Good continuations in digital image level lines

Frédéric Cao
IRISA/INRIA
Campus universitaire de Beaulieu
35000 Rennes, France
Frederic.Cao@irisa.fr

Abstract

We propose a probabilistic algorithm able to detect the curves that are unexpectedy smooth in a set of digital curves. The only parameter is a false alarme rate, influencing the detection only by its logarithm. We experiment the good continuation criterion on image level lines. One of the conclusion is that, accordingly to Gestalt Theory, one can detect egdes in a way that is widely independent of contrast. We also use the same kind of method to detect corners and junctions.

1. Introduction

Object boundaries are very often smooth curves. This is not merely coincidental since, following Gestalt Theory [24], an entity is seen as an object precisely because it has a smooth boundary. Thus, detecting smooth curves in a very robust way is an important topic in computer vision. Of course, regularity is not the only property that makes objects conspicuous. Gestaltists have tried to give the (short) list of properties, expressed since the beginning (in the 1920's) in geometrical terms, describing the vision process in a phenomenological point of view. Without being exhaustive, we can cite convexity, good continuation, symmetry, closedness, vicinity, similarity. These properties have been called *partial gestalts* by Desolneux, Moisan and Morel [9], who initiated a program to transpose Gestalt Theory to Computer Vision. The first step of this program is to detect robustly all partial gestalts. The second step, also predicted by Gestalt Theory, is to recursively group the detected features into more rigid and stable structures, by using the same elementary Gestalt laws. We refer the reader to [15] for more details. In this paper, we are only interested in the detection of a particular partial gestalt, namely good continuation. Good continuations mean that we are inclined to grouping patterns when they can be embedded in a smooth curve. It applies when there are gaps in the patterns (dotted line for instance), but also for crossing curves. In this case, perception usually prefers to choose to group the parts of the curves such that most of them are smooth. In this paper, we do not try to detected dotted lines, but address the following sub-problem: given a set of digital curves, we want to detect all their regular part, *and only them*. We introduce a definition of statistical regularity using the Helmholtz Principle described by Desolneux *et al.* [9]. This principle does not use an a priori or learned model, but a false *a contrario model*, based on a local independence assumption. The detected features are large deviations from this a contrario model: a curve is smooth when the probability that it is the conjonction of independent local regularity observations is very small. This shall be detailed below. Perceptual grouping based on good continuation was proposed with a variational approach by Guy and Medioni [11]. See also [14, 21, 23] for related works in the detection of lines or convex polygons in images. However, variational methods are not decisive algorithms and usually require a threshold to decide whether the optimal solution of the problem (segmentation, grouping, etc...) is correct. This threshold is usually data-dependent and its choice may be tedious. Our method is not variational and has a unique decision threshold which is an expected number of false detections. We shall see that the detection depends only on the logarithm of this threshold. Thus, in digital image analysis, our algorithm can be considered as parameter free.

We apply it to image level lines and we observe that we detect most edges. Since level lines give a constrast invariant representation of the image, this seems contradictory with the edge detection doctrine [3, 8, 19], which is based on contrast. Actually, the conclusion is that we should distinguish at least two types of boundaries: contrasted boundaries and good continuations. Most of the time, they coincide, but by difference, we can classify those which are due to contrast or regularity only.

The plan is as follows: in Sect. 2 we explain the a contrario model. In Sect. 3, we introduce our good continuation principle. We also introduce a method to detect good continu-

Proceedings of the Ninth IEEE International Conference on Computer Vision (ICCV 2003) 2-Volume Set
0-7695-1950-4/03 $17.00 © 2003 IEEE

ation breakings which are candidates for corners and junctions. (Sect. 4.) We display some experiments in Sect. 5 and conclude in Sect. 6. The proofs of all the announced results can be found in [4].

