltists, this disappearance is due to masking, that is to say the shape on the left is no longer the best interpretation on the right. Indeed, we interpret the local intersection as the crossing of two objects. Beyond these junctions, we are inclined to

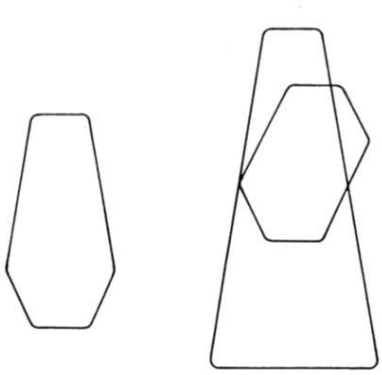

Figure 2: Masking. The figure on the left is present on the right but is not the first interpretation one usually has, since we privilege good continuations, symmetry and convexity (after Wertheimer [35])

prolongating the lines as straight as possible. Moreover, the resulting shapes are convex, and either have parallel sides or are symmetric. Put together, these properties are perceptually stronger than the first contour and make it disappear in a preattentive examination. This simple example has some important algorithmic consequences. If we run a local detection algorithm on both figures (for instance a corner detector), the detected features on the left shall also be detected (among others) on the right side, but we cannot interpret them in the same way, since new features interact with previous ones.
Even though there seems to be no way to avoid such a non local analysis, we think it necessary to first independently detect partial gestalts. This program was initiated by Desolneux, Moisan and Morel [7, 8, 9], and we shall come back on this in the following. In this paper, we are interested in the partial gestalt "good continuation". We shall also be able to define *good continuation breakings* that shall be good candidates for corners and junctions. Before going further and relating this work to existing ones, let us state our main results. We shall define a principle of good continuation allowing to find, in a set of digital curves, which are smooth or not. The method shall depend on two parameters. The first one is a scale factor which gives the size of the samples on the curves. If the curves already come from a sampling process, then this scale can be related to Shannon's rate. The second one is the average number of false detections for the whole set of curves. It shall be easy to prove that the result depends on the logarithm of this number which is therefore not a sensitive parameter. Consequently, our algorithm can be considered as "parameter free". Moreover, some perceptual arguments tend to prove that vision is widely independent of contrast. As a consequence, boundaries of objects in a digital image should correspond to the level lines of the image [4]. We try to push this argument as far as we can: if this is true, then it should be possible to detect the boundaries of objects and other geometrical structures as corners in a contrast invariant way. We shall see that this is true in a very wide extent.

The grouping of local structures in a noisy context into global ones was already studied by Guy and Medioni [15], whose purpose was to detect smooth curves by using what they called a saliency field which defines a direction and a confidence weight at any point. Points are then grouped in optimal structures which are the minimizers of a functional taking the good continuation of the curves into account. The saliency field is locally defined but the minimization is global. (See also the seminal paper by Montanari [28] for curve detection in noise). With different methods but with the same purpose, variational methods are defined in [17] and [32] to detect convex polygonal lines by grouping pieces of curves. In both cases, these features are extracted by an edge detector and grouped by a variational method. Such methods are quite general and can also be justified by Gestalt Theory. First, shapes respect a good continuation principle (or a convexity principle, which in a sense, is stronger) and are generally contrasted. Second, a common assumption is that they can be found by a variational method. Remark that the most commonly used detection methods such as image segmentation [29] or active contours [5, 20] completely follow this program. The fact that edges should be contrasted certainly finds its origin in the paradigm of the raw primal sketch of Marr [24], and the theory of edge detection [25], followed by many edge detectors among which Canny's is certainly the most famous [3]. One of the big advantage of variational methods is that they allow to select one of the best solutions to the problem. Now, there may be some false detections and it is usually necessary to introduce one or several thresholds whose value is often hard to determine automatically and are often chosen by the user. In this sense, they are not decisive algorithms. Moreover, the detection itself may depend on several parameters whose tuning is more or less ad hoc. Another common characteristic of all the methods we cited above is that they follow the edge detection doctrine, which itself seems in contradiction with the fact that image analysis should be contrast invariant. In what extent is it possible to detect the same things as these methods with a minimum number of parameters and in a way that is completely independent of contrast? What is the relative importance of contrast and good continuation? Is smoothness only a refinement to a best edge detection or is it also a first rank hypothesis? We experimentally show that

1. the hypothesis that edges are regular curve is very reasonable,
2. more unexpectedly, it is very often sufficient to define edges and no contrast information is needed!

One of the conclusions will be that contrast is one important properties *among others* that is useful to define edges. As a matter of fact, it is also known from psychophysical experiments that the geometry of contours changes the phenomenological color and contrast (see [19], Chap. 6). Contrast is not the property that masters all others since it can also be influenced by purely geometric features and the interaction between contrast and other partial gestalts is quite subtle. This does not mean that contrast is irrelevant and we do not aim at giving an edge detector uniquely based on regularity. On the contrary, if find some curves that are improbably smooth and along which the contrast is unexpectedly large, then these curves are certainly very good edges, since the probability to be mistaken is smaller than when we use

each single detector. They also resist more efficiently to masking by other gestalts. On the other hand, we shall see that regular curves and contrasted ones do not always coincide and this illustrates the need to detect all partial gestalts.
Our approach is different from Zhu's [36], who describes a procedure to learn what a "good shape" is, by examining statistics of some real shapes in a database. He then experimentally shows that random Jordan curves following these empirical distributions are potential good shapes (in the sense that they do not look so random). Conversely, a contour being given, it can be tested whether it follows the statistical laws of good shapes. Nevertheless, contours have first to be extracted and the proposed method does not cope with it.
In this paper, we shall not build a model of shapes, but use a perceptual principle called Helmholtz Principle asserting that a geometrical event is meaningful if it has a very low probability to occur by chance. Thus, the pregnance of a geometrical event is defined *a contrario* from randomness. The recent works of Desolneux, Moisan and Morel show that it can be transposed to the computational detection of simple events as alignments, modes in histograms, clusters of points or contrasted pieces of level lines [7, 8, 9, 10]. We can sum up Helmholtz Principle as follows. The purpose is to detect a structure in an image. This structure is nonlocal but is the conjunction of several local and elementary events. A certain number of these events is observed and the question is whether we should group them or not. If we assume that these atomic events are independent, we can compute the probability that they appear simultaneously, as a group. If the expected number of such a grouping is low (in general we take 1 as an upper bound) then it is considered meaningful whenever it is observed. We simply deny the independence assumption of some elementary events by showing that their conjunction is too improbable to be due only to chance. In this way, the detection threshold only depends on the size of the data, and on the precision used to compute the probabilities. Even more simply, an event is meaningful, if it has a very low probability to appear by chance in some white noise. We also point out that the structures on which we apply Helmholtz principle must be given a priori, before any observation of the image. Desolneux et al. [10] proved for some particular classes of geometrical events (alignments, clusters etc...) that the threshold for the expected number could be chosen equal to 1, and that their algorithm can then be considered as parameter free.
The plan is as follows: in Sect. 2, we define a good continuation principle for digital curves. We shall apply it to level lines of images, whose definition is recalled in Sect. 3. These lines come from a sampling process that shall be discussed in Sect. 4. In Sect. 5, we study the behavior of our detector under scale space analysis. We shall deduce, that a causality principle is satisfied and that a curve which is not meaningful at small scales does remain so at any scale. In Sect. 6, we define meaningful good continuation breakings that represent corners, junctions and terminators. We show some experiments in Sect. 7, before concluding.

2 Good continuation principle

2.1 Meaningful good continuation

Let C be a rectifiable plane curve. We sample C with a sample length equal to δ and we call p_0, ..., p_n the sampled points. At each p_i we associate the approximate direction θ_i of the tangent at p_i, computed from a chord between p_i and another curve point between p_i and p_{i+1} at fixed distance of p_i. For $1 \leq i \leq n$, we call $\Delta\theta_i = \theta_i - \theta_{i-1}$, the difference between two consecutive tangent angles. (This is also an approximation of the curvature.) Assume that C is an isotropic stationary random walk with independent increments (that is to say a discrete approximation of a Brownian motion) and that each step of the walk has a length equal to δ. Then, the $\Delta\theta_i$ are independent and identically distributed random variables, uniform in $(-\pi, \pi)$. For $\theta \in (0, \pi)$, let us compute the probability $P(\theta, n)$ that for all i, $|\Delta\theta_i| \leq \theta$. From the stationary and the independence assumption, we simply have

$$P(\theta, n) = \left(\frac{\theta}{\pi}\right)^n.$$

For a stationary random walk, this probability becomes very small, whatever θ may be, as soon as the walk is long enough. Assume that we observe a curve C with $n+1$ samples such that for all i, $|\Delta\theta_i| \leq \theta$, with $P(\theta, n)$ very small. Then there is little chance that C is a random walk. Otherwise said, there is certainly *a contrario* a better explanation to C than a simple random walk description, and this is what Helmholtz principle asserts. However, we have to precise how small $P(\theta, n)$ has to be.

Definition 1 *Let $(C_k)_k$ be a finite set of curves, and L_k be the number of samples of C_k. Let $N_c = \sum_k L_k^2$. Let Γ be a connected subcurve of one of the C_k, with $n+1$ samples and let $\theta = \max_{1 \leq i \leq n} |\Delta\theta_i|$ that we call the maximal tangent variation of Γ. We call number of false detections of Γ*

$$NF(\Gamma) = N_c \cdot P(\theta, n). \tag{1}$$

Let us comment this definition. First, N_c is the number of possible subcurves when all the C_k are closed. If they are not closed, then N_c is about twice the number of subcurves, but this difference shall be irrelevant. The number of false detections is an upper bound of the expected number of subcurves containing n samples where the tangent variation less than θ when the C_k are stationary random walks (see [9]).

Definition 2 *Let Γ be a subcurve of one the C_k. We say that Γ is a ε-meaningful good continuation if and only if*

$$NF(\Gamma) \leq \varepsilon. \tag{2}$$

As a direct consequence of the definition, a meaningful continuation must contain sufficiently many samples. More precisely, we must have

$$n \geq \frac{\log\left(\frac{\varepsilon}{N_c}\right)}{\log\left(\frac{\theta}{\pi}\right)}. \tag{3}$$

Only the logarithm of the parameters appears in this formula. In particular, ε is not as crucial as it may seem in the definition of meaningfulness (we shall check it in the experiments section). After some trials, we took $\varepsilon = 1$, and we shall see in the experiments that ε can span a large set of values with no dramatic consequences. In the same way replacing N_c by $N_c/2$ is not really important. Moreover, ε is not a completely abstract parameter since it corresponds to the average number of good continuations we observe by accident in a set of random curves. We can also interpret it this way: if an event has a number of false detections equal to ε, then we have to consider about ε^{-1} images of white noise to make it appear once. In practice, we shall detect a lot of objects. Therefore, allowing a single false detection shall not be as crucial as if there were only a few objects to be detected. Moreover, because the number of false detections bounds the expected number of good continuations from above, taking $\varepsilon = 1$ allows in practice less than one false detection.
The independence assumption of the successive tangent angle variations (since it is an approximation of the curvature, we shall also use this terminology) is clearly false, especially for most man made objects. This is exactly the aim of Helmholtz principle: we measure how bad a random walk the curve is. To this purpose, we defined a measure (in the physical meaning) on a curve, which is the size of a connected component where the variation of the tangent is less than a given threshold. A meaningful good continuation belongs to the tail of the distribution of this measure. In this way, we also do not need a precise model of what a good continuation should be, since we only use a rough model of what it should not be.

