Finding contrasted and regular edges by a contrario detection of periodic subsequences

Mariano Tepper [a],[*],[1], Pablo Musé [b], Andrés Almansa [c], Marta Mejail [d]

[a] Department of Electrical and Computer Engineering, Duke University, USA
[b] Instituto de Ingeniería Eléctrica, Facultad de Ingeniería, Universidad de la República, Uruguay
[c] CNRS - LTCI UMR5141, Telecom ParisTech, France
[d] Departamento de Computación, Facultad de Ciencias Exactas y Naturales, Universidad de Buenos Aires, Argentina

ARTICLE INFO

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, when tested with different saliency criteria, such as contrast, regularity, and the combination of both.

© 2013 Elsevier Ltd. All rights reserved.

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 [1].

We define an image as a lower (or upper) semi-continuous function $u : \mathbb{R}^2 \to \mathbb{R}$. Level sets [1], 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) [2] 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.

Edge detectors, from which the most renowned is Canny's [3], 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 [4]. From an algorithmic point of view, the main problem with the Gestalt rules is their qualitative nature. Desolneux et al. [5] 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. [5] 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. [6] and by Tepper et al. [7], which include for example, the use of regularity as the saliency measure instead of the contrast.

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. To do so, in this work we extend to the periodic case an algorithm for detecting binary subsequences.

Grompone et al. [8] proposed a method for accurately detecting straight line segments in a digital image. It is based on the

* Corresponding author. Tel.: +1 919 641 4558.
E-mail addresses: mariano.tepper@duke.edu (M. Tepper), pmuse@fing.edu.uy (P. Musé), andres.almansa@telecom-paristech.fr (A. Almansa), marta@dc.uba.ar (M. Mejail).

Helmholtz principle and hence it is 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." In the same spirit, pieces of level lines correspond to object contours and can be recovered extending this approach to the periodic case (level lines are closed curves).

The remainder of this paper is organized as follows. We begin, in Section 2, by explaining the proposed algorithm for detecting periodic subsequences in binary sequences. Then, in Section 3 we recall several variants of the meaningful boundaries algorithm [7]. In Section 4 we show how to apply the proposed algorithm to the problem of selecting salient pieces of level lines. In Section 5 we discuss the pertinence of the approach and provide some final remarks.

2. Detecting periodic subsequences

A sequence $S = (s_i)_{1 \le i \le L}$ 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 $(a^{(1)}, a^{(2)})$ with $1 \le a^{(1)} < a^{(2)} \le L$, such that $(\forall s_i, a^{(1)} \le i \le a^{(2)}) s_i \in a$. 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 $k(a) = \#\{s_i | i \in [a^{(1)}, a^{(2)}] \wedge s_i = 1\}$ and $l(a) = a^{(2)} - a^{(1)} + 1$ (i.e. the length of a). Notice that $\#\mathcal{M}(n, S) = \binom{L}{2n}$ [8].

Definition 1. (Grompone et al. [8]). 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

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

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

Proposition 1. 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. [8] for a complete proof.

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

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

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 [8]. We denote by R the number of runs in S.

Definition 2. Given a binary sequence S, its maximal ε−meaningful subsequence $(a_1, \ldots, a_n)^*$ is defined as

$$(a_1, \ldots, a_n)^* = \underset{\substack{1 \le n \le R \\ (a_1, \ldots, a_n) \in \mathcal{M}(n, S)}}{\mathrm{argmin}} \ \mathrm{NFA}(a_1, \ldots, a_n).$$

We propose now to extend the above definitions to support periodic binary sequences. A binary sequence $S = (s_i)_{1 \le i \le L}$ is made periodic by considering L as 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 $(a^{(1)}, a^{(2)})$, can belong to one of the two different types:

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

inter–subsequences: if $a^{(1)} > a^{(2)} (\forall s_i, 1 \le s_i \le a^{(2)} \vee a^{(1)} \le s_i \le L) s_i \in a$.

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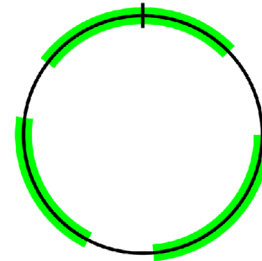
Fig. 1. A periodic sequence where runs are represented in green (light gray). If treated as a non-periodic sequence, any subsequence detector would detect four subsequences at best, when in fact the desired result is to detect three subsequences.

$a_i \subseteq S$. The set of all n-subsequences in S will be denoted by $\mathcal{M}(n, S)$.

We define $k(a) = \#\{s_i | i \in [a^{(1)}, a^{(2)}] \wedge s_i = 1\}$ 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 3. Given a periodic binary sequence S of period L, an n-subsequence $(a_1, \ldots, a_n)$ in $\mathcal{M}(n, S)$ is said ε−meaningful if

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

where $p = \mathrm{Pr}(s_i = 1)$, $1 \le i \le L$. 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 periodic binary sequence is smaller than ε.