2. The a contrario model

The detection principle we use is called Helmholtz principle and was introduced by Desolneux, Moisan and Morel [9]. We can formulate it as follows. Assume that we observe a *finite* collection $\mathcal{C} = (c_1, \ldots, c_N)$ of N local objects sharing a given quality Q (for instance color, or orientation, or low curvature values...). Assume also that we observe N_T possible groups of these objects (that is to say a given subset G of all the parts of $\mathcal{C}$). Which among those N_T groups are perceptually meaningful? Following Helmholtz [13], such groups are conspicuous because they are not *generic* (or random) configurations. In order to quantify this deviation to randomness, let us define the number of false alarms (NFA)

$$NFA = \mathbb{E}(\#\{g \in G,\ \text{sharing } Q\}),$$

to be the expectation of the number of group candidates $g = (x_1, \ldots, x_k) \in G$ sharing Q when we *mentally* assume, that, anything else held equal, the quality Q is independently and uniformly distributed on $(x_1, \ldots, x_k)$. This independence assumption is an a contrario assumption that has to be rejected for meaningful groups. Let us now assume that $NFA < \varepsilon$. If $g = (x_1, \ldots, x_k)$ share Q then we say that g is ε-meaningful. It means that, in average, we can observe less than ε such groups by chance. If ε is very small (we will see how to choose it), this allows to reject the independence assumption, and validate the detection.

3. Good continuation principle

3.1. Meaningful good continuation

Let C be a rectifiable plane curve. We sample C with a sample length equal to δ and we assume that C has $n+1$ sampled points p_0, ..., p_n. 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 denote by $k_i = \theta_i - \theta_{i-1}$, the difference between two consecutive tangent angles, called (by extension) curvature in the following. Let us define our a contrario model. Assume that C is an isotropic stationary random walk with independent increments. Then, the k_i are independent and identically distributed random variables, uniform in $(-\pi, \pi)$. For $\kappa \in (0, \pi)$, let us compute the probability $P(\kappa, n)$ that for all $1 \leq i \leq n$, $|k_i| \leq \kappa$. From the stationarity and the independence assumptions, we simply have

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

If $P(\kappa, n)$ is very small, then there is little chance that C is a random walk. However, we have to make clear how small $P(\kappa, n)$ has to be. It is natural to assume that a good continuation does not contain any right angles, and we require that $\kappa \leq \kappa_{\text{max}} = \frac{\pi}{2}$. Let N_κ be a positive integer. We set for all $1 \leq i \leq N_\kappa$, $\kappa_i = \frac{i}{N_\kappa}\kappa_{\text{max}}$.
Let us now consider a finite family of sampled curves $(C_j)_{j \in J}$, and let L_j the number of curvature samples of C_j. The total number of subcurves of all the C_j is less than $N_c = \sum_{j \in J} L_j^2$ (with equality if the curves are closed).

Definition 1 *Let Γ be a connected subcurve of one of the C_k, with n curvature samples $(k_i)_{1 \leq i \leq n}$. Let $k = \max_{1 \leq i \leq n} |k_i|$ be the maximal curvature of Γ. Let also $\kappa = \min_{1 \leq i \leq N_\kappa} \{kappa_i, k < \kappa_i\}$. We call number of false alarms of Γ*

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

if $k < \kappa_{max}$ and $NF(\Gamma) = N_c N_\kappa$ else.

In order to make notations shorter, we shall denote "Γ is a ε-g.c" for Γ is a ε-meaningful good continuation.

In the definition above, we simply multiply the probability (in the a contrario model) to observe a curve with n curvature samples smaller than one of the fixed values κ_i, by the total number of tested configurations $N_c N_\kappa$ (the number of subcurves times the number of tested curvatures). The meaning of this definition and the link with the general principle of Sect. 2 is given by the following result.

Proposition 1 *Assume that the curves C_k are stationary random walks with independent increments and known length L_k. Then the expected number of ε-g.c. is smaller than ε.*

The proof only relies on the additivity of expectation.

