Finding Edges by a Contrario Detection of Periodic Subsequences

Mariano Tepper[1,*], Pablo Musé[2], Andrés Almansa[3], and Marta Mejail[1]

[1] Departamento de Computación, Facultad de Ciencias Exactas y Naturales, Universidad de Buenos Aires, Argentina
{mtepper,marta}@dc.uba.ar

[2] Instituto de Ingeniería Eléctrica, Facultad de Ingeniería, Universidad de la República, Uruguay
pmuse@fing.edu.uy

[3] CNRS - LTCI UMR5141, Telecom ParisTech, France
andres.almansa@telecom-paristech.fr

Abstract. A new method to detect salient pieces of boundaries in an image is presented. After detecting perceptually meaningful level lines, periodic binary sequences are built by labeling each point in close curves as salient or non-salient. We propose a general and automatic method to detect meaningful subsequences within these binary sequences. Experimental results show its good performance.

Keywords: topographic maps, level lines, periodic binary sequences, edge detection, Helmholtz principle.

1 Introduction

Shape plays a key role in our cognitive system: in the perception of shape lies the beginning of concept formation. Formally, shapes in an image can be defined by extracting contours from solid objects. Shapes can be represented and analyzed as the locus of an infinite number of points, which leads to level-sets methods [7].

We define an image as a function $u : \mathbb{R}^2 \to \mathbb{R}$ with continuous first derivatives. Level sets [7], or level lines, provide a complete and contrast-invariant image description. We define the boundaries of the connected components of a level set as a level line. These level lines have the following properties: (1) level lines are closed Jordan curves; (2) level lines at different levels are disjoint; (3) by topological inclusion, level lines form a partially ordered set.

We call the collection of level lines (along with their level) a topographic map. The inclusion relation allows to embed the topographic map in a tree-like representation. For extracting the level lines of a digital image we use the Fast Level Set Transform (FLST) [6] which computes level lines by bilinear interpolation. In general, the topographic map is an infinite set and so only quantized grey levels are considered, ensuring that the set is finite.

* Present address: Department of Electrical and Computer Engineering, University of Minnesota.

L. Alvarez et al. (Eds.): CIARP 2012, LNCS 7441, pp. 773–780, 2012.

Edge detectors, from which the most renowned is Canny's [1], rely on the fact that information is concentrated along contours (regions where contrast changes abruptly). From one side, only a subset of the topographic map is necessary to obtain a *perceptually* complete description. Going to a deeper level, perceptually important level lines, in general, are so because they contain contrasted *pieces*. In summary, we have to prune the topographic map and then prune inside the level lines themselves.

The search for perceptually important contours will focus on unexpected configurations, rising from the perceptual laws of Gestalt Theory [5]. From an algorithmic point of view, the main problem with the Gestalt rules is their qualitative nature. Desolneux et al. [3] developed the Computational Gestalt detection theory which seeks to provide a quantitative assessment of gestalts. It is primarily based on the Helmholtz principle which states that conspicuous structures may be viewed as exceptions to randomness. In this approach, there is no need to characterize the elements one wishes to detect but contrarily, the elements one wishes to avoid detecting, i.e., the background model. When an element sufficiently deviates from the background model, it is considered meaningful and thus, detected.

Within this framework, Desolneux et al. [3] proposed an algorithm to detect contrasted level lines in grey level images, called meaningful boundaries. Further improvements to this algorithm were proposed by Cao et al. [2] and by Tepper et al. [8,9]. In this work we address the dissection of meaningful boundaries, developing an algorithm to select salient pieces contained in them. Each level line is considered as a periodic binary sequence where, following a partial saliency model, each point is labeled as salient or non-salient. Then, the goal is to extract meaningful subsequences of salient points. In order to do so, we extend to the periodic case an algorithm for binary subsequence detection proposed by Grompone et al [4].

The remainder of this paper is organized as follows. In Section 2 we recall the definition by Tepper of meaningful boundaries[8]. In Section 3 we describe the proposed algorithm. In Section 4 we show examples that prove the pertinence of the approach and provide some final remarks.

2 Meaningful Contrasted Boundaries

We begin by formally explaining the meaningful boundaries algorithm, as defined by Tepper et al. [8,9].