Proposition 1 *Let Γ and Γ' be two curves with maximal tangent variation θ, θ' in $(0, \pi)$. Assume that Γ and Γ' have the same number of samples. Then*

$$\theta \leq \theta' \Rightarrow NF(\Gamma) \leq NF(\Gamma'). \tag{4}$$

Let Γ and Γ' be two curves such that $|\Delta\theta| < \theta$ along Γ and Γ'. Then

$$\Gamma \subset \Gamma' \Rightarrow NF(\Gamma') \leq NF(\Gamma). \tag{5}$$

These properties directly follow from (1) but their meaning is interesting. First, when the length of a curve is fixed, a curve is more meaningful when it is as straight as possible. Second, when θ is fixed, a smooth curve is more meaningful when it is longer, and we always gain meaningfulness by adding some points with a low curvature. On the other hand, we see that whatever θ may be, a curve may become meaningful if it is long enough. A possible solution is to fix a maximal value of θ. We shall see that by forbidding curves with right angles, many curves shall be eliminated. Another consequence of the properties above is that a very

meaningful curve may contain a lot of smaller meaningful ones. These curves are in general masked by the largest one (this is the simplest manifestation of masking since it occurs within a single partial gestalt). It seems reasonable to eliminate these curvelets, and in the same time, this reduces the amount of stored information. The following definition, classically adapted from [7], formalizes this masking phenomenon.

Definition 3 *Let Γ be a ε-good continuation in a curve C. We say that Γ is maximal meaningful, if*

- $\forall \Gamma'$ *ε-meaningful,* $\Gamma' \subseteq \Gamma \Rightarrow NF(\Gamma) \leq NF(\Gamma')$.
- $\forall \Gamma'$ *ε-meaningful,* $\Gamma \subsetneq \Gamma' \Rightarrow NF(\Gamma) < NF(\Gamma')$.

This means that a maximal curve does not contain a strictly more meaningful curve and is not strictly included in a more meaningful one.

Proposition 2 *Let Γ_1, Γ_2 be two different maximal meaningful good continuations on the same curve. Then $\Gamma_1 \cap \Gamma_2 = \emptyset$.*

Proof. Assume that Γ_1 and Γ_2 are ε-maximal meaningful and that $\Gamma_1 \cap \Gamma_2 \neq \emptyset$. Let us call θ_1 and θ_2 their respective maximal tangent variation and $\theta = \max(\theta_1, \theta_2)$. Then $\Gamma = \Gamma_1 \cup \Gamma_2$ is connected, and by Proposition 1 it is also ε-meaningful since

$$NF(\Gamma) \leq \max(NF(\Gamma_1), NF(\Gamma_2)) \leq \varepsilon.$$

Since Γ_1 and Γ_2 are maximal meaningful, $\Gamma_1 = \Gamma_2 = \Gamma_1 \cup \Gamma_2$. □

3 Topographic map

In what follows, we apply the good continuation principle to curves in grey level images. It is natural to use the topographic map which provides a complete representation of the image [4]. Let u be a gray level image. For $\lambda \in \mathbb{R}$, we consider the upper level set

$$\chi_\lambda(u) = \{x \text{ s.t. } u(x) \geq \lambda\},$$

and the lower level-set

$$\chi^\lambda(u) = \{x \text{ s.t. } u(x) \leq \lambda\}.$$

The topographic map of u is the collection of the level lines that are the boundaries of connected components of level sets. There are upper and lower level lines, depending on the fact that they are boundaries of upper or lower level sets. The level lines, indexed by the associated gray level, entirely determine u. More interestingly, they can be embedded in a tree structure that determines u up to a monotone contrast change [14]. This means that the topographic map is independent of contrast. We shall apply the good continuation detection on level lines and shall forget any gray level information, keeping only their geometry. The tree

of level lines can be computed by the Fast Level Set Transform [23, 27]. In practice, the gray level takes a finite number of values and we can choose the quantization step such that the topographic map densely covers the image. If needed, new gray levels are created by bilinear interpolation of the original data. This method does not create new branches in the tree structure since it satisfies a maximum principle. Bilinear interpolation also partially removes the pixellization effect and the following section details the implications of interpolation.

4 Sampling and interpolation

We shall directly apply the definition of good continuation to level lines of images. We consider samples of length δ varying between 1 and 2, and estimate the direction of the tangent on a length l less or equal to δ. In practice, we always choose $l = \delta$ since we experimentally checked that l had indeed a small influence. This means that we estimate the curvature on pieces of curve of length from 2 to 4. The process is not completely invariant with respect to translations along the curve since the starting point of the parameterization may have some influence. When we extract level lines, we can choose the approximate number of points per pixel, and we usually take it about 5. We compute the approximate curvatures by changing the starting point, and apply the good continuation principle to the largest value when the offset varies.
Bilinear interpolation reduces the effect of the pixellization. Instead of horizontal and vertical segments, level lines become concatenations of pieces of hyperbolae, see Fig. 3. In particular, a corner at right angle produces level lines whose tangent varies less rapidly. Moreover, the

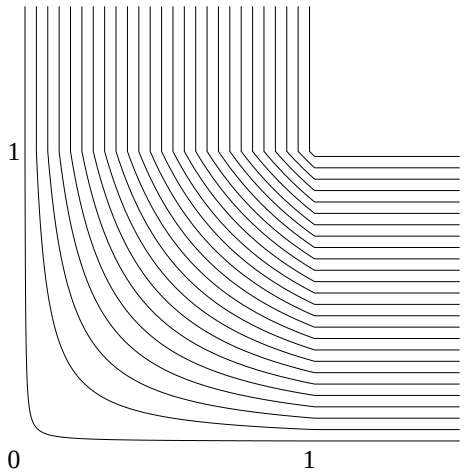

Figure 3: Level lines in a corner after bilinear interpolation. The interpolated level lines are pieces of hyperbolae and the variation of the tangent are less brutal than a right angle

sampling and the approximation of the tangent by chord angles is equivalent to regularize the

tangent angle. As a consequence, an irregular curve may appear smooth at some intermediary scale. We admit the following axiom: "a right angle cannot be a good continuation". Therefore, we must choose test values of θ small enough so that a sampled and interpolated right angle cannot be detected as a good continuation. On Fig. 3 are plotted the level lines in a right angle in the square $(0,1)^2$. None of these lines should be a good continuation, and in particular the one passing through the point $(\frac{1}{2},\frac{1}{2})$. A simple computation shows that the maximal value of $\Delta\theta$ on this line is approximately given by

$$\theta_{max} = \arccos\left(0.5 \cdot \frac{-0.066+\delta}{0.062+(-0.066+\delta)^2}\right),$$

and we can reasonably take $\theta_{max} = \arccos(\frac{1}{2\delta})$. This function is increasing with respect to δ and is equal to $\frac{\pi}{3}$ for $\delta = 1$ and 1.3 for $\delta = 2$. This limits the value of $\Delta\theta$ in the research of good continuations.

5 Detection performance and shape scale space

We now focus on the choice of the sample length δ, which plays the role of a scale parameter. A classical solution would be to perform a multiscale analysis and define a good continuation as a curve which satisfies our principle for some value of the scale. Of course, the number of tested scales should be incorporated in the definition of the number of false detections. This still leaves a range of possible scales. If we want to detect the smallest structures, then we should choose δ as small as possible, but it cannot be taken arbitrarily small. Indeed, in order to apply Helmholtz principle, we have to compute the probability that an event (in our case good continuation) occurs in a set of random curves. Alternatively, we can compute the probability that it occurs in some white noise. The digital model of white noise is a set a i.i.d values on a grid. By Shannon's theory [33], we can associate to this set of samples a unique band-limited function whose frequencies are below twice the sampling rate, so that a white noise can then be viewed as an analytic function! By using this interpolation, it is obvious that if we look some interpolated digital image at very low scales, it is as smooth as we want (and so are almost all its level lines by Sard's theorem). Obviously, this does not make sense, and we cannot interpret the data at arbitrary small scales. Therefore, the size of the samples has to be related to Shannon's rate. In the previous paragraph, we used the bilinear interpolation but the conclusions remain unchanged. Another way to see this, is to reexamine Helmholtz principle. We want that an image of white noise does not contain more than one meaningful structure by chance. If we now smooth a white noise, this has to remain true since the smoothing should not create meaningful structures in the image (causality principle). On the other hand, since the smoothing has killed high frequencies, the samples we consider must be longer, and their size has to be related to the size of the smoothing "kernel". (This denomination may be inappropriate since we can use a nonlinear smoothing.) The arguments we develop in this section lead to this conclusion: we can choose a sample length only depending on the smoothing scale such that we can apply Helmholtz principle at any

scale for the same definition of good continuations. (i.e. we assume that the curves are trials of stationary random walks with a given sample length.) Moreover, a random curve which is not a good continuation at the initial scale, does not become a good continuation at larger scales. This yields that it is useless to smooth curves before detecting good continuations.
One of the simplest and most intrinsic curve smoothing method is the mean curvature motion [11, 13, 1, 31], which is described by the partial differential equation

$$\frac{\partial C}{\partial t} = \frac{\partial^2 C}{\partial s^2}, \tag{6}$$

where t is the smoothing parameter and s the arc-length parameter of the curve at scale t. Let us call θ the angle of the tangent with respect to a fixed direction. We use the following lemma proved for instance in [11].

Lemma 1 *Let Γ_t evolving by the mean curvature motion. Then θ satisfies the same equation*

$$\frac{\partial \theta}{\partial t} = \frac{\partial^2 \theta}{\partial s^2}. \tag{7}$$

These intrinsic heat equations (6) and (7) are nonlinear because the length parameter s depends on t. (In particular, the partial derivatives with respect to s and t do not commute.) Thus, it is impossible to directly solve (7) above, since it first requires to determine the arc-length parameter which can be found only by solving (6). Nevertheless, in a first approximation, we can neglect this dependence, and the equation becomes the classical linear heat equation. This can be justified by the fact that the nonlinear heat equation can be asymptotically retrieved by reparameterizing the curve at regular scale step tending to 0. (See [26].) Thus, in the following, we shall use (7) as if t and s were independent variables. We shall denote the parameterization by u (belonging to a fixed interval) to remind that u is not the arc-length. We can then define a curve scale space as follows.

1. Parameterize the curve by its arc-length.
2. Compute the tangent angle θ as a function of the arc-length.
3. Apply the linear heat equation to θ.

The smoothed curve is retrieved up to a isometry by integrating θ. Contrary to the mean curvature motion, this method leaves the length unchanged, but a closed curve does not remain closed in general. (We shall not be interested on border effects.) At least, this evolution is consistent with the mean curvature motion for small scales. For infinite random curves, we apply the heat equation to the angle variation between two fixed points to avoid problems of definition.

5.1 Angle distribution

Let Γ be a random walk with unit length samples, that is a continuous random process in the plane, such that, for all $n \in \mathbb{Z}$ the tangent to Γ in $(n, n+1)$ is equal to $(\cos X_n, \sin X_n)$ where the X_n are i.i.d variables, uniform in $(-\pi, \pi)$. Let u be the arc-length at time $t = 0$, and θ the angle of the tangent with a fixed direction. Let make θ evolve by

$$\frac{\partial \theta}{\partial t} = \frac{\partial^2 \theta}{\partial u^2}.$$

It is well known that the solution of the heat equation at time t is obtained by the convolution of the initial data with the gaussian with variance $2t$. In what follows, we shall denote by G_σ the gaussian with variance σ^2, that is

$$G_\sigma(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{x^2}{2\sigma^2}}.$$

Let

$$\begin{aligned} \Delta\theta(s, d, 0) &= \theta\left(s + \frac{d}{2}, 0\right) - \theta\left(s - \frac{d}{2}, 0\right), \\ &= \sum_{k \in \mathbb{Z}} \mathbf{1}_{[k-\frac{d}{2}, k+\frac{d}{2})}(s)\, X_k. \end{aligned}$$

The variation of the tangent angle between the points with parameters $s - \frac{d}{2}$ and $s + \frac{d}{2}$ is $\Delta\theta(s, d, 0) \pmod{2\pi}$. By linearity, for all d, $\Delta\theta(s, d, t)$ (the variation at time t) is solution of the heat equation. This allows to prove the following result, asserting that the "curvature" distribution concentrates around 0.