3.2. On the choice of ε and δ

As a direct consequence of the definition, a meaningful continuation containing n angle samples must satisfy

$$\exists i \leq N_\kappa \text{ such that } n \geq \frac{\log\left(\frac{\varepsilon}{N_c N_\kappa}\right)}{\log\left(\frac{\kappa_i}{\pi}\right)}. \tag{2}$$

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, and idem for N_κ. Even though level lines in images are not random walks (since they are at least self avoiding), we experimentally checked that the number of detections in an image of white noise is linearly increasing with ε and that, in practice, there are no detections for $\varepsilon = 1$.
The sampling length δ must be chosen conformly to Shannon's sampling theory. If the curves are extracted from digital images, there is a minimal scale under which the curves are smooth, but this obviously must not be considered as meaningful. Thus, we do not consider curves under Nyquist's distance, namely two pixels. It is possible to make the algorithm scale invariant by using several values for δ (larger than Nyquist's distance). The procedure stays unchanged except that the number of false alarms is multiplied by the number of tested δ.

3.3. Maximality

A very meaningful curve may contain a lot of smaller meaningful ones. In general, we only see the largest one (masking phenomenon). Thus, it seems reasonable to keep only this curve. The following definition (first introduced in [7]) aims at this.

Definition 2 *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,* $\Gamma \subsetneq \Gamma' \Rightarrow NF(\Gamma) < NF(\Gamma')$.

Maximal good continuations form individual objects as asserted by the next result.

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

4. Corners, junctions and terminators

Corners and junctions are usually geometrical strong cues of shapes, and give information on the relative depth of objects. Detecting them has been a subject of constant interest for twenty years (see [1, 6, 12] among many others). Most used algorithms are local and have several parameters. They usually define a corner as a point where the direction of the gradient has rapid variations. This requires a smoothing and one or several thresholds.
As in [17], we consider that corners are rapid changes of the direction of a curve, but that the curve also has to be flat enough on both sides of the corner. We also use an a contrario model. Therefore, the detection threshold, formulated in terms of false alarms will be given a fixed value, independent of the image.

4.1. Good continuation breaking

Let us consider an ideal corner made by two non collinear segments with a common endpoint. Let us denote by C the resulting (sampled) curve and i_0 the (sampled) abscissa of the corner. For an ideal corner, the histogram of the tangent angle θ_i at the point $C(i_0 + i)$ has exactly two distinct values α_1 and α_2. For a more realistic curve, we expect to observe a peak around each value and a gap between them.
We take the same angle quantization as in Sect. 3. The interval $(-\pi, \pi)$ is cut into $N_d = 4N_\kappa$ bins of the same size $\delta\theta = \frac{\pi}{2N_\kappa}$. For $l > 0$, we set

$$\alpha_1(l) = \min_{-l \leq i < 0} \left[\frac{\theta_i}{\delta\theta}\right], \quad \beta_1(l) = \max_{-l \leq i < 0} \left[\frac{\theta_i}{\delta\theta}\right]$$

and

$$\alpha_2(l) = \min_{0 < i < l} \left[\frac{\theta_i}{\delta\theta}\right], \quad \beta_2(l) = \max_{0 < i < l} \left[\frac{\theta_i}{\delta\theta}\right],$$

where the angles θ_i are measured relatively to the tangent at the corner point $C(i_0)$ and the brackets stand for the integral part. Of course, the number of points in the histogram is equal to the length of the piece of curves we consider, namely $2l$. A gap in the histogram is meaningful if its relative size is large enough. Thus, we shall impose that the relative size of the histogram in $(0, 2\pi)$ is large and that the relative size of the intervals (α_1, β_1) and (α_2, β_2) in (α_1, β_2) is small, for some l.
To quantify these properties, we adopt the a contrario model that the values of the angles are independent and uniformly distributed in $(-\pi, \pi)$. As in the case of good continuation, we shall define a finite number of events. The number of false alarms is defined as the product probability of the event and the total number of considered events. An ε-meaningful event is such that its number of false alarms is less than ε. Let us fix $l > 0$. We compute the probability that

- the spanned interval of angles (normalized by $\delta\theta$) is larger than s and smaller than N_d (we do not want the curve to make a whole loop).
- the spanned intervals on each side of the corner candidate are respectively smaller than m_1 and m_2,