Let C be a level line of the image u and let us denote by $\{x_i\}_{i=0\ldots n-1}$ the n regularly sampled points of C, with arc-length two pixels, which in the a contrario noise model are assumed to be independent. In particular the gradients at these points are independent random variables (the image gradient norm $|Du|$ can be computed by standard finite differences on a 2×2 neighborhood). We note by μ_k $(0 \leq k < n)$ the k-th value of the values $|Du|(x_i)$ sorted in ascending order.

The detection algorithm consists in rejecting the null hypothesis $\mathcal{H}_0$: *the line C with contrasts $\{\mu_k\}_{k=0\ldots n-1}$ is observed only by chance.* For this we assume that the values of $|Du|$ are i.i.d., extracted from a noise image with the same gradient histogram as the image u itself.

Desolneux et al. [3] present a thorough study of the binomial tail $\mathcal{B}(n, k; p)$ and its use in the detection of geometric structures. The regularized incomplete beta function, defined by $I(x; a, b)$ is an interpolation $\widetilde{\mathcal{B}}$ of the binomial tail to the continuous domain $\widetilde{\mathcal{B}}(n, k; p) = I(p; k, n-k+1)$ where $n, k \in \mathbb{R}$ [3]. Additionally the regularized incomplete beta function can be computed very efficiently.

Let $H_c(\mu) = \mathrm{P}(|Du| > \mu)$. For a given line of length l, the probability under $\mathcal{H}_0$ that, some parts with total length greater or equal than $l_{(s,n)}(n-k)$ have a contrast greater than μ can be modeled by $\widetilde{\mathcal{B}}(n \cdot l_{(s,n)}, k \cdot l_{(s,n)}; H_c(\mu))$, where $l_{(s,n)} = \frac{l}{s \cdot n}$ acts as a normalization factor [8,9].

Definition 1. *Let $\mathcal{C}$ be a finite set of N_{ll} level lines of u. A level line $C \in \mathcal{C}$ is a ε-meaningful boundary if $N_{ll} \cdot K \cdot \min_{k<K} \widetilde{\mathcal{B}}(n \cdot l_{(2,n)}, k \cdot l_{(2,n)}; H_c(\mu_k)) < \varepsilon$ where K is a parameter of the algorithm. We also note*

$$k_{\min} = \underset{k<K}{\arg\min}\, \widetilde{\mathcal{B}}(n \cdot l_{(2,n)}, k \cdot l_{(2,n)}; H_c(\mu_k)) \tag{1}$$

The parameter K controls the number of points that we allow to be likely generated by noise, that is a line must have no more than K points with a "high" probability of belonging to the background model. It is simply chosen as a percentile of the total number of points in the line.

Def. 1 is motivated by the following proposition (we refer to the work by Tepper [8] for a complete proof).

Proposition 1. *The expected number of ε-meaningful boundaries, in a finite set of random level lines is smaller than ε.*

3 Boundary Clean-up by Detecting Meaningful Periodic Subsequences

Prop. 1 asserts that if a level line is a meaningful boundary, then it cannot be entirely generated in white noise (up to ε false detections on the average) but it can have parts that are likely to be contained in noise.

Cao et al. [2] propose to give an upper bound to the size of those parts. Assume that C is a piece of level line with L independent points, contained in a non-edge part, described by the noise model. The probability that L is larger than $l > 0$ needs to be estimated, knowing that $|Du| \geq \mu$. This is exactly the a posteriori length distribution $p(\mu; l) \stackrel{\text{def}}{=} P(L \geq l \mid |Du| \geq \mu)$. The estimation of this distribution was studied by Cao et al. [2].

Let us now consider an image u with N_{ll} (quantized) level lines. Let us also denote by N_l the number of all possible sampled subcurves of these level lines. ($N_l = \sum_{i=1}^{N_{ll}} n_i(n_i - 1)/2$, where n_i is the number of independent points in line i). As in Prop. 1, it can be proved that $N_l \cdot p(\mu; l)$ is an upper bound of the expected number of pieces of lines of length larger than l with gradient larger than μ. For a fixed μ, let be l such that $N_l \cdot p(\mu; l) < \varepsilon$. Then, we know that on a white noise image, on the average, we cannot observe more than ε pieces of level line with a length larger than l and a gradient everywhere larger than μ. Then one can define $\mathcal{L}(\mu) = \inf\{l, N_l \cdot p(\mu; l) < \varepsilon\}$ and keep every subcurve of any meaningful boundary with length equal or greater than $\mathcal{L}(\mu)$, where $|Du| \geq \mu$.