Proposition 3 *Let θ be solution of the heat equation. Then*

$$Var(\Delta\theta(s, 1, t)) \leq \frac{1}{6}\left(\frac{\pi^3}{t}\right)^{1/2}. \tag{8}$$

Proof. Let us apply the heat equation to $\Delta\theta$. To simplify the notations, we set $\sigma^2 = 2t$. By linearity, we obtain

$$\Delta\theta(s, 1, t) = \sum_{k \in \mathbb{Z}} s_k X_k \quad \text{with} \quad s_k = \frac{1}{\sigma\sqrt{2\pi}} \int_{k-1/2}^{k+1/2} e^{-\frac{u^2}{2\sigma^2}}\, du. \tag{9}$$

Since the X_k are independent with variance $\frac{\pi^2}{3}$, the variance of $\Delta\theta$ is

$$Var\Delta\theta(0, 1, t) = \frac{\pi^2}{3} \sum_{k \in \mathbb{Z}} s_k^2. \tag{10}$$

Moreover,

$$\begin{aligned} 2\pi\sigma^2 \sum_{k\in\mathbb{Z}} s_k^2 &= \sum_{k\in\mathbb{Z}} \left(\int_{k-1/2}^{k+1/2} e^{-\frac{u^2}{2\sigma^2}}\, du \right)^2, \\ &\leq \sum_{k\in\mathbb{Z}} \int_{k-1/2}^{k+1/2} e^{-\frac{u^2}{2\sigma^2}}\, du \quad \text{since each term is smaller than 1,} \\ &= \sigma\sqrt{2\pi}. \end{aligned}$$

By using (10), this inequality gives the result. □

In fact, $\sqrt{t}\Delta\theta(s,1,t)$ certainly converges in law toward a normal distribution, but we shall not use it explicitly. The result above only means that, as expected, the heat equation has smoothed the approximated Brownian motion and the prior distribution of the angle variation on unit length samples is no longer uniform in $(-\pi,\pi)$. Then, if we want to apply Helmholtz principle, we should compute large variations with respect to this new distribution and not to the uniform one. More importantly, the samples are no longer independent. In the following, we shall see that a possibility to retrieve independent variations of θ uniformly distributed in $(-\pi,\pi)$ is to sample the curve at a lower rate. Indeed, downsampling the curve is equivalent to consider $\Delta\theta(s,d,t)$ with $d>1$. This is also equivalent to add d terms of the type $\Delta\theta(s_i,1,t)$ and take the result modulo 2π. We have seen that each of these terms has a variance proportional to $t^{-1/2}$. Thus we can hope that if we sum sufficiently many terms, the variance shall be large enough to compensate the smoothing effect and the correlation between samples.

Definition 4 *Let X be a real random variable. We denote by $\widetilde{X}$ the periodized value of X in $(-\pi,\pi)$ defined as the unique random variable in $(-\pi,\pi)$ such that $\widetilde{X} \equiv X \pmod{2\pi}$.*

We shall only consider variables defined by density distribution and shall denote by f_X the density of X.

Proposition 4 *Let X be a random variable and $\widetilde{X}$ its periodized variable. Then*

$$f_{\widetilde{X}}(x) = \frac{1}{2\pi} \sum_{n\in\mathbb{Z}} \widehat{f_X}(n) e^{inx} \mathbf{1}_{(-\pi,\pi)}(x). \tag{11}$$

Proof. If I is an interval included in $(-\pi,\pi)$, then

$$P(\widetilde{X}\in I) = \sum_{k\in\mathbb{Z}} P(X \in I + 2k\pi),$$

since the intervals $I+2k\pi$ are disjoint. Thus, for $x\in(-\pi,\pi)$

$$f_{\widetilde{X}}(x) = \sum_{k\in\mathbb{Z}} f_X(x+2k\pi).$$

By applying Poisson sommation formula, we obtain for $x \in (-\pi, \pi)$

$$f_{\widetilde{X}}(x) = \frac{1}{2\pi} \sum_{k \in \mathbb{Z}} \widehat{f_X}(n) e^{inx}. \qquad \square$$

We now take samples with length equal to $2\sigma + 1$ (with still $\sigma = \sqrt{2t}$). This is equivalent to replace the angle variation $\Delta\theta(s, 1, t)$ by $\sum_{j=-\sigma}^{\sigma} \Delta\theta(s + j, 1, t) \pmod{2\pi}$. We denote by

$$Y_p = \sum_{-\sigma \leq j \leq \sigma} \Delta\theta((2\sigma + 1)p + j, 1, t) \tag{12}$$

these values obtained by sampling the curve at multiples of $2\sigma + 1$.
In the following, we shall write $Y \xrightarrow{\mathcal{L}} X$ if Y converges in law to a variable X.

Proposition 5 *Let $\widetilde{Y_p}$ the periodized value of Y_p in $(-\pi, \pi)$. Then, when $\sigma \to +\infty$, $\widetilde{Y_p} \xrightarrow{\mathcal{L}} U_p$ where U_p is a uniform variable in $(-\pi, \pi)$.*

Proof. By translation invariance, we can take $p = 0$ and we shall write Y instead of Y_0. By definition of Y, we have

$$Y = \sum_{-\sigma \leq j \leq \sigma} \sum_{k \in \mathbb{Z}} s_k X_{j+k},$$

where s_k is as in (9). We can rewrite this sum as

$$Y = \sum_{k \in \mathbb{Z}} w_k X_k \quad \text{where } w_k = \sum_{-\sigma \leq j \leq \sigma} s_{k-j}. \tag{13}$$

The new weights w_k are also given by the expression

$$w_k = \int_{\frac{k-\sigma-1/2}{\sigma}}^{\frac{k+\sigma+1/2}{\sigma}} G_1(u)\, du, \tag{14}$$

where G_1 is still the gaussian with unit mass and variance 1. In this expression, we have made the assumption that σ was an integer, in order not to write integer values in the integral bounds, but it only aims at simplifying the notations, and does not change the result for large values of σ. Whereas s_k tend to 0 with $h \to +\infty$, this is not the case of w_k for a larger and larger set of values of k, roughly corresponding to $-\sigma \leq k \leq \sigma$. As a consequence, Y can take arbitrary large values when σ becomes large. After periodization, these values shall uniformly fill the interval $(-\pi, \pi)$. More precisely, by Prop. 4

$$f_{\widetilde{Y}}(x) = \frac{1}{2\pi} \sum_{n \in \mathbb{Z}} \widehat{f_Y}(n) e^{inx} \mathbf{1}_{(-\pi,\pi)}(x),$$

hence

$$\widehat{f_{\widetilde{Y}}}(\xi) = \sum_{n \in \mathbb{Z}} \widehat{f_Y}(n) \operatorname{sinc}(\pi(\xi - n)). \tag{15}$$

But since the X_k are i.i.d, uniform in $(-\pi, \pi)$, we also have

$$\widehat{f_Y}(\xi) = \prod_{k\in\mathbb{Z}} \operatorname{sinc}(\pi w_k \xi). \tag{16}$$

We are going to prove that for all ξ, $\widehat{f_{\widetilde{Y}}}(\xi)$ tends to $\operatorname{sinc}\xi$ which is the Fourier transform of $\frac{1}{2\pi}\mathbf{1}_{(-1,1)}$. To this purpose, we calculate bounds for $\widehat{f_Y}(n)$ if $n \neq 0$ and when σ tends to $+\infty$. For k such that $-\sigma \leq k \leq \sigma$, we have $\pi w_k \geq \int_0^2 G_1(u)\,du = \beta > 1$. Let n, with $|n| > 1$. Then, for $-\sigma \leq k \leq \sigma$,

$$|\operatorname{sinc}(\pi w_k n)| \leq \frac{1}{\pi w_k n} \leq \frac{1}{\beta n}.$$

Hence, for $|n| \geq 1$,

$$|\widehat{f_Y}(n)| \leq \frac{1}{(\beta n)^{2\sigma+1}}.$$

This implies

$$\begin{aligned} |\widehat{f_{\widetilde{Y}}}(\xi) - \operatorname{sinc}(\xi)| &\leq \sum_{n\neq 0} |\widehat{f_Y}(n)| \\ &\leq \sum_{|n|>0} \frac{1}{(\beta n)^{2\sigma+1}}. \end{aligned}$$

We bound this sum by

$$\sum_{|n|>0} \frac{1}{(\beta n)^{2\sigma+1}} \leq 2\int_1^\infty \frac{du}{(\beta u)^{2\sigma+1}} = \frac{1}{\sigma\beta^{2\sigma+1}},$$

and, since $\beta > 1$, this term tends to 0 when $\sigma \to +\infty$. Hence,

$$\forall \xi \in \mathbb{R}, \quad \lim_{\sigma\to\infty} \widehat{f_{\widetilde{Y}}}(\xi) = \operatorname{sinc}(\xi).$$

This exactly means that $\widetilde{Y}$ converges in law to a uniform variable in $(-\pi, \pi)$. □

This result can be checked experimentally. We take a random walk with 2 million samples and apply a convolution with standard deviation $\sigma = 5$. As can be seen in Fig. 4, if we sample the resulting curve with the same rate, then the angle distribution is clearly nonuniform since small values are privileged. As expected, if we now sample the smoothed curve at larger rate, the distribution is nearly uniform.

5.2 Samples independence

In the previous section we saw that the angle variations of a smoothed random walk approximately followed a uniform distribution in $(-\pi, \pi)$ as soon as the sample length is suitably

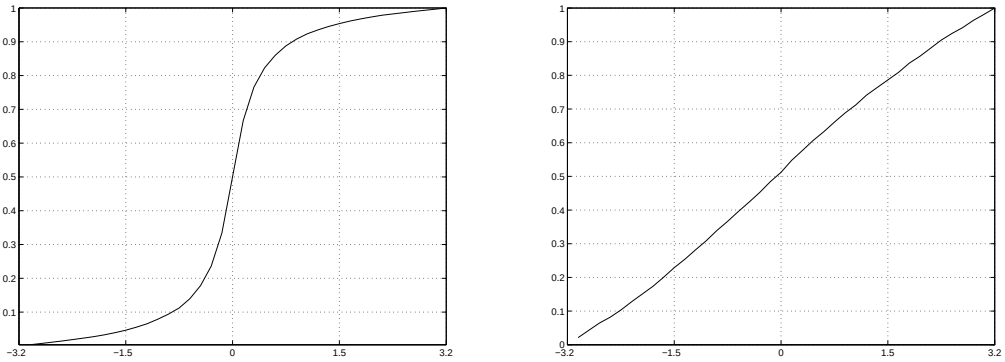

Figure 4: We applied the linear heat equation (convolution with a gaussian with standard deviation equal to 5) to a random walk of length $2 \cdot 10^6$. Left: empirical repartition function of angle variation with sample length equal to 1. Right: same thing with sample length equal to 11, the distribution is nearly uniform.

chosen. In order to apply Helmholtz principle, we need to prove that samples are also independent. The intuition is that if we take very long samples, then they are more and more independent, because even if the gaussian kernel is not compactly supported, its decay is exponential. Thus, the problem is to choose the minimal length such that this is still true. In view of Prop. 3, we cannot take less than $O(\sqrt{t})$ samples. In this section, we prove that this number allows to retrieve a quasi-independence of samples. It is sufficient to prove it for two consecutive increments, the same proof can be applied in the general case. Let Y_0 and Y_1 be two consecutive angles variations of the smoothed subsampled curve. By using the same notations as above, we have

$$Y_0 = \sum_{k \in \mathbb{Z}} w_k X_k, \quad Y_1 = \sum_{k \in \mathbb{Z}} w_k X_{k+2\sigma+1},$$

with

$$w_k = \int_{\frac{k-\sigma-1/2}{\sigma}}^{\frac{k+\sigma+1/2}{\sigma}} G_1(u)\, du.$$