Let us fix s, m_1, m_2 and the lower bound of the spanned interval. The probability that the width of the interval is more than s is

$$1 - \left(\frac{s}{N_d}\right)^{2l-1}.$$

Knowing this, points are uniformly distributed in this interval. The probability that the tangent on each side of the corner spans angles smaller than m_1 and m_2 is $\left(\frac{m_1 m_2}{s^2}\right)^l$. Thus, if we fix the lower bound of the interval, the probability of the above event is

$$p_{\text{angle}}(s, m_1, m_2, l) = \left(\frac{m_1 m_2}{s^2}\right)^l \left(1 - \left(\frac{s}{N_d}\right)^{2l-1}\right). \tag{3}$$

In order to define a number of false alarms, we now have to count the number of possible configurations. Elementary calculations show that the number of possible choices for (m_1, m_2, s) is

$$\sum_{s=2}^{N_d} \sum_{j=2}^{s} (j-1) = \frac{1}{6}(N_d - 1)N_d(N_d + 1).$$

There are also N_d choices for the lower bound of the interval spanned by the angles. Each point shall be tested; we denote by N_s the total number of samples. Finally, we have to fix the maximal number of values for l (the length of the angle). Let us denote this maximal length by L. The total number of tested configurations is equal to

$$N_T = \frac{1}{6} L N_s (N_d - 1) N_d^2 (N_d + 1).$$

Let us now fix the value of L. We choose it as the minimal integer such that

$$N_T \cdot p_{\text{angle}}(3\kappa_{\max}, \kappa_{\max}, \kappa_{\max}, L) < \varepsilon. \tag{4}$$

Remark that this definition makes sense, since this quantity tends to 0 when L tends to $+\infty$. The interpretation is the following. If $l > L$, a histogram with length $3\kappa_{\max}$ with two modes and a gap of relative size $\frac{1}{3}$ (i.e. corresponding each to an angle $\kappa_{\max}$) forms a meaningful angle. (See definition below.)

Remark 1 *For $\delta = 2$, $N_\kappa = 10$, $N_s = 50.000$ (which is typically what is observed) and $\varepsilon = 1$, we get $N_d = 48$, and $L = 13$.*

We now have all the ingredients at hand to define an ε-meaningful good continuation breaking.

Definition 3 *Consider a set of curves, with a total number of samples equal to N_s. Let C a curve in this set and $C(i_0)$ a point of this curve. We consider the piece of curve $\Gamma_{i_0,l}$, whose points are the $C(i_0 + i)$, for $-l \le i \le l$. We say that C has a ε-meaningful angle centered at i_0 if there exists l, $1 \le l \le L$, and nonnegative integers s, m_1 and m_2 such that $s \le N_d$ and $m_1 + m_2 \le s$ satisfying the following requirements:*

1. $(\alpha_1, \beta_1) \le \cap(\alpha_2, \beta_2) = \emptyset$,
2. $s \le \beta_2 - \alpha_1 < N_d$,
3. $\beta_1 - \alpha_1 \le m_1$, $\beta_2 - \alpha_2 \le m_2$,
4. $\alpha_2 - \beta_1 + 1 \ge \max(\beta_1 - \alpha_1 + 1, \beta_2 - \alpha_2 + 1).$,
5. *the number of false alarms*

$$nf(i_0, l) \equiv N_T \cdot p_{\text{angle}}(s, m_1, m_2, l) \tag{5}$$

is less than ε.

If C has a meaningful angle at i_0, we set

$$l_0 = \min\{l, \, nf(i_0, l) < \varepsilon\}.$$

We call angle at i_0 the piece of curve containing the points with index between $i_0 - l_0$ and $i_0 + l_0$. We call number of false detections of the angle at i_0 the number $NF_a(i_0) = nf(i_0, l_0)$.

As for good continuation, this choice is motivated by the following result.

Proposition 3 *If $(C_j)_{j \in J}$ is a set of random walk with independent and isotropic increments, the expected number of ε-meaningful angles is less than ε.*

4.2. Corners

A meaningful corner is simply a good continuation breaking connecting two good continuations.

Definition 4 *We say that i_0 is a ε-meaningful corner, if*

- *there is a ε-maximal meaningful angle centered at i_0,*
- *and the endpoints the angle belong to a good continuation.*

Since ε-corners are ε-angles, Proposition 3 obviously holds for ε-meaningful corners.