Proof. This proof follows closely the one by Grompone et al. [8] but adapted to periodic sequences. The expected number of ε−meaningful n-subsequences is given by

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

$\mathrm{NFA}(a_1, \ldots, a_n) < \varepsilon$ 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 $\mathrm{P}_U^\alpha = \mathrm{P}(\prod_{i=1}^{n} U_i < \alpha)$. Then

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

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

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

Using the classical lemma $\mathrm{P}(U_i < \alpha) < \alpha$, that $\mathrm{P}(U_2 = u_2, \ldots, U_n = u_n) \le \mathrm{P}(U_2 \le u_2, \ldots, U_n \le u_n)$, 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 $\alpha = \varepsilon / 2\binom{L}{2n} \prod_{i=1}^{n} (l(a_i) + 1)$ gives the wanted result. □

The maximality rule from Definition 2 holds unchanged in the periodic case.

On the implementation side, Grompone et al. [8] 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.

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.

3. Meaningful boundaries

We briefly reproduce the formal definitions that lead to the meaningful boundaries algorithm, as defined by Tepper et al. [7]. Several criteria are used to determine the saliency of a boundary: contrast, regularity and the combination of both. All these variants will be used to test the proposed algorithm.

3.1. Contrasted meaningful boundaries

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 of 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. [5] 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 $\tilde{\mathcal{B}}$ of the binomial tail to the continuous domain $\tilde{\mathcal{B}}(n, k; p) = I(p; k, n-k+1)$ where $n, k \in \mathbb{R}$ [5]. Additionally the regularized incomplete beta function can be computed very efficiently.

Let $H_c(\mu) = P(|Du| > \mu)$. 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 $\tilde{\mathcal{B}}(n \cdot l_{(s,n)}, k \cdot l_{(s,n)}; H_c(\mu))$, where $l_{(s,n)} = l/s \cdot n$ acts as a normalization factor [7].

Definition 4. 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

$$\text{NFA}_K(C) = N_{ll} \cdot K \cdot \min_{k < K} \tilde{\mathcal{B}}(n \cdot l_{(2,n)}, k \cdot l_{(2,n)}; H_c(\mu_k)) < \varepsilon, \qquad (2)$$

where K is a parameter of the algorithm. We also note

$$k_{\min} = \operatorname*{argmin}_{k < K} \tilde{\mathcal{B}}(n \cdot l_{(2,n)}, k \cdot l_{(2,n)}; H_c(\mu_k)). \qquad (3)$$

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.

Definition 4 is motivated by the following proposition (we refer to the work by Tepper et al. [7] for a complete proof).

Proposition 3. *The expected number of ε−meaningful boundaries in a finite set of random level lines is smaller than ε.*

3.2. Combining contrast and smoothness

As already stated, in natural images contrasted boundaries often locally coincide with object edges. Thus, they are also incidentally smooth. Active contours [9] rely on this combination of good contrast and smoothness to provide well localized contours. In this section, we reproduce results that use smoothness as the saliency measure in the a contrario detection process [7].

Let C be a rectifiable planar curve, parameterized by its length. Let l be the length of C and $x = C(\tau) \in C$. Without loss of generality, we assume that $\tau = 0$.

Definition 5. (Cao et al. [10]). Let $s > 0$ be a fixed positive value such that $2s < l$. We call regularity of C at x (at scale s) the quantity

$$R_s(x) = \frac{\max(|x - C(-s)|, |x - C(s)|)}{s}, \qquad (4)$$

where $|x_i - x_j|$ represents the Euclidean distance between x_i and x_j.

Fig. 2 visually explains the pertinence of this definition. Only when one of the subcurves $C((-s, 0))$ or $C((0, s))$ is a line segment, $R_s(x) = 1$; in all other cases $R_s(x) < 1$. When s is small enough, regularity is inversely proportional to the curve's curvature around x [10].

The question about the choice of s arises naturally and was studied in detail by Cao et al. [10] and Musé [11]. We will limit ourselves to state that a larger value of s (thus at less local scale of analysis) is more robust to noise. On the other side, s should not be too large either. In practice, and following Cao et al. [10] one may safely set $s = 5$, which is the value we use in our experiments.
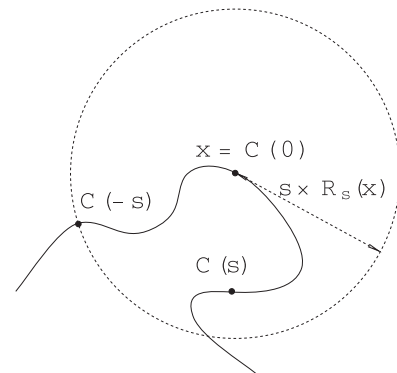
Fig. 2. Reproduced from the work by Cao et al. [10]. The regularity at x is obtained by comparing the radius of the circle with s. The radius is equal to s if and only if the curve is a straight line. If the curve has a large curvature, the radius will be small compared to s.
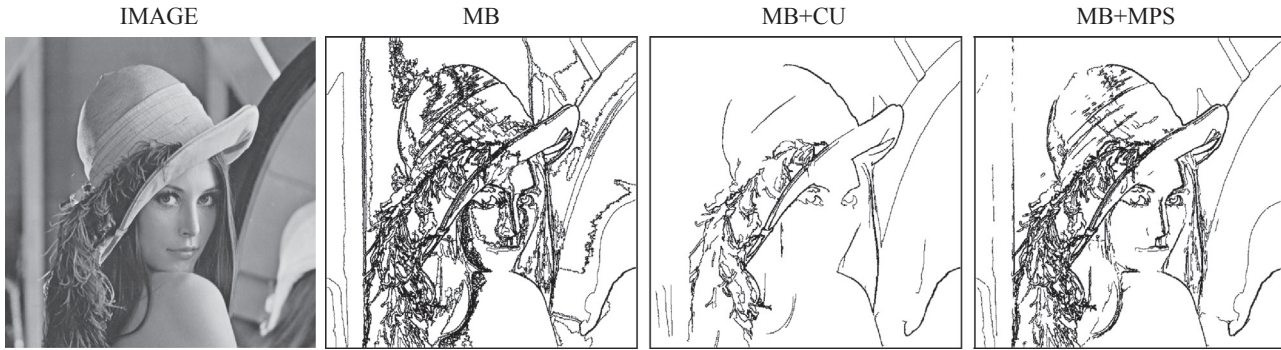
IMAGE   MB   MB+CU   MB+MPS