Proposition 6 *When $\sigma \to +\infty$, $\widetilde{Y_0} \xrightarrow{\mathcal{L}} U_0$ and $\widetilde{Y_1} \xrightarrow{\mathcal{L}} U_1$, where U_0 and U_1 are two independent uniform variables in $(-\pi, \pi)$.*

Proof. By Prop. 5 above, we already know that $\widetilde{Y_0} \xrightarrow{\mathcal{L}} U_0$ and $\widetilde{Y_1} \xrightarrow{\mathcal{L}} U_1$, where U_0 and U_1 are uniform in $(-\pi, \pi)$. The goal is here to prove that U_0 and U_1 are independent. First, $f_{\widetilde{Y_0},\widetilde{Y_1}}$ is given by

$$f_{\widetilde{Y_0},\widetilde{Y_1}}(x, y) = \sum_{(k,l) \in \mathbb{Z}^2} f_{Y_0,Y_1}(x + 2k\pi, y + 2l\pi) \mathbf{1}_{(-\pi,\pi)^2}(x, y).$$

By using Poisson formula in two dimensions, we obtain

$$f_{\widetilde{Y_0},\widetilde{Y_1}}(x,y) = \frac{1}{2\pi} \sum_{(k,l)\in\mathbb{Z}^2} \widehat{f_{Y_0,Y_1}}(k,l) e^{i(kx+ly)} \mathbf{1}_{(-\pi,\pi)^2}(x,y). \tag{17}$$

Since Y_0 and Y_1 are not independent, their joint distribution is not simply the tensorial product of their marginal density. However

$$\begin{aligned}
\widehat{f_{Y_0,Y_1}}(k,l) &= E\left(e^{i(kY_0+lY_1)}\right) \\
&= E\left(e^{i\sum_{n\in\mathbb{Z}}(kw_n+lw_{n-2\sigma-1})X_n}\right) \\
&= \prod_{n\in\mathbb{Z}} \operatorname{sinc}(\pi(kw_n + lw_{n-2\sigma-1}))
\end{aligned}$$

since the X_k are i.i.d. and uniform in $(-\pi,\pi)$. Let us denote by $a(k,l)$ these coefficients (that depend on σ). The characteristic function of $(\widetilde{Y_0},\widetilde{Y_1})$ is given by

$$\widehat{f_{\widetilde{Y_0},\widetilde{Y_1}}}(x,y) = \sum_{(k,l)\in\mathbb{Z}^2} a(k,l) \operatorname{sinc}(\pi(x-k)) \operatorname{sinc}(\pi(y-l)).$$

We prove that for all $(x,y)\in\mathbb{R}^2$, this sum converges to the unique term corresponding to $a(0,0)=1$. If $k=0$ or $l=0$, we can apply Prop. 5 to obtain an estimate on $a(k,l)$ as a function of σ.
Let first k and l be positive. For any n such that $-\sqrt{\sigma} \le n \le \sqrt{\sigma}$,

$$w_n = \int_{-1}^{1} G_1(u)\, du + O\left(\frac{1}{\sqrt{\sigma}}\right),$$

with an error term independent of n. Since, $\pi \int_{-1}^{1} G_1(u)\, du > 1$, if σ is large enough, we also have

$$\forall n, \quad 0 \le n \le \sigma, \quad \pi w_n > \alpha > 1,$$

where α can be chosen uniformly over n. In the same way, for $2\sigma \le n \le 2\sigma+\sqrt{\sigma}$, we have $\pi w_{n-2\sigma-1} \ge \alpha > 1$. Hence, since k and l are positive,

$$\begin{aligned}
|a(k,l)| &= \left| \prod_{n\in\mathbb{Z}} \operatorname{sinc}(\pi(kw_n + lw_{n-2\sigma-1})) \right| \\
&\le \left(\prod_{n=0}^{\sqrt{\sigma}} \frac{1}{kw_n + lw_{n-2\sigma-1}} \right) \cdot \left(\prod_{n=2\sigma}^{2\sigma+\sqrt{\sigma}} \frac{1}{kw_n + lw_{n-2\sigma-1}} \right) \qquad (18) \\
&\le \frac{1}{(kl\alpha^2)^{\sqrt{\sigma}}}. \qquad (19)
\end{aligned}$$

Inequality (19) implies

$$\lim_{\sigma \to +\infty} \sum_{k,l=1}^{\infty} |a(k,l)| = 0. \tag{20}$$

By symmetry, these arguments are also valid for k and l negative. Let us now assume that $k > 0$ and $l < 0$. We cannot use (18) directly since the denominator is not the sum of two positive terms. However, for $0 \leq n \leq \sqrt{\sigma}$, $w_{n-2\sigma-1}$ tends to $\int_{-3}^{-1} G_1(u)\, du$, the error being uniform with respect to n. Assume now that $k \geq -l$. Then

$$kw_n + lw_{n-2\sigma-1} \geq k(w_n - w_{n-2\sigma-1}).$$

Since $\pi \left(\int_{-1}^{1} G_1(u)\, du - \int_{-3}^{-1} G_1(u)\, du \right) > 1$, we can choose σ large enough such that for all n such that $0 \leq n \leq \sqrt{\sigma}$,

$$\pi(w_n - w_{n-2\sigma-1}) \geq \beta > 1.$$

Therefore

$$\begin{aligned}
\sum_{0<-l\leq k} |a(k,l)| &= \sum_{0<-l\leq k} \prod_{n\in\mathbb{Z}} |\operatorname{sinc}(\pi(kw_n + lw_{n-2\sigma-1}))| \\
&\leq \sum_{0<-l\leq k} \prod_{0\leq n\leq\sqrt{\sigma}} |\operatorname{sinc}(\pi(kw_n + lw_{n-2\sigma-1}))| \\
&\leq \sum_{0<-l\leq k} \prod_{0\leq n\leq\sqrt{\sigma}} \left| \frac{1}{\pi(kw_n + lw_{n-2\sigma-1})} \right| \\
&\leq \sum_{0<-l\leq k} \prod_{0\leq n\leq\sqrt{\sigma}} \frac{1}{\pi(k(w_n - w_{n-2\sigma-1}))} \\
&\leq \sum_{0<-l\leq k} \frac{1}{(k\beta)^{\sqrt{\sigma}}} \\
&\leq \sum_{k>0} \frac{k}{(k\beta)^{\sqrt{\sigma}}}.
\end{aligned}$$

This sum tends to 0 when σ tends to $+\infty$. We can fix the case $-l \geq k$ in the same way, by considering the w_n for $2\sigma \leq n \leq \sqrt{\sigma} + 2\sigma + 1$. By symmetry, we can also apply the same arguments for $k < 0 < l$. Finally, we conclude that for $(x, y) \in \mathbb{R}^2$,

$$\lim_{\sigma \to +\infty} \widehat{f_{\widetilde{Y_0},\widetilde{Y_1}}}(x,y) = \frac{1}{(2\pi)^2} \operatorname{sinc}(\pi x) \operatorname{sinc}(\pi y),$$

which is the characteristic function of two independent variables, uniform in $(-\pi, \pi)$, and this concludes the proof. □

Let us sum up the results of both previous sections. Let Γ be an isotropic random walk with unit length samples and angle variation uniform in $(-\pi, \pi)$. Let u be the arc-length of the curve and θ the angle of the tangent with a fixed direction. We then apply the heat equation to θ, i.e. we solve

$$\frac{\partial \theta}{\partial t} = \frac{\partial^2 \theta}{\partial u^2}.$$

For any $t > 0$, we obtain a new curve whose arc length is u and the tangent angle is θ at scale t. Then, if we sample this new curve at unit rate, the angle variation are more and more concentrated around 0, meaning that the curve has been smoothed. Moreover, angle variations at different sample are correlated. Hence, if we want to apply Helmholtz principle to this model, we have to take those two factors into account. A simpler method consists in subsampling the curves until their statistics get close to the ones of a stationary random walks with independent increments. We then saw that the smallest possible sampling rate was proportional to $\sqrt{t}$, and we can consider that the new curve is still a random walk with independent increments and angle variations uniform in $(-\pi, \pi)$. We have rigorously proved our result for a shape scale space that is not one of the most classical and that does not have the nicest properties of regularity and invariance. Moreover, the independence is obvious if we apply a convolution with a compactly supported kernel. However, it is well known that an iterated convolution is equivalent to a convolution with a gaussian, and this makes the case of the gaussian central. We do not know how to lead the same kind of calculations for intrinsic evolutions, but the results above always make sense in a qualitative point of view. One of the main difference is that the relation between the scale and the sample length is certainly not as simple as in our linear case, but it can be evaluated empirically and once for all on random curves. This sample length should be chosen such that the downsampled curve has still increments that are approximately uniform in $(-\pi, \pi)$. In the calculations above, we saw that this also implied (for our model of smoothing) that the increments were also nearly independent. For general scale spaces, independence should be tested on subsampled random walks after other scales spaces (specially intrinsic ones) to check if this independence hypothesis is still sound.

5.3 Influence of the number of samples

In the previous section, we saw that a subsampling of curves after smoothing still allowed to apply Helmholtz principle by comparing the empirical distribution of the tangent angle along the smooth curve with the uniform distribution on $(-\pi, \pi)$. However, this subsampling decreases the total number of subcurves we consider (since it is inversely proportional to the square of the sample length). This tends to decrease the number of false detections of a curve. On the other hand, a curve contains less samples after downsampling, and as the number of false detection is a product, this prevents the product from being very small. In this section, we show that the second effect (increase) is always stronger than the first one. We should also take in consideration the fact that the curve length also decrease during the evolution. If we

assume that the length of all curves have the same rate of decay, then we can incorporate it in the sample size.

Proposition 7 *Assume that the angle variation distribution is uniform in* $(-\pi, \pi)$ *for all sample length. Let* $NF(d)$ *the number of false detections of a curve* Γ *sampled at rate* $d > 1$. *Then, if* $d < \frac{L}{4}$, $NF(d) > NF(1)$.

Proof. The assumption that the angle distribution is uniform in $(-\pi, \pi)$ at all scale is licit thanks to the results of the previous section. Let us study $N_\alpha(d) = d^{-2}P(|\Delta\theta| < \alpha \text{ along } \Gamma)$ for $d > 1$ and $\alpha \leq \frac{\pi}{2}$. (We still forbid right angles.) Since the distribution of the angles are uniform on $(-\pi, \pi)$,

$$\rho(d) \equiv \frac{NF(d)}{NF(1)} = \frac{N_\alpha(d)}{N_\alpha(1)} = \left(\frac{\alpha}{\pi}\right)^{L(\frac{1}{d}-1)} d^{-2}.$$

Thus

$$(\ln \rho)'(d) = \frac{1}{d^2}\left(L \ln \frac{\pi}{\alpha} - 2d\right),$$

which is positive for $d < \frac{L}{2} \ln \frac{\pi}{\alpha}$. If $d < \frac{L}{4} < \frac{L}{2} \ln 2$, the number of false detections cannot decrease. We shall never subsample a curve at this rate since the downsampled curve would only contain four points! □

5.4 Experimental verification

In order to illustrate these calculations, we tested the good continuation detections on an image of white noise progressively smoothed by the mean curvature motion with time equal to $t = 0$, 1, 2 and 3, with corresponding sample lengths $1 + 2\sqrt{2t} = 1$, 3.8, 5 and 5.9. The first false detections appeared for ε respectively equal to $10^{2.0}$, $10^{2.8}$, $10^{2.2}$ and $10^{1.1}$. In these experiments, we can choose the sample size since we know the sampling rate in the original image and the amount of smoothing. If an image is naturally blurred, then there is no a priori reason to choose large curve samples, and some good continuations will then be detected. This is not a failure of the method, since the features are indeed present. The true conclusion, is that we need a further description of the image that shall give the cause of the detected features. For instance, edges will contain a very large number of parallel good continuations and the width of edges can give and indication of the amount of blurring [21] related to the "true scale" of the image as an input of good continuation detection. Another interesting experiments is to compute the number of good continuations in a white noise as a function of ε, and to check that it grows linearly in $\log \varepsilon$.