4.3. Terminators

By the definition above, corners form a particular class of breakings. Among corners, we can also distinguish a subclass that we call *terminators*. Roughly speaking, they are the points where curves make a U-turn. They are not considered as corner but as the endpoint of thin objects. The denomination is due to Bergen and Julesz [2] who studied the terminators in texture perception.
Assume that C is a random walk with isotropic and independant increments. Then,

$$\mathbb{E}(|C(i) - C(-i)|^2) = 2i\delta^2,$$

where $\mathbb{E}$ is the expectation relatively to the random walk model. We take the same notations as in Def. 3 and Def. 4. Let i_0 be a meaningful corner with length l. Let

$$d_{\max}(l) = \max_{1 \leq k \leq l} |C(i_0 + k) - C(i_0 - k)|$$

and

$$d_{\min}(l) = \min_{1 \leq k \leq l} |C(i_0 + k) - C(i_0 - k)|.$$

Definition 5 *We say that i_0 is a ε-terminator, if it is a ε-corner and if*

$$d_{max}(l) - d_{min}(l) \leq \delta \sqrt{2l},$$

where l is the length of the breaking at i_0.

If we change a little bit this value, some corners may become terminators and vice-versa but it still 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 needle with negligible width. In the same way, a very sharp corner may be seen as a terminator and the algorithm will reproduce this indetermination. A definitive decision can be made only by considering more global and high level cues.

5 Experiments

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 [5]. 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. The topographic map gives a global contrast invariant representation of the image [10] that can be embedded in a tree structure. It can be efficiently computed by an algorithm called Fast Level Set Transform [18, 20]. The images are quantized such that the level lines densely covers the image.

On Fig. 1, we represent the level lines of an image (multiple of 10) and the good continuations (computed on all level lines). The influence of the false alarms rate is weak: for $\varepsilon = 1$ we get about 7,800 detections while we get 4,100 of them for $\varepsilon = 10^{-5}$. The top right subfigure is the results of meaningful edges of Desolneux, Moisan and Morel [8], which are contrast based. As expected, the tree on the right side of the image is detected as an edge, and is logically not detected as a good continuation since it has a very irregular boundary.

In the next experiment, we detect good continuations, corners and terminators on a painting by Kandinsky (Fig. 2). We emphasize that the image is not smooth at all, that level lines are really noisy and fill the image. As edges contain many level lines, corners and terminators appear in clusters, and this should also give them some relevance. This corner and junction detection is less local than most of existing methods. As an exemple, we compare our result with the classical Harris detector [12]. 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. Moreover, we insist on the fact that there are no critical parameters that have to be tuned by the user. For instance Harris' detector has four parameters: the standard deviation of a smoothing kernel, the size of the neighborhood to compute the location of the maxima of the cornerness function, a threshold on these maxima; the last parameter appears in the definition of cornerness, and its value has been suggested by Harris and Stephen [12]. We also remark that there is no distinction between corners and terminators, which are perceptually quite different. CPU time on this 1090×755 image is 61s.

In both experiments, we can check that most of edges are good continuations. But it is even more interesting to notice that, in spite of the large number of curves, there are only a very small number of good continuations which are not edges.

6. 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 random walk. The algorithm has one parameter: the maximal number of false detections. It does not need a precise tuning and in practice, it may be taken equal to 1. When dealing with level lines, the sampling should be taken minimal while staying in agreement with Shannon's sampling theory. In fact, several sample lengths can be tested; the same method still applies if we simply multiply the number of false detections by the number of test lengths. The algorithm is decisive since we do not need to fix thresholds a posteriori. It does not use any a priori learning, except the number of curves in the image and the qualitative *a contrario* model we use. Moreover, it does not require any smoothing. It can be shown that detection is not improved by smoothing. This shall be exposed in further works.