Fig. 3. Comparison of the results obtained with both clean-up algorithms. The one by Cao et al. (CU) [6] produces underdetection; this is corrected by using MPS.
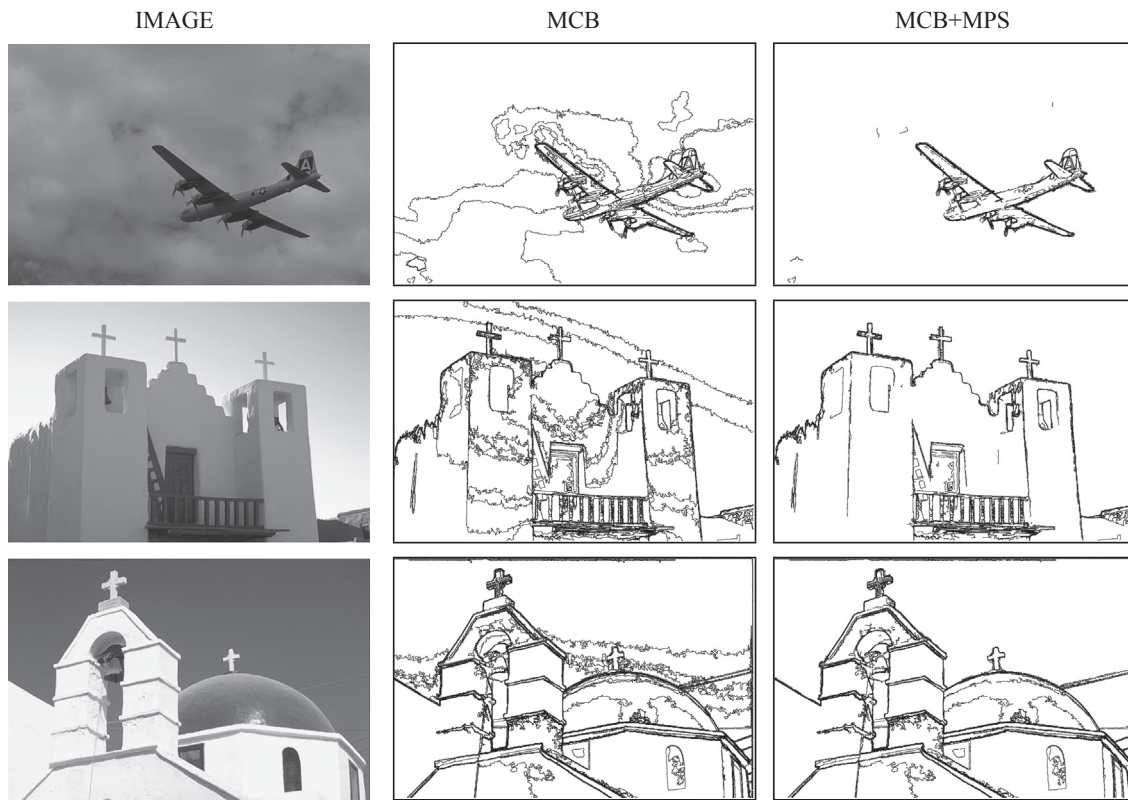
IMAGE   MCB   MCB+MPS

Fig. 4. Results of the presented clean-up algorithm on MCB. MPS eliminates the vast majority of the unwanted pieces of level line.