6 Corners, junctions and terminators

The detection of corners and junctions have also been a topic of interest for more than twenty years. Corners are usually geometrical strong cues of closed shapes, whereas junctions appear at occlusions where the boundaries of two objects meet. A further study of the configuration

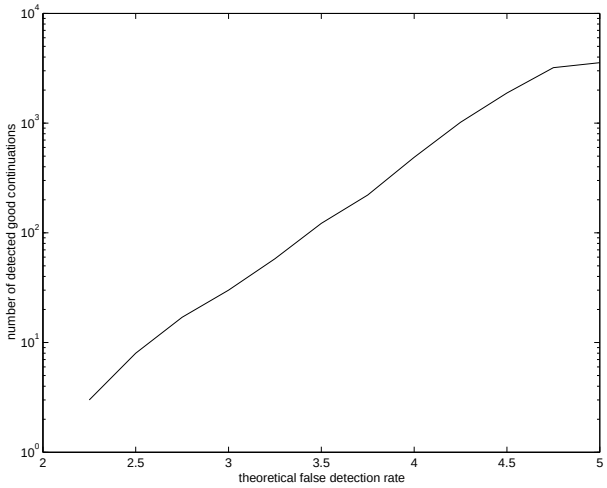

Figure 5: Number of detected good continuations in white noise as a function of ε (log scale). Even though the level lines are not random walks, the number of detection seems to be also linear as a function of $\log \varepsilon$

of a junction gives a hint on the local relative depth of objects. Most used algorithms consider that corners and junctions are points where the normal of the boundaries of objects varies fastly. As a consequence, many authors have tried to define a "cornerness" density function whose maximum points are corners. Harris and Stephens [16] took the points where both eigenvalues of $G_\sigma * (Du\, Du^t)$ are large, where G_σ is a gaussian smoothing with standard variation equal to σ. Without gaussian smoothing ($\sigma = 0$), the orthogonal to Du is an eigen direction is associated to the eigenvalue 0. Since at a corner, this direction varies much, the gaussian blurring mixes these directions and the matrix may have two nonnull eigenvalues. Other classical methods define corners as points where a function of the curvature of level lines is maximal (see for instance, the paper of Deriche and Giraudon giving a review of the first corners extraction algorithms with an extension to junctions [6]). Unfortunately, curvature is a second order differential operator and, as such, is sensitive to noise. To eliminate most false detections, a preliminary multiscale smoothing (generally a gaussian blurring) is applied, with the side effect to move corners. One can also impose that corners and junctions belong to edges detected by classical algorithms (see for instance in [22]), thus making the corner detection more robust since non purely local.
Our approach will be prolongate this last argument, since we define corners and junctions by considering that the tangent angle of the curve must be significantly different on each side of the corner and that the curve is flat enough on each side. We then impose that a corner is close to a good continuation. Our main purpose is to understand the level at which a corner can be defined by geometrical and contrast independent information alone, and with no threshold on the curvature.

6.1 Good continuation breaking

Let us consider the model of corner made by two non collinear segments with a common endpoint. Let us denote by C the resulting (sampled) curve and assume that the corner corresponds to parameter i_0. Let us examine the tangent angle θ_i at the point $C(i_0 - i)$. For a perfect corner, there are two distinct values α_1 and α_2 such that for negative i, $\theta_i = \alpha_1$, and for positive i, $\theta_i = \alpha_2$. If we now compute the histogram of the θ_i (for $i \neq 0$), we obtain two Dirac masses at the values α_1 and α_2. If the curve is a bit noisy and the two subcurves are not completely flat then, we expect that we shall have two peaks around the values α_1 and α_2 and a gap between them. We shall only impose that the gap between the values θ_i for both $i < 0$ and $i > 0$ is large enough. In [7], Desolneux proposed to detect peaks and gaps in histograms by a method based on numbers of false detections. The method we propose for corner detection is a particular case of her work.
Let

$$\alpha_1(l) = \min_{-l \leq i < 0} \theta_i, \quad \beta_1(l) = \max_{-l \leq i < 0} \theta_i$$

and

$$\alpha_2(l) = \min_{0 < i < l} \theta_i, \quad \beta_2(l) = \max_{0 < i < l} \theta_i,$$

where the angles are measured relatively to the tangent at the corner point ($C(i)$. Assume that $(\alpha_1, \beta_1) \cap (\alpha_2, \beta_2) = \emptyset$. In this case, we can also suppose that $\beta_1 \leq \alpha_2$. Let us denote

$$r(l) = \frac{\alpha_2 - \beta_1}{\beta_2 - \alpha_1}$$

which is the relative size of the gap in the set of the spanned angle.

Definition 5 *Consider a set of curves, with a total number of samples equal to N_s. Let C a curve in this set and $C(i)$ a point of this curve. We say that $C(i)$ is a ε-meaningful breaking if and only if for some $l > 1$,*

- $(\alpha_1, \beta_1) \cap (\alpha_2, \beta_2) = \emptyset$,
- $\beta_2 - \alpha_1 < 2\pi$,
- $r \geq \frac{1}{2}$,
- $N_s \cdot (1 - r)^{2l} \leq \varepsilon$.

Let us explain the four conditions of this definition. The first one says that the pieces of curves spans a disjoint set of tangent angle. The second one is only technical and asserts that the curve does not make a complete loop. The third one asks that the gap between each angle set is large, since it is at least half the total interval of the spanned angles. Finally, in the last condition, the left hand term is the product of total number of samples and a probability, which is the probability that none of $2l$ random values, independent and uniformly distributed

in (α_1, β_2) fall in the interval (β_1, α_2). Remark also that this definition is not purely local since the gap must separate a sufficient number of points before being meaningful. We can associate a size to a breaking, which is the number of samples we need before it becomes meaningful.

6.2 Terminators

The preceding definition allows (see the experiments section) to detect where a curve contains two smooth parts separated by an sudden change. This does not always correspond to corners or junctions. Indeed, it is not rare that a single level line contains two parallel good continuations connected by a rapid U-turn, for instance, the boundary of a digital segment. In the same idea, when we scan a pen stroke, it is very likely that very long and thin level sets are contained in the stroke. They do not necessarily follow the whole boundary of the stroke since a local extremum may cut the level set in several parts. For instance, Fig. 6 represents some level lines of a part a Kandinsky's painting (see on Fig. 15 the black circle in the middle below the dark dot). Because of the blurring and interpolation, there are many parallel level lines in the circle (not all represented). Local extrema cut level lines into several parts. On Fig. 6, at

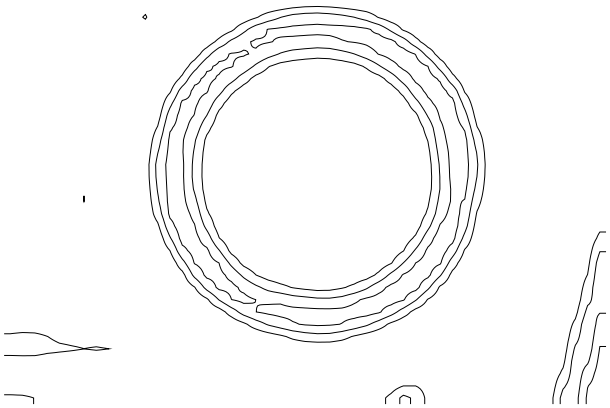

Figure 6: Level lines of a part of Kandinsky's painting (see Fig. 15). At least four terminators appear inside the circles but they are not conspicuous. If we look at this figure rapidly without any precise examination we only see concentric circles

least four terminators (U-turns between two good continuations) appear. If we isolated the corresponding level lines, we would interpret them as the end of a thin structure and this is a very strong geometrical feature. Instead of this, broken curves are connected afterward to form two other good continuations (two circles). Then, the created circles are parallel to the circles that already exist, and this new structure (six parallel circles) is very pregnant since it is closed, convex, symmetric, etc...
The endpoints have been called *terminators* by Bergen and Julesz [2], and are particularly important in the perception of textures. We shall use the same denomination although the

terminators we shall detect shall not correspond to low scale textures, since our definition requires that terminators connect pieces of curve that have to be long enough, and this requirement is generally too strong for small textures. Moreover, the example of Fig. 6 shows that many of our terminators shall also require a further masking analysis. In any case, it is important to classify continuation breakings a bit further by giving a definition of terminators, since they do not correspond to junctions or corners. We take the same notations as in Def. 5. Let $d_k = |C(i_0 + k) - C(i_0 - k)|$, where i_0 is the index of the breaking candidate.

Definition 6 *We say that a point is a ε-terminator, if it is a ε-breaking and if $\max_{k \leq l} d_k \leq c_0 \delta$, where c_0 is a constant only depending on the sample size.*

This definitions only means that a terminator joins two pieces of curves that stay close to each other. For technical reasons due to the sampling of the curve, we take in practice $c_0 = \sqrt{5}/2$, which corresponds to a stroke with a width equal to δ and up to an offset due to the parameterization. If we change a little bit this value, some corners may become terminators and vice-versa but it makes sense. Indeed, if a stroke is thick enough or if we zoom in it, then we can see it as a rectangle, and if it is thin enough (or unzoom it), it may appear as a line with no width. In the same way, a very sharp corner may be seen as a terminator.

In practice, there are many detected meaningful breakings in images which do not correspond to anything, simply because the level lines which do not correspond to edges are noisy. Because of numerical artifacts (quantization, blocking effect due to compression, aliasing and ringing along edges) they may also contain short straight parts. Thus we impose that a breaking is close to a good continuation or another breaking. This is made easy be using the fact that a breaking is associated with a piece of curve. (See the definition of breakings.) We also impose that a good continuation does not contain any breaking. This does not affect the contrast invariance of the algorithms. Concerning terminators, we expect to find a lot of them in the width of contours and we want to eliminate them. We propose to do this as follows. We first impose that a terminator is closed to a good continuation. We can also compute the direction of a terminator which is the common direction of the pieces of curves starting from it. Then, if when we follow this direction we still fall in a pixel containing a good continuation, this means that the terminator is in the middle of some good continuation and we do not keep it.

7 Experiments

Let us first apply the detection to Kanizsa's figures (Fig. 7). They have been scanned directly from the book and are not binary images. In particular, they are very noisy and some structures of the page verso partially appear by transparency. We choose the quantization step such that any pixel is crossed at least by one level line. As can be expected, the noise due to the paper texture and the scan has no geometrical structure and good continuations exactly coincide with edges. Corners are suitably detected (see Fig. 8), except one in the

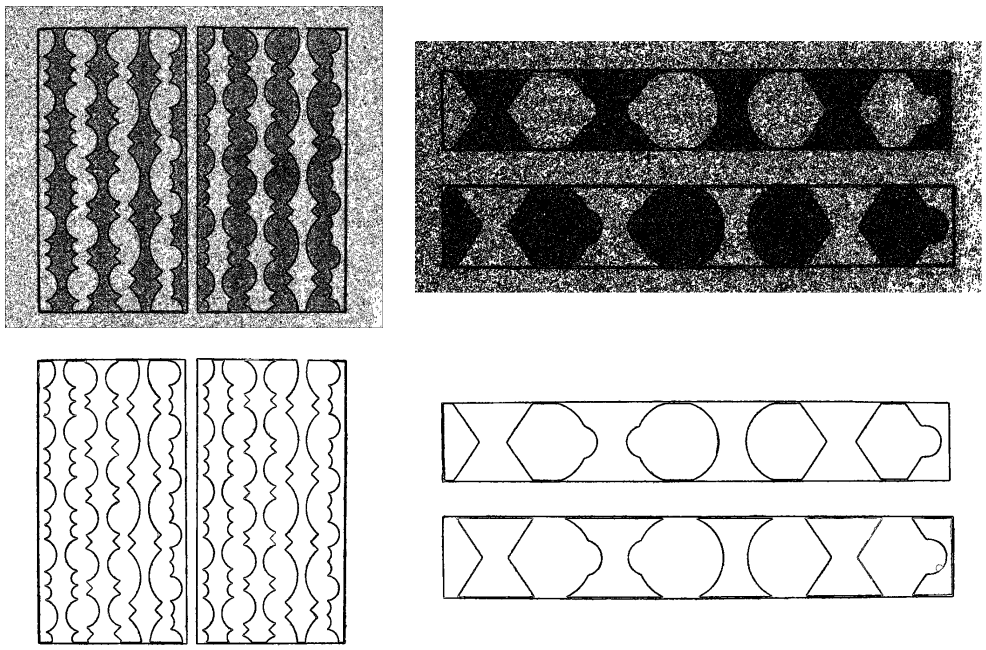

Figure 7: Detection of good continuation on Kanizsa's figure. Upper row: level lines with level multiple of 5. Lower row: good continuations with $\varepsilon = 1$. As expected, edges are piecewise smooth and contained in the union of the good continuations. The real thing is that no other level line is a good continuation

bottom-center. There is another detection on the top of the right part, but it is in fact a terminator detected as a corner because of the stroke width. If we display all the detected terminators without using the procedure described at the end of Sect. 6, we see that, as expected, most of them are located in the boundaries of the objects. After elimination, there are, in this case, no terminators left.