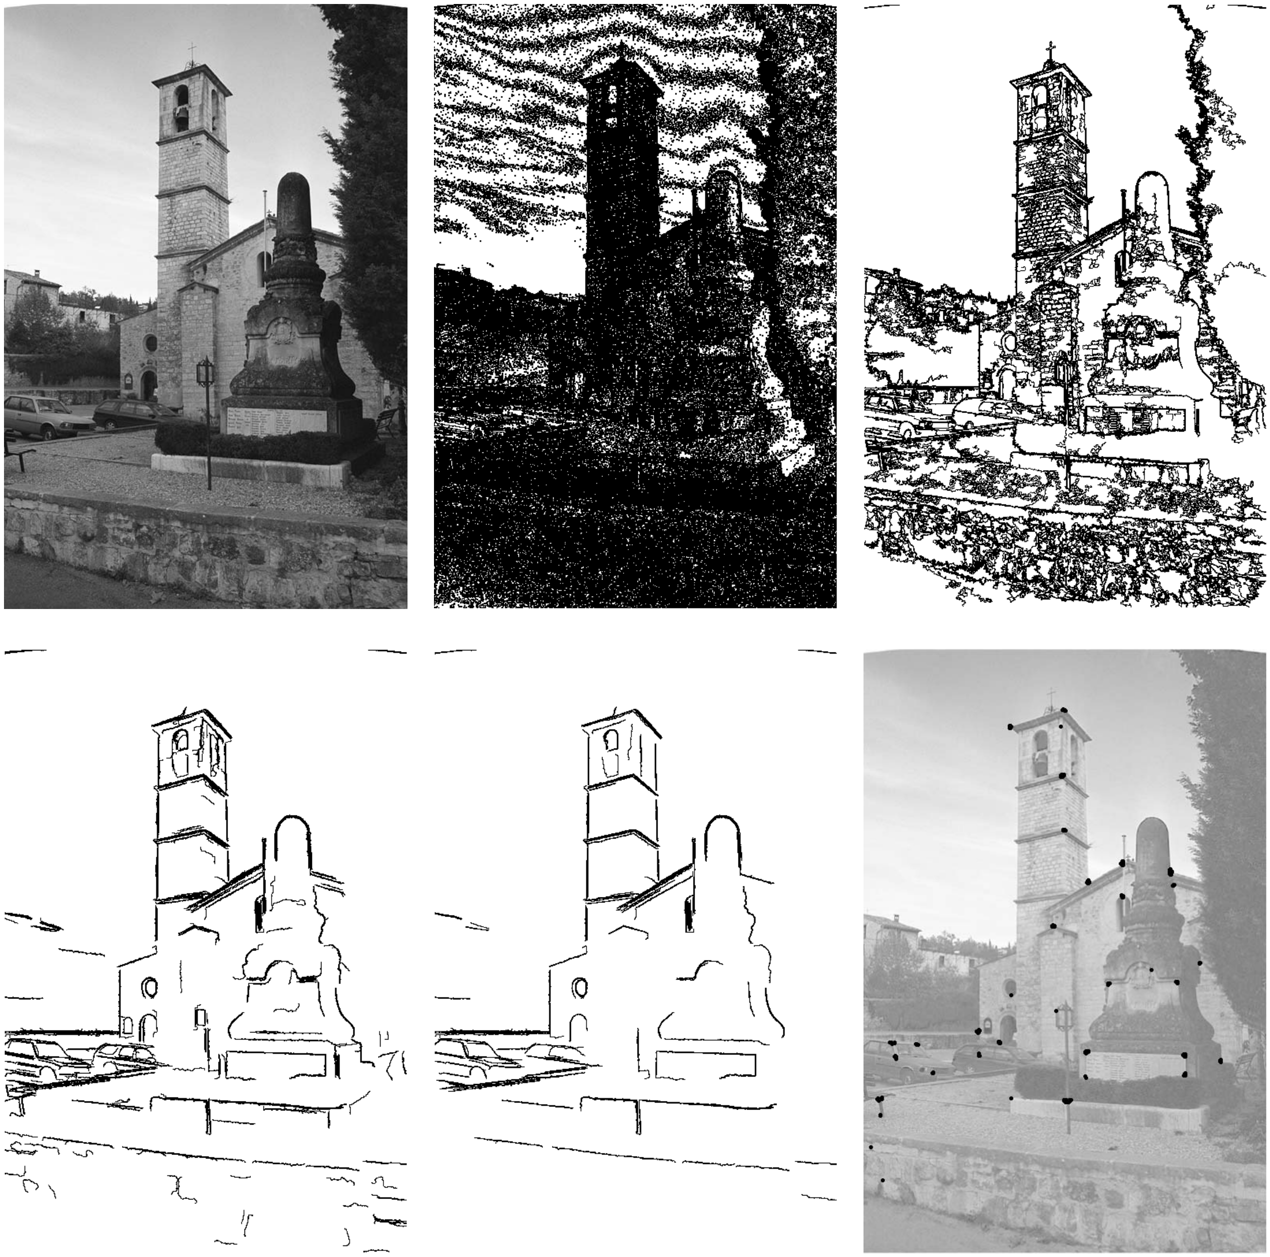

Figure 1. Photograph of Valbonne church . Top: on the left, the original image (512×768). On the right, its level lines (quantized each 10 levels), and Helmholtz edges, based on contrast [8]. Bottom, from left to right. Good continuations with $\varepsilon = 1$ and $\varepsilon = 10^{-5}$, and 1-meaningful corners. The main features are stable with respect to the false detection rate since there are 7771 good continuations for $\varepsilon = 1$ and 4101 for $\varepsilon = 10^{-5}$. The most meaningful curves in a computational point of view, are also the most visuallly meaningful. Textures can be contrasted but are not regular in general. This is the case of the tree contour on the right side of the image. There are initially, $523, 303$ complete level lines and $3 \cdot 10^7$ possible subcurves. CPU time is 20s on a Pentium IV 2.4GHz.