MCB+MPS
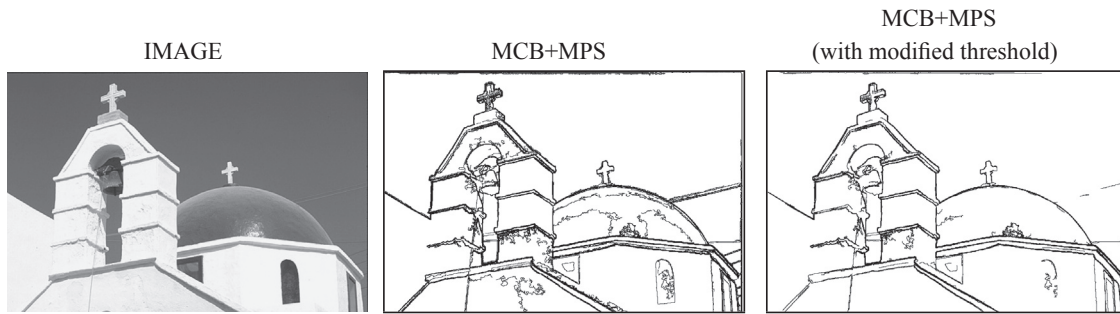IMAGE   MCB+MPS   (with modified threshold)

Fig. 5. Altering the binarization threshold might help, in some cases, to eliminate spurious structures. We modified it in a simple way that allows to visualize its effect, setting for the result on the right the threshold $p_2 = \min(6 \ast p_1, 1)$, where p_1 is the threshold for the result on the center. Notice that the window is better recovered with the original threshold.

Let us denote by $H_s(r)$ the distribution of the regularity in white noise level lines, i.e.,

$$H_s(r) = P(R_s(x) > r, x \in C, C \text{ is a white noise level line}), \quad (5)$$

which depends only on s and can be empirically estimated.

Again, the curve detection algorithm consists of adequately rejecting the null hypothesis $\mathcal{H}_0$: *the values of $|R_s|$ are i.i.d., extracted from a noise image.* We assume that, in the background model, contrast and regularity are independent.

Definition 6. Let $\mathcal{C}$ be a finite set of N_{ll} level lines of u. A level line $C \in \mathcal{C}$ is a $\varepsilon-$meaningful regular boundary (MRB) if

$$\text{NFA}_K^R(C) = N_{ll} K_s \min_{k < K_s} \tilde{\mathcal{B}}(n \cdot l_{(2s,n)}, k \cdot l_{(2s,n)}; H_s(\rho_k)) < \varepsilon, \quad (6)$$

and K_s is a parameter of the algorithm. We also note

$$k_{\min} = \underset{k < K_s}{\text{argmin}}\, \tilde{\mathcal{B}}(n \cdot l_{(2s,n)}, k \cdot l_{(2s,n)}; H_s(\rho_k)). \quad (7)$$

Proposition 4. *The expected number of $\varepsilon-$meaningful contrasted regular boundaries in a finite set E of random curves is smaller than ε.*

A proof is given in [7]. We can also combine the two criteria, as done by Tepper et al. [7], obtaining the following formulation.

Definition 7. Let $\mathcal{C}$ be a finite set of N_{ll} level lines of u. A level line $C \in \mathcal{C}$ is a $\varepsilon-$meaningful contrasted regular boundary (MCRB) if

$$\text{NFA}_K^{CR}(C) = N_{ll}\, K_c\, K_s \max\begin{pmatrix} \min_{k < K_c} I_c(C, k)^2 \\ \min_{k < K_s} I_s(C, k)^2 \end{pmatrix} < \varepsilon, \quad (8)$$

where

$$I_c(C, k) = \tilde{\mathcal{B}}(n \cdot l_{(2,n)}, k \cdot l_{(2,n)}; H_c(\mu_k)),$$
$$I_s(C, k) = \tilde{\mathcal{B}}(n \cdot l_{(2s,n)}, k \cdot l_{(2s,n)}; H_s(\rho_k)),$$

and K_c and K_s are parameters of the algorithm. In this case, $k_{\min}$ is defined as in Eq. (3) or as in Eq. (7), depending on whether contrast or regularity attains the maximum.

Here K_c and K_s have the same meaning as K in Definition 4 and they are also set as a percentile of the total number of points in the curve.