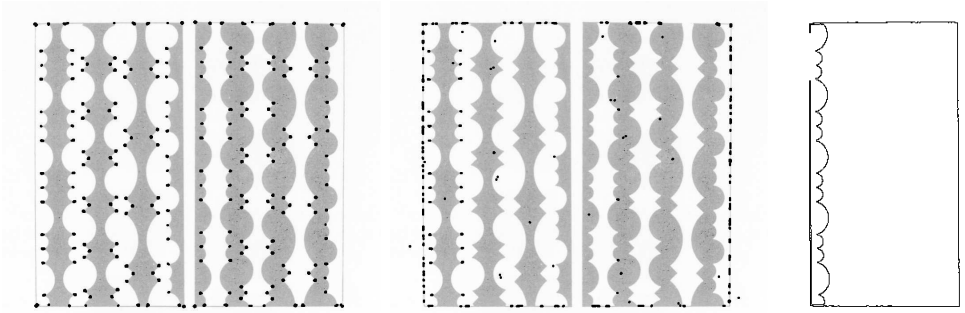

Figure 8: Detection of corners. Left: most meaninfgul corners are as expected. Middle: terminators before the clearing procedure. The terminators in the frame are aligned with the edge and point into good continuation. Terminators coinciding with corners correspond to inner level lines (since the angle is sharper) while corners correspond to outer level lines. Therefore, even if termintors point outward, since they are located in the inner boundary the point to the outer good continuations and are also eliminated. After the clearing procedure, there are no more terminators. Right: a level line of the left hand figure. There are two terminators which are both automatically eliminated

On Fig. 9, the coincidence of perceptual edges with good continuations is striking. It is also worth noticing that the whole window panes are detected as good continuations (and not only their sides), because their boundary is long enough and because of blur and interpolation, the corners are smoothed. Since level lines accumulate around edges, we can visually check that they are indeed smooth. But there are also many level lines in low contrasted regions (as the window panes), but only a very few of them satisfy the good continuation principle, and this is the real point.

On Fig. 10, we see that edges and good continuations often coincide agin, and moreover, we examine the influence of the sampling length δ, for which we choose three values $\delta = 1$, 1.5 and 2. With the minimum length, a few short insignificant good continuations are detected, while most of them disappear above the critical sampling. On the contrary, large samples may absorb some small oscillations (see for example the right most pillar). Nevertheless, the most meaningful structures cross the scales with no modification.

Good continuations do not always coincide with contrasted edges. For instance, on Fig. 12, the tree on the right has not a smooth boundary (it is nearly fractal) while it is contrasted. Clouds are neither contrasted nor smooth, but if we zoom out the image in agreement with Shannon's theory (i.e. we apply an appropriate blurring before downsampling), then clouds

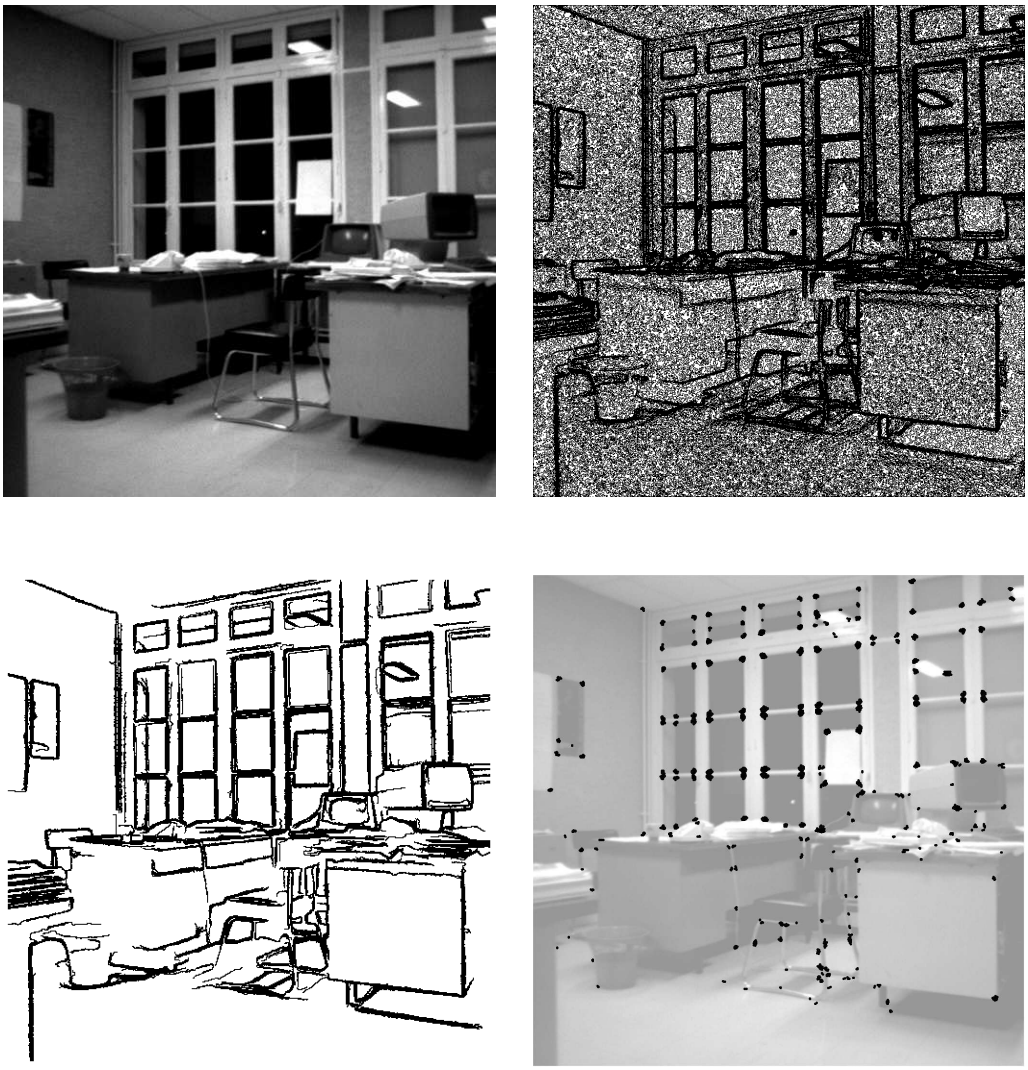

Figure 9: INRIA desk. From top to down and left to right: original image, level lines (multiples of 5), meaningful good continuations, meaningful corners and junctions

Figure 10: The Annunciation (Fra Angelico). Influence of sampling. From top to down and left to right. Original image, then $\delta = 2$, 3 and 4

become good continuations while the tree is still not meaningful. If we take $\varepsilon = 10^{-5}$ instead of $\varepsilon = 1$, we obtain only twice less good continuations, in agreement with the fact that the detection depends on $\log \varepsilon$. We also remark that we obtain nearly the same results by detecting good continuations among contrasted edges and good continuations alone.

Figure 11: Original image of the Valbonne church (courtesy of Robotvis group)

In Fig. 13, level lines correspond to edges only on short distances since edges are often mixed up with the surrounding textures. Nevertheless, good continuations still coincide with meaningful contrasted edges [9] in a very surprising way. This time we also find a large number of curves surrounding the cheetah and that are not contrasted edges. It is also worth noticing that they all have a common direction, and this is not a coincidence, since they are due to the grass orientation. The posterior association of these directions shall give some information on the texture. (See [7] for a grouping of parallel directions.) We also tested the influence of the maximal number of false detections by taking $\varepsilon = 1$ and $\varepsilon = 10^{-5}$. As expected, some good continuations disappear but the largest shapes are preserved.
On Fig. 14, the image is a texture made by straw wisps. In this precise example, there is a high density of terminators which has to be studied further.

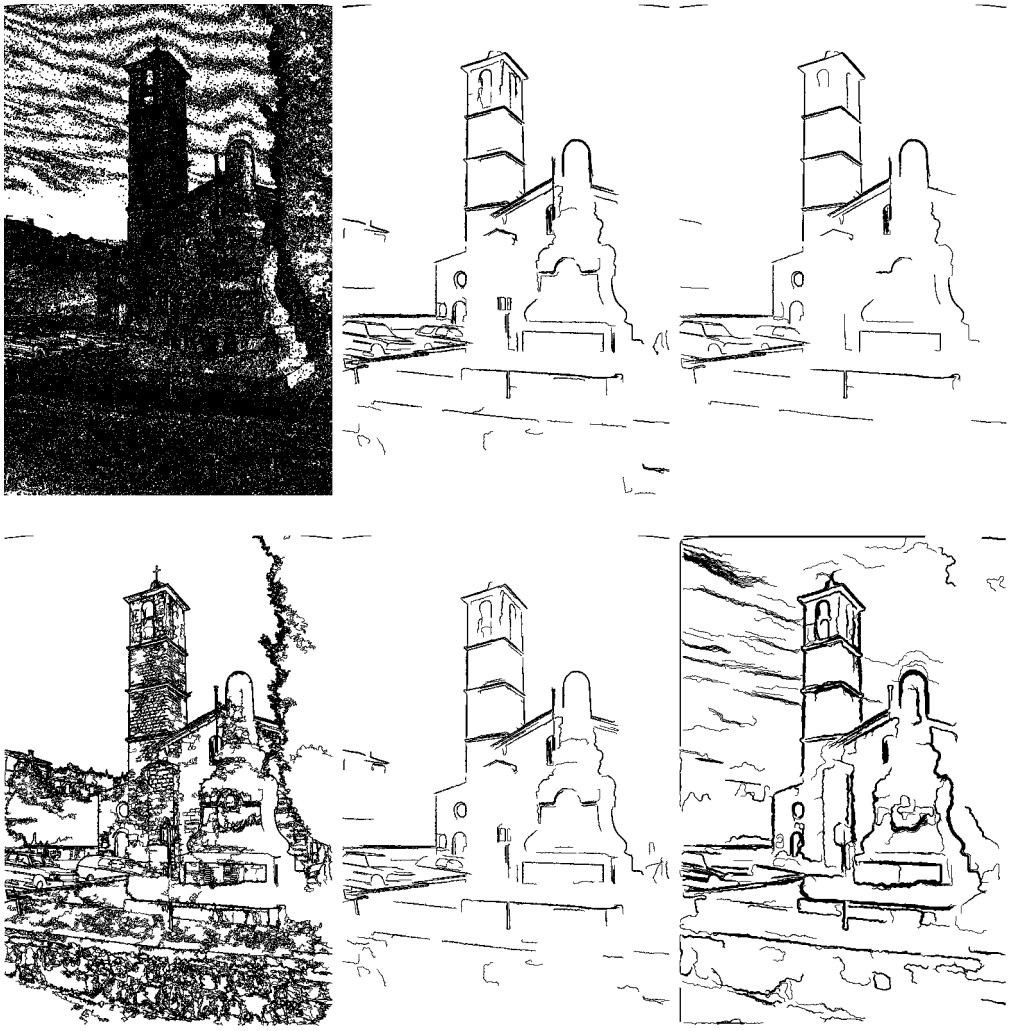

Figure 12: Photograph of Valbonne church . Top, from left to right: level lines (quantized each 10 levels). Good continuations with $\varepsilon = 1$ and $\varepsilon = 10^{-5}$. The main features are stable with respect to the false detection rate since there are 6540 good continuations for $\varepsilon = 0$ and 3326 for $\varepsilon = 10^{-5}$. Bottom, from left to right. Helmholtz edges [9]. Bottom, good continuations among meaningful edges, and good continuations in the zoomed out image (by a factor 2). In this last case, the clouds are detected as good continuations. These curves should then be grouped by parallelism