The value of μ can be seen as a new parameter of the method. Its value can be fixed arbitrarily using a conservative approach [2]: letting $|Du|$ be less than 1, means that edges with an accuracy less than one pixel may be detected. Thus, taking $\mu = 1$ is the least restrictive choice. For μ about 1, values of $\mathcal{L}(\mu)$ less than a few hundreds are obtained.

Since $\mathcal{L}(\mu)$ is a decreasing function of μ, fixing it at a small value produces large lengths. We are imposing that the contrasted pieces have to be very large and this is not always the case, as argued before. Furthermore the probability distribution $p(\mu; l)$ has to be estimated. We propose to take a different path to remove non-contrasted boundary parts.

In Def. 1, pieces of a meaningful boundary are explicitly allowed to be generated in white noise. We are certainly not interested in these pieces and this relaxation responds to the fact that we want to retrieve the remaining pieces of that boundary (i.e. edge region). The desired detection of contrasted parts in a boundary is very close in spirit to periodic subsequence detection.

3.1 Detecting Periodic Subsequences

Grompone et al. [4] proposed a method for accurately detecting straight line segments in a digital image. It is based on the Helmholtz principle and hence parameterless. In the authors' words, "at the core of the work lies a new way to interpret binary sequences in terms of unions of segments".

A sequence $S = (s_i)_{1 \leq i \leq L}$ of length L is binary if $\forall i,\ s_i \in \{0, 1\}$. A subsequence $a \subseteq S$ is defined by a pair of indices $\left(a^{(1)}, a^{(2)}\right)$ with $1 \leq a^{(1)} < a^{(2)} \leq L$, such that $\left(\forall s_i,\ a^{(1)} \leq i \leq a^{(2)}\right)\ s_i \in a$. Given a binary sequence S of length L, an n-subsequence is an n-tuple $(a_1, \ldots, a_n)$ of n disjoints subsequences $a_i \subseteq S$. The set of all n-subsequences in S will be denoted by $\mathcal{M}(n, S)$. We define $k(a) = \#\{s_i \mid i \in [a^{(1)}, a^{(2)}] \wedge s_i = 1\}$ and $l(a) = a^{(2)} - a^{(1)} + 1$ (i.e. the length of a). Notice that $\#\mathcal{M}(n, S) = \binom{L}{2n}$ [4].

Definition 2. (Grompone et al. [4]) *Given a binary sequence S of length L, an n-subsequence $(a_1, \ldots, a_n)$ in $\mathcal{M}(n, S)$ is said ε-meaningful if*

$$\mathrm{NFA}(a_1, \ldots, a_n) \stackrel{\mathrm{def}}{=} \binom{L}{2n} \prod_{i=1}^{n} (l(a_i) + 1)\mathcal{B}(l(a_i), k(a_i), p) < \varepsilon \tag{2}$$

where $p = \Pr(s_i = 1), 1 \leq i \leq L$. This number is called number of false alarms (NFA) of $(a_1, \ldots, a_n)$.

Proposition 2. *The expected number of ε-meaningful n-subsequences in a random binary sequence is smaller than ε.*

We refer to the work by Grompone et al. [4] for a complete proof.

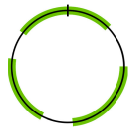

Fig. 1. A periodic sequence where runs are represented in green. If treated as a non-periodic sequence, any subsequence detector would detect four subsequences at best, when in fact the desired result is to detect three subsequences.

A run in S is a maximal subsequence only containing ones, i.e.

$$\left(\forall i \in [a^{(1)}, a^{(2)}],\ s_i = 1\right) \wedge \left(a^{(1)} = 1 \vee s_{a^{(1)}-1} = 0\right) \wedge \left(a^{(2)} = L \vee s_{a^{(2)}+1} = 0\right).$$

One can restrict the search for n-subsequences to the ones where each of the n subsequences starts at a run start and ends at a run end [4]. We denote by R the number of runs in S.