Proposition 5. *The expected number of $\varepsilon-$meaningful contrasted regular boundaries in a finite set E of random curves is smaller than ε.*
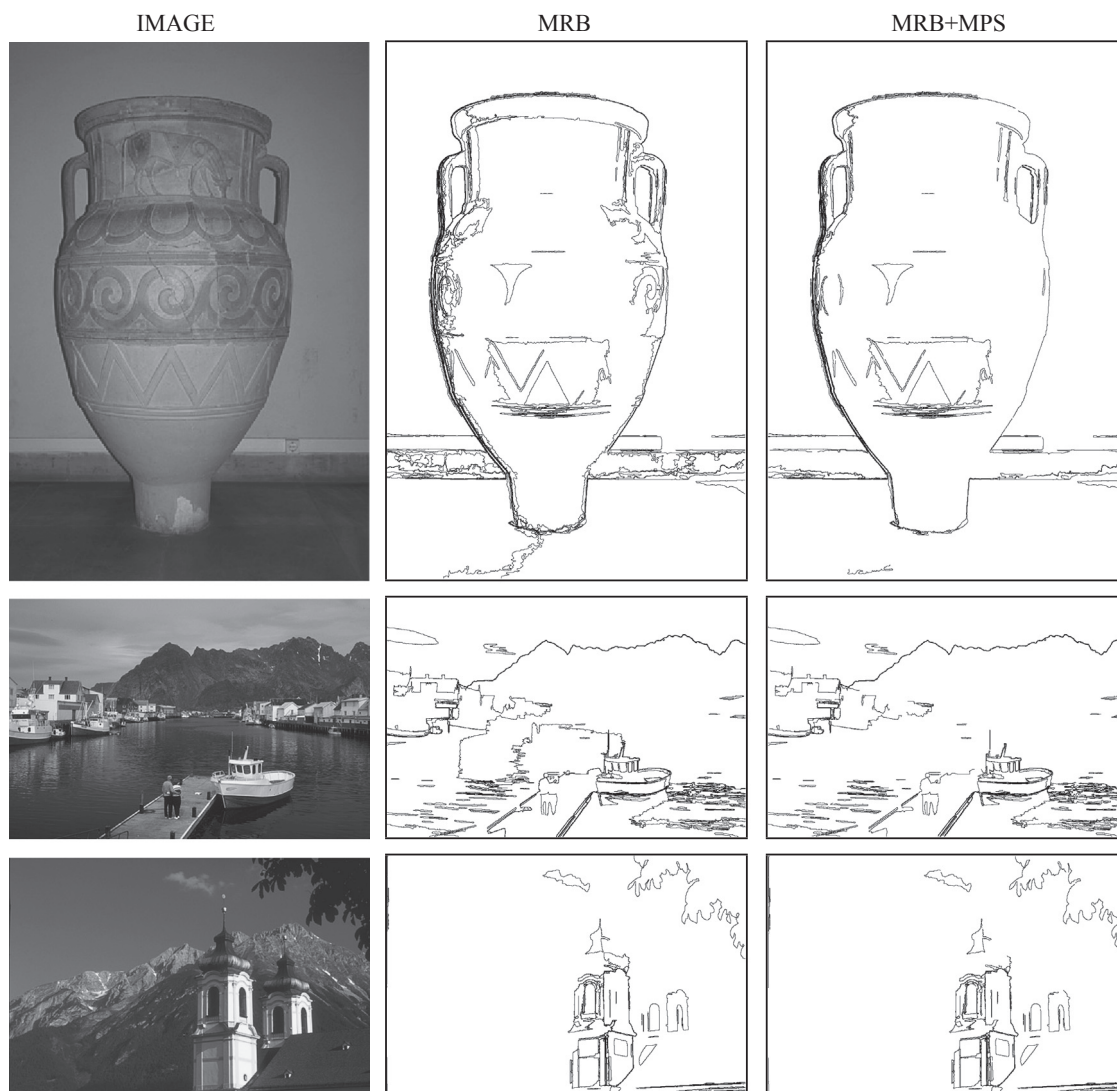

| IMAGE | MRB | MRB+MPS |

Fig. 6. Results of the presented clean-up algorithm on MRB. MPS eliminates the vast majority of the unwanted pieces of level line.

A proof is given in [7]. Definition 7 exhibits some interesting properties:

- A contrasted but irregular curve will not be detected.
- A regular but non-contrasted curve will not be detected.
- An irregular and non-contrasted curve will not be detected.
- A regular and contrasted curve will be detected.

Both gestalts, i.e., contrast and good continuation, interact in a novel way: they compete for the "control" of the curve.

Meaningful boundaries are further pruned to eliminate the packet effect (i.e., a bundle of level lines passing through a very salient edge) using the maximality procedure detailed in [7].

4. Boundary clean-up

Proposition 3 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. [6] 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 $p(\mu; l) = P(L \geq l \mid |Du| \geq \mu)$. The estimation of this distribution was studied by Cao et al. [6].

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 Proposition 3, 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 $\mathcal{L}(\mu) = \inf_l \{l, N_l \cdot p(\mu; l) < \varepsilon\}$ 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 [6]: 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 Definition 4, 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.

4.1. Boundary clean-up by detecting meaningful periodic subsequences

We now explain how to adapt the detector of meaningful periodic subsequences (MPS) for cleaning-up boundaries.