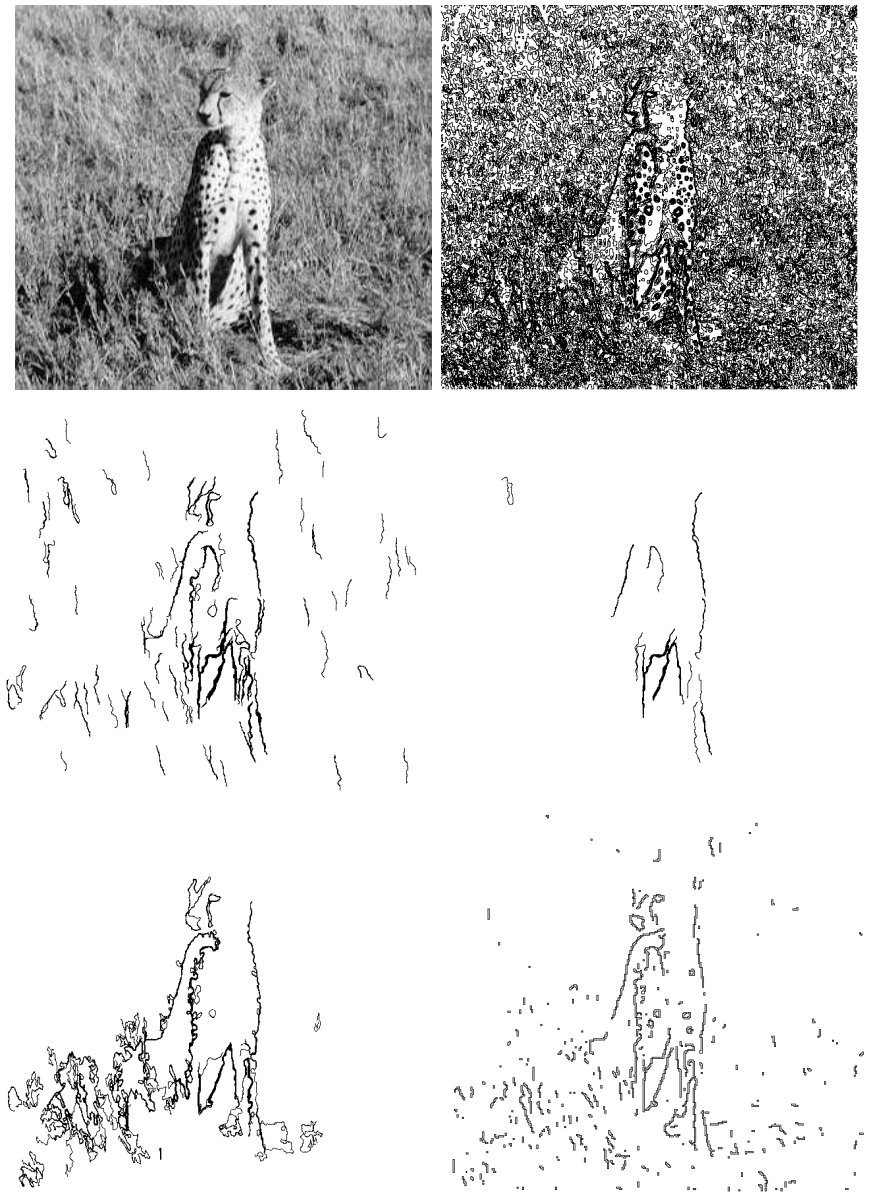

Figure 13: Detection of convex level lines on a textured image. Top: original image, and level lines (quantized each 30 levels). Middle row. Left: meaningful good continuations with $\varepsilon = 1$. The segments corresponding the grass should be grouped with respect to orientation. Right: meaningful good continuations with $\varepsilon = 10^{-5}$. As can be seen, the detection is stable with respect to ε. Bottom row. Left: "Helmholtz edges" (see [9]). Right: Canny's edge detector

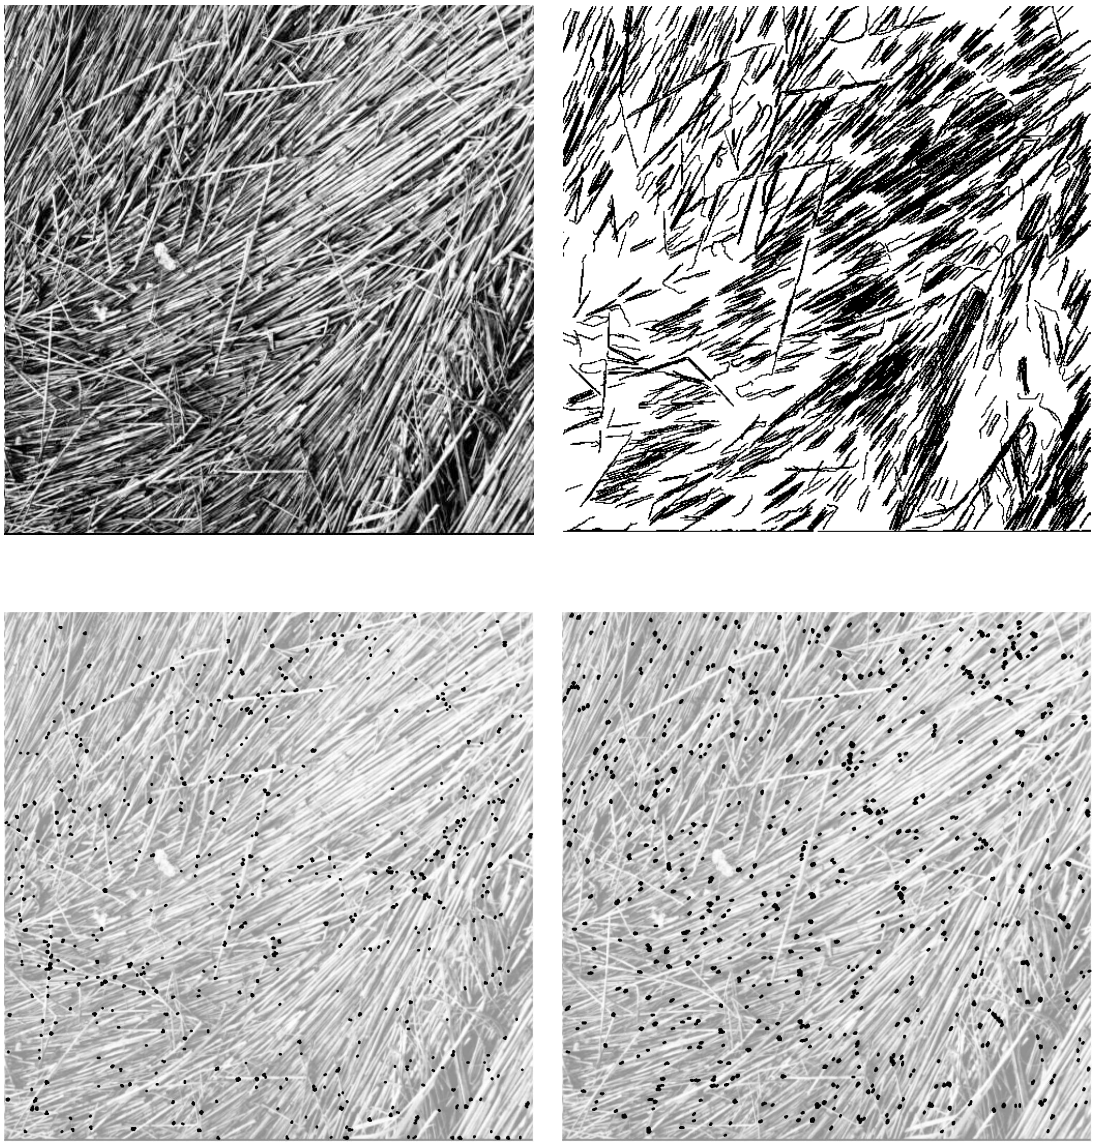

Figure 14: Detection of convex level lines on an image of straw (top left). Good continuations give the direction of the texture (top right). There are corners/junctions (bottom left), but terminators are also very numerous (bottom right). A further analysis should analyse whether the continuations have a common direction and the terminators density (textons density [2])

Finally, we display an experiment of good continuation breaking on a painting by Kandinsky (Fig. 15). We emphasize that the image is not smooth at all and that level lines are really noisy and fill the image. As above, we eliminate terminators which are surrounded by good continuations. Moreover, we eliminate corners and terminators which are isolated (there are no other ones in the same pixel). As edges contain many level lines, corners and terminators appear in clusters, and this should also give them some relevance. (See [10] for cluster detection.) This corner and junction detection is less local than most of existing methods (A recent paper by Sojka [34] gives a benchmark between different classical algorithms.) As an exemple, we compare our result with the classical Harris detector [16]. Let us point out that our method also gives a quantitative meaning of the points we find in terms of false detection rate, which could be even stabler if we found point clusters. For applications, for instance image registration, it is useful to find as many features as possible, but it is also important that these features are robust. Our method is certainly not able to recover all the corners, since curves have to be long enough, but this also corresponds to what a real corner should be. However, since the false detection rates are controlled both for good continuation detection and breakings, and since corners can further be grouped in cluster, our information is certainly more reliable, and this reliability is measured by a number of false detections. Moreover, we insist on the fact that there are no critical parameters that have to be tuned by the user. It is not the case for Harris' detector for which we have to specify the standard deviation of a smoothing kernel, the size of the neighborhood to compute the location of the maxima of the cornerness function and a threshold on these maxima. Another parameter appearing in the definition of cornerness has been suggested by Harris and Stephen [16]. We also remark that there is no distinction between corners and terminators, which are perceptually quite different.