Definition 3. *Given a binary sequence S, its maximal ε-meaningful subsequence $(a_1, \ldots, a_n)^*$ is defined as*

$$(a_1, \ldots, a_n)^* \stackrel{\text{def}}{=} \underset{\substack{1 \leq n \leq R \\ (a_1, \ldots, a_n) \in \mathcal{M}(n,S)}}{\arg\min} \text{NFA}(a_1, \ldots, a_n).$$

We propose now to extend the above definitions to support periodic binary sequences. A binary sequence $S = (s_i)_{1 \leq i \leq L}$ is made periodic by considering L its period. Periodic sequences are different in nature from their non-periodic counterparts, see Fig. 1. A definition suitable for the periodic case is needed.

In the periodic case, a subsequence must be defined more carefully. Now a subsequence $a \subseteq S$, defined by a pair of indices $\left(a^{(1)}, a^{(2)}\right)$, can belong to one of two different types:

Intra-subsequences: if $a^{(1)} < a^{(2)}$ then the non-periodic definition holds, i.e., $1 \leq a^{(1)} < a^{(2)} \leq L$, and $\left(\forall s_i,\ a^{(1)} \leq s_i \leq a^{(2)}\right)\ s_i \in a$.

Inter-subsequences: if $a^{(1)} > a^{(2)}$ $\left(\forall s_i,\ 1 \leq s_i \leq a^{(2)} \vee a^{(1)} \leq s_i \leq L\right)\ s_i \in a$.

Runs are modified accordingly to also cover inter-subsequences. Given a periodic binary sequence S of period L, a periodic n-subsequence is an n-tuple $(a_1, \ldots, a_n)$ of n disjoints subsequences $a_i \subseteq S$. The set of all n-subsequences in S will be denoted by $\mathcal{M}(n, S)$.

We define $k(a) = \#\{s_i \mid i \in [a^{(1)}, a^{(2)}] \wedge s_i = 1\}$ and the length of a as

$$l(a) = \begin{cases} a^{(2)} - a^{(1)} + 1, & \text{if } a \text{ is an intra-subsequence;} \\ a^{(2)} + L - a^{(1)} + 1, & \text{if } a \text{ is an inter-subsequence.} \end{cases}$$

Notice that $\#\mathcal{M}(n, S) = 2\ \binom{L}{2n}$ since from each pair of points in S two subsequences can be constructed.

Definition 4. *Given a periodic binary sequence S of period L, an n-subsequence $(a_1, \dots, a_n)$ in $\mathcal{M}(n, S)$ is said ε-meaningful if*

$$\mathrm{NFA}(a_1, \dots, a_n) \stackrel{\text{def}}{=} 2 \binom{L}{2n} \prod_{i=1}^{n} (l(a_i) + 1)\ \mathcal{B}(l(a_i), k(a_i), p) < \varepsilon$$

where $p = \Pr(s_i = 1), 1 \leq i \leq L$. This number is called number of false alarms (NFA) of $(a_1, \dots, a_n)$.

Proposition 3. *The expected number of ε-meaningful n-subsequences in a random periodic binary sequence is smaller than ε.*

Proof. This proof follows closely the one by Grompone et al. [4] but adapted to periodic sequences. The expected number of ε-meaningful n-subsequences is given by

$$\mathbb{E}\left(\sum_{(a_1,\dots,a_n) \in \mathcal{M}(n,S)} \mathbf{1}_{\mathrm{NFA}(a_1,\dots,a_n) < \varepsilon} \right) = \sum_{(a_1,\dots,a_n) \in \mathcal{M}(n,S)} \mathrm{P}\left(\mathrm{NFA}(a_1, \dots, a_n) < \varepsilon\right).$$

$\mathrm{NFA}(a_1, \dots, a_n) < \varepsilon$ implies that $\prod_{i=1}^{n} \mathcal{B}(l(a_i), k(a_i), p) < \frac{\varepsilon}{2 \binom{L}{2n} \prod_{i=1}^{n}(l(a_i)+1)}$.