Before applying the detector 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 $\mu_{k_{\min}}$ and $\rho_{k_{\min}}$ are natural choices (see Definitions 4, 6, and 7). A maximal ε boundary
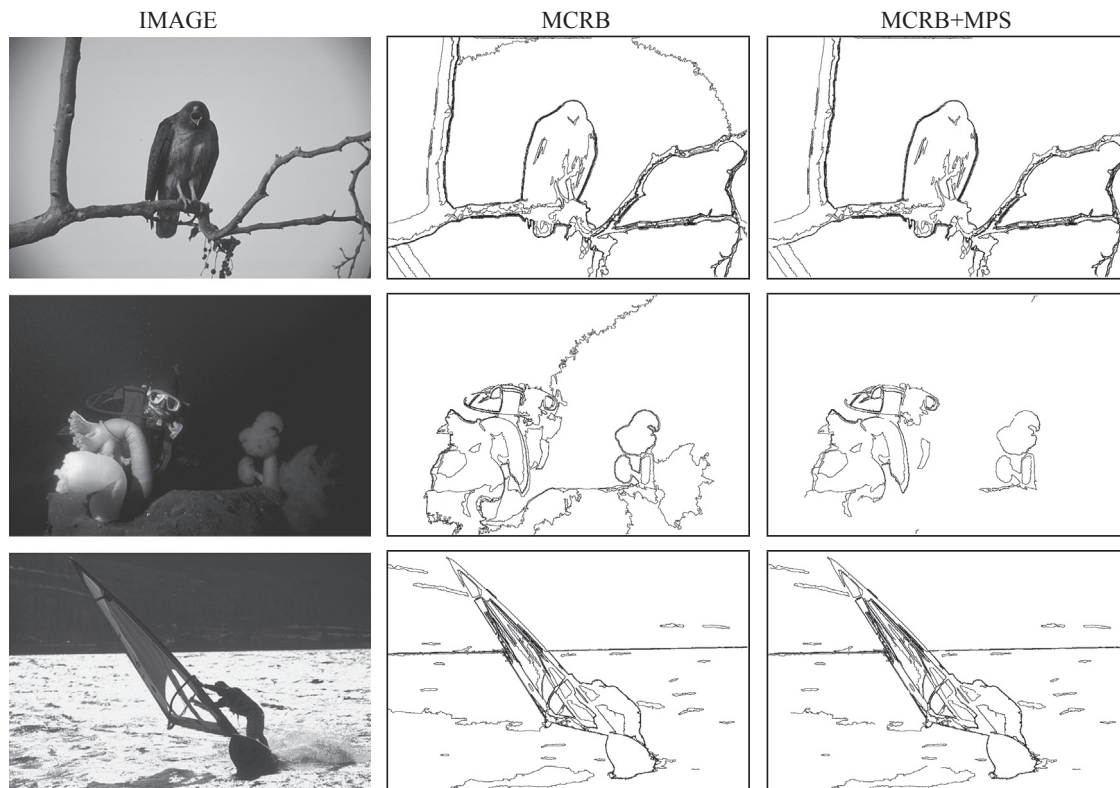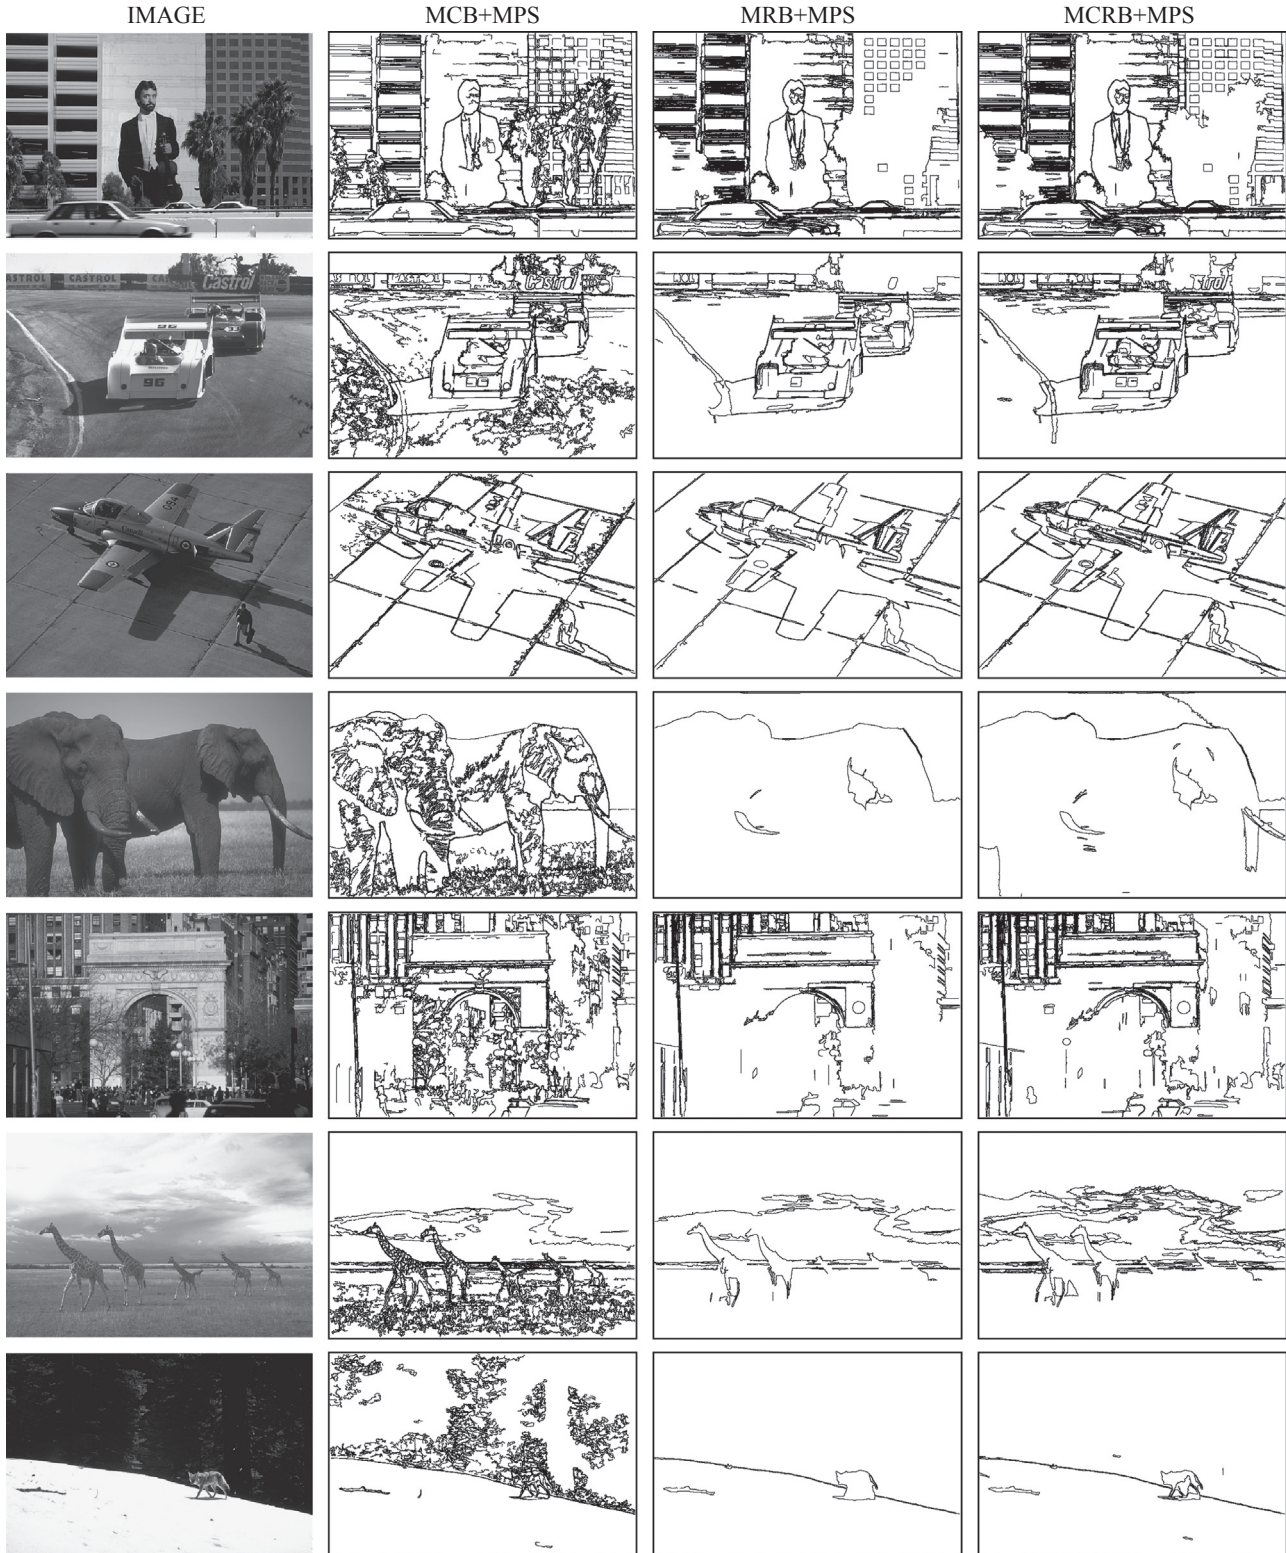