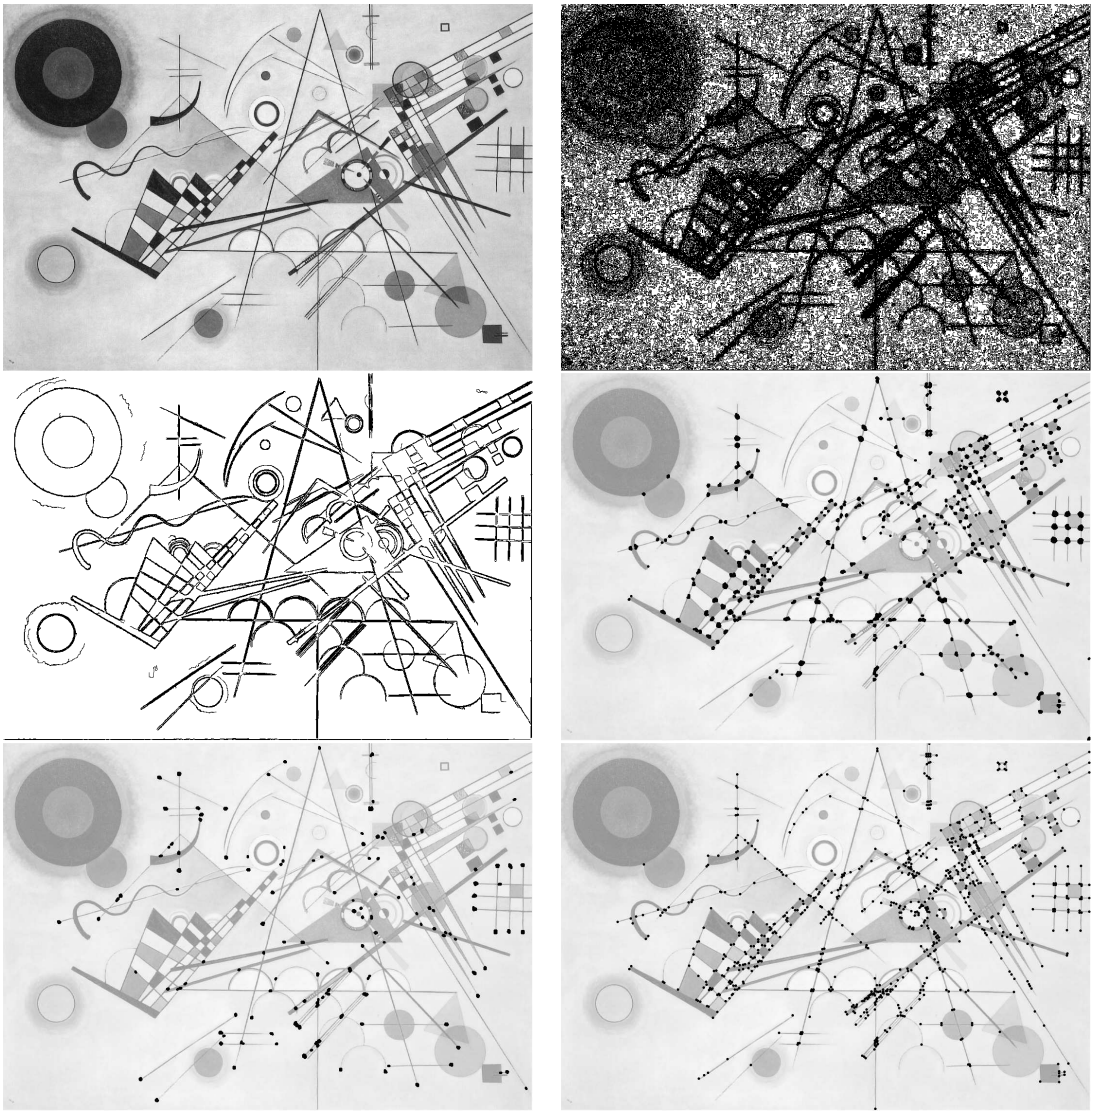

Figure 15: Good continuation and breaking detection. Top: Original image, level lines (multiples of 3). Middle: good continuations, and corner/junctions. Bottom row: meaningful terminators, and corners detected by Harris [16] detector. We have removed terminators as described in Sect. 6 and kept only non-isolated breakings. The only important parameter is the scale (or sampling length). In Harris detector, we have to fix three parameters: a smoothing scale, a neighborhood size for maxima computation and a threshold for the cornerness function. In addition, another parameter is used for the cornerness definition and its value is suggested by Harris

8 Conclusion, discussion and perspectives

In this paper, we proposed a method to define the approximate regularity of digital curves. This does not use a model of "good shape" but measures *a contrario* how a curve differs from a discrete Brownian motion. The detection thresholds are then fixed by Helmholtz principle. The algorithm has two parameters: a sampling distance and a maximal number of false detections. None of them has to be tuned precisely. In practice, the number of false detections may be taken equal to 1, and the sampling should be taken minimal while staying in agreement with Shannon's sampling theory. This algorithm is decisive since we do not need to fix thresholds a posteriori. A fact which is well known in signal processing is that the price to pay for controlling the false detection rate is to accept not to detect true features. In practice, we saw that most important smooth curves are detected whatever the value of the false detection rate, which was predicted by the log dependence of the method with respect to ε. Our algorithm does not use any a priori learning, except the number of curves in the image and the *a contrario* model we use. Moreover, it does not require any a priori smoothing, and we showed that under a particular shape scale space, the detection is not improved by smoothing, when the sample length is suitably chosen, such that it makes sense to apply Helmholtz principle to smooth curves.
We applied this algorithm to detect smooth pieces of level lines in images. It is then checked that contours are almost always improbably smooth curves. This justifies a posteriori many segmentations algorithms which take for granted that edges are smooth (as Mumford-Shah segmentation [29] or active contours [20]). We can also notice that in natural images, the number of good continuations is small compared to the total number of level lines in the image, and this is also in agreement with the fact that natural images are very irregular [12]. Our detection does not depend upon any contrast information and the coincidence of the detected curve with edges is surprising. However, we stress that our algorithm is not an edge detector and not meant to replace them. We a priori did not aim at finding edges, but only curves whose variations are too small to occur only by chance. The conclusion is that edges satisfy many partial properties and either contrast or the good continuation principle are almost sufficient to find them. Of course, regularity and contrast are not always equivalent: contrasted parts may be irregular and smooth parts may have a low contrast. This implies that *all* partial gestalts have to be examined. This, of course, increases the computational time, but a feature that is found by several independent detectors is much more reliable. Such an exhaustive detection is only a preliminary step since grouping and masking are crucial in making structures conspicuous. These interactions are nonlocal and represent both a theoretical and computational challenge.
Before this, we certainly should reduce the number of good continuations or the computational cost will be too important when we use combinations of partial gestalts. A first remark we already made is that edges usually contain several good continuations since level lines are very close to each other. A very simple possibility is to select, among the good continuations crossing a pixel, the one with the smallest number of false detections. This simple procedure allows to divide the number of detected curves by a factor up to 5. In order to minimize the

number of corners and terminators, we should group them in clusters.
Contrary to Guy and Medioni [15], we have not tried to detect lines from sets of dots, and only tested the good continuation principle on level lines of images. Now, we can apply it to any type of curves (not only level lines). Once we have detected some local structures, we may group them with respect to proximity and resemblance. We can then evaluate the smoothness of any curve passing through a set of these features. Indeed, the detection threshold is function only of the total number of possible curves and a scale parameter giving the typical distance between objects. We may also introduce an a priori knowledge by computing the empirical distribution of the angle variation and take it as a prior for each set of curves.
In a short term view, good continuations provide an efficient shape detector for image and shape matching [30]. We also plan to use curve smoothness for motion analysis. Indeed, trajectories of objects correspond to highly coherent spatiotemporal structures. The coherence of motion also certainly interacts with the facility to group dots in a noisy environment. Contrary to the approach we developed in this paper, such a method would first require to extract some information in a more subtle way than simple level lines.

Acknowledgments. I would like to thank Jean-Michel Morel, Agnès Desolneux and Lionel Moisan for introducing the subject of Gestalt detection to me.

References

[1] L. Alvarez, F. Guichard, P.L. Lions, and J.M. Morel. Axioms and fundamental equations of image processing. *Arch. Ration. Mech. Anal.*, 123(3):199–257, 1993.

[2] Bergen and Julesz. Textons, the fundamental elements in preattentive vision and perception of textures. *Bell-Syst. Techn. J.*, 62(6):1619–1645, 1983.

[3] J. Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(6):679–698, 1986.

[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, R. Kimmel, and G. Sapiro. Geodesic active contours. *International Journal of Computer Vision*, 22(1):61–79, 1997.

[6] R. Deriche and G. Giraudon. A computational approach for corner and vertex detection. *Int. Jour. of Comp. Vis.*, 10(2):101–124, 1993.

[7] A. Desolneux. *Événements significatifs et applications à l'analyse d'image.* PhD thesis, ENS-Cachan, 2000.

[8] A. Desolneux, L. Moisan, and J.M. Morel. Meaningful alignments. *International Journal of Computer Vision*, 40(1):7–23, 2000.

[9] A. Desolneux, L. Moisan, and J.M. Morel. Edge detection by Helmholtz principle. *Journal of Mathematical Imaging and Vision*, 14(3):271–284, 2001.

[10] A. Desolneux, L. Moisan, and J.M. Morel. Partial gestalts. http://www.cmla.ens-cachan.fr. Submitted to IEEE. Trans. on PAMI, 2001.

[11] M. Gage and R.S. Hamilton. The heat equation shrinking convex plane curves. *J. Diff. Geom.*, 23:69–96, 1986.

[12] Y. Gousseau and J.M. Morel. Are natural images of bounded variation? *SIAM J. of Math. Anal.*, 33(3):634–648, 2001.

[13] M.A. Grayson. The heat equation shrinks embedded plane curves to round points. *J. Diff. Geom.*, 26:285–314, 1987.

[14] F. Guichard and J.M. Morel. *Iterative Image Filtering.* To appear, 2002.

[15] G. Guy and G. Medioni. Inferring global perceptual contours from local features. *int. Journal of Comp. Vision*, 20(1):113–133, 1996.

[16] C.G. Harris and M. Stephens. A combined corner and edge detector. In *4th Alvey Vision Conference, Manchester*, pages 189–192, 1988.

[17] D.W. Jacobs. Robust and efficient detection of salient convex groups. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(1):23–37, 1996.

[18] G. Kanizsa. *Vedere e Pensare.* Il Mulino, Bologna, 1991.

[19] G. Kanizsa. *La Grammaire du Voir.* Diderot, 1996. Original title: *Grammatica del vedere.* French translation from Italian.

[20] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1:321–331, 1987.

[21] S. Ladjal. Personal communication, 2001.

[22] T. Lindeberg and M.X. Li. Segmentation and classification of edges using minimum description length approximation and complementary junction cues. *Computer Vision and Image Understanding*, 67(1):88–98, 1997.

[23] J.L. Lisani, P. Monasse, and L. Rudin. Fast shape extraction and application. Preprint 16, CMLA, ENS-Cachan. Available at http://www.cmla.ens-cachan.fr, 2001.

[24] D. Marr. *Vision.* N.York, W.H. and Co, 1982.

[25] D. Marr and E. Hildreth. Theory of edge detection. *Proceeding of Royal Society of London*, 207:187–207, 1980.

[26] F. Mokhtarian and A. K. Mackworth. A theory of multi-scale, curvature-based shape representation for planar curves. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 14(8):789–805, 1992.

[27] P. Monasse and F. Guichard. Fast computation of a contrast invariant representation. *IEEE Transactions of Image Processing*, 9(5):860–872, 2000.

[28] U. Montanari. On the optimal detection of curves in noisy pictures. *Communications of the ACM*, 14(5):335–345, 1971.

[29] D. Mumford and J. Shah. Optimal approximation by piecewise smooth functions and associated variational problems. *Communication on Pure and Applied Mathematics*, XLI-I(4), 1989.

[30] P. Musé, F. Sur, and J.M. Morel. Recherche dans les grandes bases de formes. Technical Report 2002-02, ENS Cachan, 2002. In French.

[31] P.J. Olver, G. Sapiro, and A. Tannenbaum. Differential invariant signatures and flows in computer vision: a symmetry group approach. In B. M. Ter Haar Romeny, editor, *Geometry-Driven Diffusion in Computer Vision*, pages 255–306. Kluwer Acad. Publ., 1994.

[32] H. Pao, D. Geiger, and N. Rubin. Measuring convexity for figure/ground separation. In *International Conference of Computer Vision, ICCV 99.*, volume 2, pages 948–955, 1999.

[33] C.E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423, 623–656, 1948.

[34] E. Sojka. A new algorithm for detecting corners in digital images. In *18th Spring Conference on Computer Graphics*, 2001.

[35] M. Wertheimer. Untersuchungen zur Lehre der Gestalt, II. *Psychologische Forschung*, 4:301–350, 1923.

[36] S.C. Zhu. Embedding Gestalt laws in markow random fields. *IEEE Trans. on PAMI*, 21(11):1170–1187, 1999.

Unité de recherche INRIA Rennes
IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)
Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38330 Montbonnot-St-Martin (France)
Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)
Unité de recherche INRIA Sophia Antipolis : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

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