Let $U_i = \mathcal{B}(l(a_i), k(a_i), p)$ be a random variable, let $\alpha \in \mathbb{R}^+$, and let $\mathrm{P}_U^\alpha = \mathrm{P}\left(\prod_{i=1}^{n} U_i < \alpha\right)$. Then,

$$\mathrm{P}_U^\alpha = \sum_{u_2,\dots,u_n} \mathrm{P}\left(\prod_{i=1}^{n} U_i < \alpha \middle| U_2 = u_2, \dots, U_n = u_n \right) \mathrm{P}\left(U_2 = u_2, \dots, U_n = u_n\right).$$

Since the a_i are disjoint, the U_i are independent. Then

$$\mathrm{P}_U^\alpha = \sum_{u_2,\dots,u_n} \mathrm{P}\left(\prod_{i=1}^{n} U_i < \frac{\alpha}{u_2 \dots u_n} \right) \cdot \mathrm{P}\left(U_2 = u_2, \dots, U_n = u_n\right).$$

Using the classical lemma $\mathrm{P}(U_i < \alpha) < \alpha$, that $\mathrm{P}\left(U_2 = u_2, \dots, U_n = u_n\right) \leq \mathrm{P}\left(U_2 \leq u_2, \dots, U_n \leq u_n\right)$, and that there are $l(a_i) + 1$ possible values for U_i,

$$\mathrm{P}\left(\prod_{i=1}^{n} U_i < \alpha \right) < \prod_{i=2}^{n} (l(a_i) + 1)\ \alpha < \prod_{i=1}^{n} (l(a_i) + 1)\ \alpha.$$

Let us recall that $\#\mathcal{M}(n, S) = 2 \binom{L}{2n}$, then setting $\alpha = \frac{\varepsilon}{2 \binom{L}{2n} \prod_{i=1}^{n}(l(a_i)+1)}$ gives the wanted result. □

The maximality rule from Def. 3 holds unchanged in the periodic case.

On the implementation side, Grompone et al. [4] describe a dynamic programming scheme for the non-periodic case that eases the heavy computational burden. We show now that implementing the algorithm for detecting periodic subsequences is indeed straightforward.

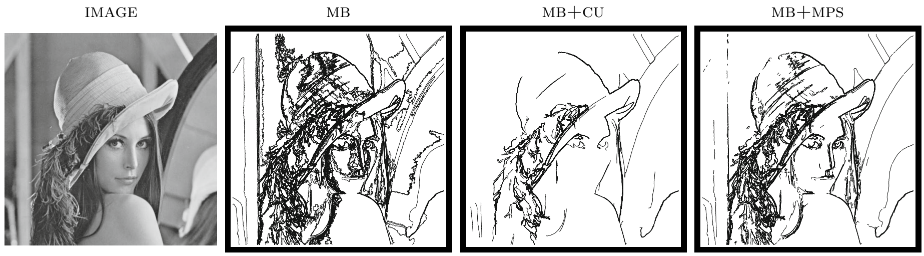

Fig. 2. Comparison of the results obtained with both clean-up algorithms. The one by Cao et al. (CU) [2] produces underdetection; this is corrected by using MPS.

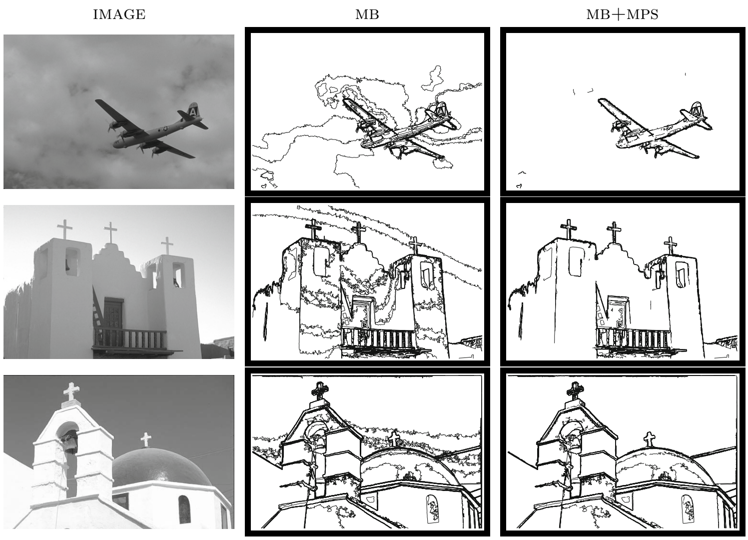

Fig. 3. Results of the presented clean-up algorithm. MPS eliminates the vast majority of the unwanted pieces of level line.