| IMAGE | MCRB | MCRB+MPS |

Fig. 7. Results of the presented clean-up algorithm on MCRB. MPS eliminates the vast majority of the unwanted pieces of level line.

is thus converted into a periodic binary sequence. We want to apply the periodic subsequence detection algorithm from Definitions 3 and 2 to that sequence. The only parameter left is $p = \Pr(s_i = 1), 1 \leq i \leq L$ and it is straightforward defined, for each boundary, as $p = H_c(\mu_{k_{\min}})$ or $p = H_s(\rho_{k_{\min}})$, depending on the case.

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

Algorithm 1 shows a possible procedure to obtain all ε−meaningful boundaries and then apply the proposed clean-up method to them.

IMAGE	MCB+MPS	MRB+MPS	MCRB+MPS

Fig. 8. Comparison of the results obtained with the presented clean-up algorithm on the different versions of MB. MPS eliminates the vast majority of the unwanted pieces of level line. We recommend to zoom-in the results in the electronic version (they are vectorial graphics) for better visualization.

Algorithm 1. Computation of ε-meaningful boundaries in image u and application of the proposed clean-up procedure.

```
input  : An image u and a scalar ε.
output : A set of closed curves S_res.
S ← FLST(u);                          // Compute the set of level lines
S_res ← ∅;
for C ∈ S do
    nfa_C ← compute NFA_K(C), NFA_K^R(C), or NFA_K^CR(C);
    if nfa_C < ε then
        S_res ← S_res ∪ {C};
Discard non-maximal boundaries from S_res [7];
S_final ← ∅;
for C ∈ S_res do
    C̃ ← threshold C as detailed in Section 4.1;
    Add the maximal 1-meaningful subsequence (a_1, ..., a_n)* of C̃ to S_final;
return S_final;
```

5. Discussion

The proposed 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. 3 shows an example of the benefits of the proposed clean-up method over the one by Cao et al. [6]. 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.

Figs. 4, 6, and 7 show examples on images from the BSD database [12]. In all cases MPS eliminates the vast majority of spurious pieces of level lines.