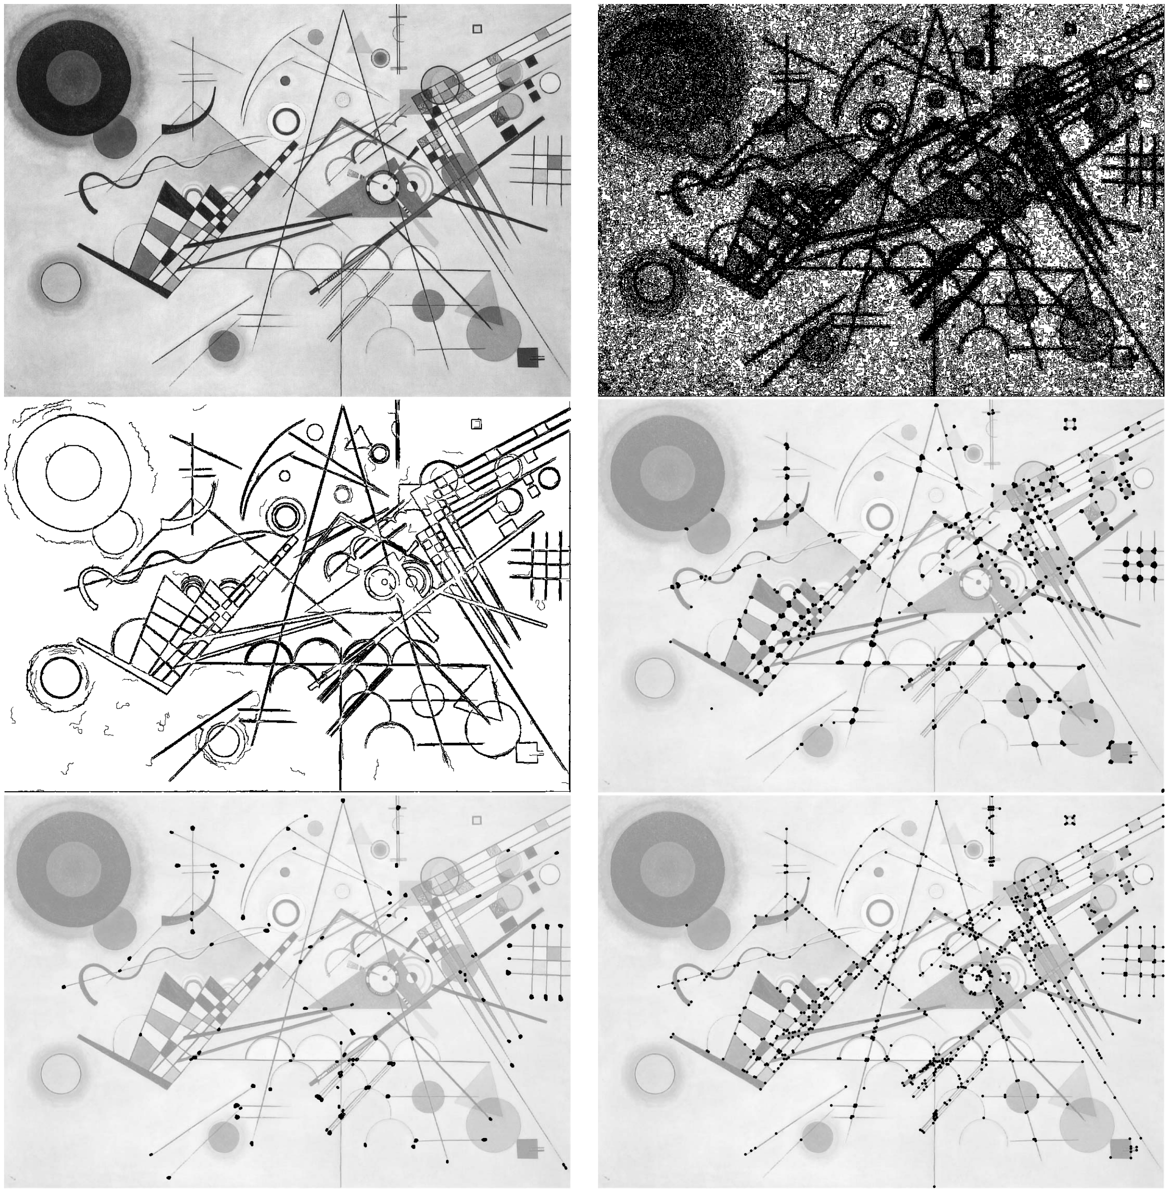

Figure 2. Good continuation and breaking detection. Top: Original image (1090×755), level lines (multiples of 3 for display). Middle: good continuations, and corner/junctions. Bottom row: meaningful terminators, and corners detected by Harris [12] detector. In Harris detector, we have to fix three parameters: 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. With Harris' detector, corners are detected near edges, probably because of ringing effects.

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 [22] or active contours [16]).

The 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.

Acknowledgments

The author sincerily thanks Agnès Desolneux, Lionel Moisan and Jean-Michel Morel for introducing him to the "a contrario detection theory", for their advice and all valuable conversation.

References

[1] L. Alvarez and F. Morales. Affine morphological multiscale analysis of corners and multiple junctions. *International Journal of Computer Vision*, 25(2):95–107, 1997.

[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] F. Cao. Contrast invariant, parameterless detection of good continuations, corners and terminators. Technical Report PI-1487, IRISA, 2002.

[5] V. Caselles, B. Coll, and J. Morel. A Kanizsa program. In *Progress in Nonlinear Differential Equations and their Applications*, volume 25, pages 35–55, 1996.

[6] R. Deriche and G. Giraudon. A computational approach for corner and vertex detection. *International Journal of Computer Vision*, 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. Morel. Edge detection by Helmholtz principle. *Journal of Mathematical Imaging and Vision*, 14(3):271–284, 2001.

[9] A. Desolneux, L. Moisan, and J. Morel. Partial gestalts. http://www.cmla.ens-cachan.fr. To appear in IEEE. Trans. on PAMI, 2001.

[10] F. Guichard and J. Morel. *Iterative Image Filtering*. To appear, 2003.

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

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

[13] H. Helmholtz. *Handbuch der Physiologischen Optik*. Allgemeine Encyklopädie der Physik, 1867.

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

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

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

[17] T. Lindeberg and M. 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.

[18] J. 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.

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

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

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

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

[23] 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.

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