We begin by shifting the periodic sequence S (with R runs), to transform inter-subsequences into intra-subsequences. A circular shift to the left is used. We then form a non-periodic sequence $S^{(2)}$ of length $2L$ from two periods of the periodic sequence S of period L. Let $R^{(2)}$ be the number of runs in $S^{(2)}$. Two key tricks allow us to solve the problem: (1) restrict the number of tested subsequences. In the non-periodic case, we test for n-subsequences for $S^{(2)}$ where $1 \leq n \leq R^{(2)}$. In the periodic case, we only test for n-subsequences where $1 \leq n \leq R$; (2) subsequences longer than L are not tested. With these two restrictions, one can simply detect non-periodic subsequences in non-periodic sequence $S^{(2)}$ and the result will be optimal.

4 Results and Final Remarks

Before applying the detector of meaningful periodic subsequences (MPS) to any boundary, we need to binarize it since its contrast (or its regularity) takes on real values. This former problem is solved by thresholding on the contrast (or on the regularity). In this direction, we claim that a natural choice is $\mu_{k_{\min}}$ (see Def. 1, p. 775). A maximal ε-meaningful boundary is thus converted into a periodic binary sequence. We want to apply the periodic subsequence detection algorithm from Def. 4 and 3 to that sequence. The only parameter left is $p = \Pr(s_i = 1), 1 \leq i \leq L$ and it is straightforward defined as $p \stackrel{\text{def}}{=} H_c(\mu_{k_{\min}})$.

We finally define the following clean-up rule: *For any meaningful boundary, keep every subcurve belonging to its maximal 1-meaningful subsequence.*

This clean-up mechanism does not impose a minimal length to contrasted parts. The length is adjusted automatically, by choosing the most meaningful subsequence in the level line. As an additional advantage, there is no need to estimate any probability distribution. Fig. 2 shows an example of the benefits of the proposed clean-up method over the one by Cao et al. [2]. Their version clearly produces underdetection: visually important structures are missed (notice the face in the third image). The proposed algorithm produces a very mild overdetection: some small noisy parts are not eliminated but no important structure is lost. Fig. 3 shows two more examples on images from the BSD database. Notice that, on the last row, MPS does not remove a few pieces of lines that should be removed (e.g., the lower wall and the roof). This does not occur because of a failure in MPS, but because of a faulty binarization, that is, the $\mu_{k_{\min}}$ was not optimal in those cases.

In summary, we presented a general and fully automatic algorithm to detect meaningful subsequences within periodic binary sequences. We apply it to select salient pieces of level lines in an image, showing good results on natural images.

References

1. Canny, J.: A Computational Approach to Edge Detection. IEEE Transactions on Pattern Analysis and Machine Intelligence 8(6), 679–698 (1986)
2. Cao, F., Lisani, J.L., Morel, J.M., Musé, P., Sur, F.: A Theory of Shape Identification. Lecture Notes in Mathematics, vol. 1948. Springer (2008)
3. Desolneux, A., Moisan, L., Morel, J.M.: From Gestalt Theory to Image Analysis, vol. 34. Springer (2008)
4. Grompone, R., Jakubowicz, J., Morel, J.M., Randall, G.: On Straight Line Segment Detection. Journal of Mathematical Imaging and Vision 32(3), 313–347 (2008)
5. Kanizsa, G.: Organization in Vision: Essays on Gestalt Perception. Praeger (1979)
6. Monasse, P., Guichard, F.: Fast Computation of a Contrast Invariant Image Representation. IEEE Transactions on Image Processing 9(5), 860–872 (2000)
7. Serra, J.: Image Analysis and Mathematical Morphology. Academic Press (1983)
8. Tepper, M.: Detecting clusters and boundaries: a twofold study on shape representation. Ph.D. thesis, Universidad de Buenos Aires (March 2011)
9. Tepper, M., Gómez, F., Musé, P., Almansa, A., Mejail, M.: Morphological Shape Context: Semi-locality and Robust Matching in Shape Recognition. In: Bayro-Corrochano, E., Eklundh, J.-O. (eds.) CIARP 2009. LNCS, vol. 5856, pp. 129–136. Springer, Heidelberg (2009)