Notice that, in the last row of Fig. 4, 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. Fig. 5 depicts the same example, with a more aggressive binarization threshold. These spurious structures are now discarded. In any case, there is no universal threshold for all cases and it should be tuned for the application at hand.

Fig. 8 presents a comparison of the results obtained with the different meaningful boundaries algorithms.

Before concluding, we need to address an important objection. Instead of applying MPS to the meaningful boundaries, one could directly apply MPS to the raw collection of level lines. For each level lines, we could test several thresholds, as we do in MB, and select the periodic subsequence for which the NFA is minimized. This procedure would avoid using MB as an intermediary step. There are two main reasons for not following this more straightforward approach. First, we have a computational reason: applying MPS with different thresholds would be prohibitive; MB is much faster than MPS, allowing to render a practical algorithm. Second, since the maximality constraint [7] can be applied thanks to the inclusion properties between level lines. If we dissect the level lines before applying the maximality rule, we cannot apply it afterwards. A different, and much more complex, mechanism would have to be used to solve the packet effect.

In summary, we presented a general and fully automatic algorithm to detect meaningful subsequences within periodic binary sequences. As a useful application, we used it to select salient pieces of level lines in an image. We have also shown that different saliency criteria can be used, such as contrast, regularity or a combination of both. In all cases the proposed algorithm exhibits good results on natural images.

Conflict of Interest

None declared.

References

[1] J. Serra, Image Analysis and Mathematical Morphology, Academic Press, 1983.
[2] P. Monasse, F. Guichard, Fast computation of a contrast invariant image representation, IEEE Transactions on Image Processing 9 (2000) 860–872.
[3] J. Canny, A computational approach to edge detection, IEEE Transactions on Pattern Analysis and Machine Intelligence 8 (1986) 679–698.
[4] G. Kanizsa, Organization in Vision: Essays on Gestalt Perception, Praeger, 1979.
[5] A. Desolneux, L. Moisan, J.M. Morel, From Gestalt Theory to Image Analysis, vol. 34, Springer-Verlag, 2008.
[6] F. Cao, J.L. Lisani, J.M. Morel, P. Musé, F. Sur, A theory of shape identification, in: Lecture Notes in Mathematics, vol. 1948, Springer, 2008.
[7] M. Tepper, P. Musé, A. Almansa, On the role of contrast and regularity in perceptual boundary saliency, Journal of Mathematical Imaging and Vision (2013) 1–17.
[8] R. Grompone, J. Jakubowicz, J.-M. Morel, G. Randall, On straight line segment detection, Journal of Mathematical Imaging and Vision 32 (2008) 313–347.
[9] M. Kass, A. Witkin, D. Terzopoulos, Snakes: active contour models, International Journal of Computer Vision 1 (1988) 321–331.
[10] F. Cao, P. Musé, F. Sur, Extracting meaningful curves from images, Journal of Mathematical Imaging and Vision 22 (2005) 159–181.
[11] P. Musé, On the definition and recognition of planar shapes in digital images, Ph.D. Thesis, École Normal Supérieure de Cachan, 2004.
[12] D. Martin, C. Fowlkes, D. Tal, J. Malik, A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics, in: ICCV, 2001.

Mariano Tepper received the Ph.D. degree from the Universidad de Buenos Aires, Argentina, in 2011. He obtained the M.Sc. degree in Applied Mathematics from the ENS Cachan, France, in 2006 and the Licentiate degree in Computer Science from the Universidad de Buenos Aires, in 2007. From 2011 to 2012 he was a Postdoctoral Research Associate at the University of Minnesota, USA, and currently he is a postdoctoral associate at Duke University, USA. His research interests include image analysis and processing, computer vision and pattern recognition.

Pablo Musé received the Electrical Engineering degree from the Universidad de la República, Uruguay, in 1999, the Diplôme d'Etudes Approfondies degree in mathematics, vision and learning and the Ph.D. degree in applied mathematics from ENS de Cachan, France, in 2001 and 2004, respectively. From 2005 to 2006 he was with Cognitech, Inc., Pasadena, CA, where he worked on computer vision and image processing applications. In 2006 and 2007 he was a Postdoctoral Scholar with the Division of Geological and Planetary Sciences, California Institute of Technology, Pasadena. Since 2008 he has been with the Division of Electrical Engineering, School of Engineering, Universidad de la República, where he is currently an Associate Professor for signal processing.

Andrés Almansa received his HDR, Ph.D. and M.Sc./Engineering degrees in Applied Mathematics and Computer Science from Université Paris-Descartes, ENS Cachan (France) and Universidad de la República (Uruguay), respectively, where he is an Associate Professor since 2004. His current interests as a CNRS Research Scientist at Telecom ParisTech include image restoration and analysis, subpixel stereovision and applications to earth observation, high quality digital photography and film restoration.

Marta Mejail (Ph.D. in Computer Sciences) is a Professor at the Department of Computer Sciences of the University of Buenos Aires, Argentina. Since 1993 she is one of the Directors of Image Processing and Computer Vision Group of the Department of Computer Sciences. Her research areas are image processing, computer vision, SAR image processing, and biometrics recognition.