

A Nonparametric Approach for Histogram Segmentation

Julie Delon, Agnès Desolneux, José-Luis Lisani, and Ana Belén Petro

Abstract—In this work, we propose a method to segment a 1-D histogram without *a priori* assumptions about the underlying density function. Our approach considers a rigorous definition of an admissible segmentation, avoiding over and under segmentation problems. A fast algorithm leading to such a segmentation is proposed. The approach is tested both with synthetic and real data. An application to the segmentation of written documents is also presented. We shall see that this application requires the detection of very small histogram modes, which can be accurately detected with the proposed method.

Index Terms—Document analysis, histogram analysis, histogram segmentation, image segmentation, parameter-free method.

I. INTRODUCTION

HISTOGRAMS have been extensively used in image analysis and, more generally, in data analysis, mainly for two reasons: they provide a compact representation of large amounts of data, and it is often possible to infer global properties of the data from the behavior of their histogram. One of the features that better describes a 1-D histogram is the list of its *modes*, i.e., the intervals of values around the data is concentrated. For example, the histogram of hues or intensities of an image made of different regions shall exhibit several peaks, each one ideally corresponding to a different region in the image. In this case, a proper segmentation of the image can be obtained by determining appropriate thresholds separating the modes in the histogram. However, it is not always easy to quantify the amount of “data concentration” in an interval and, hence, to separate modes.

Among the algorithms proposed for 1-D histogram segmentation, we can distinguish between parametric and nonparametric approaches. The first ones (see [13]) assume the set of data as samples of mixtures of k random variables of given distributions, as in the Gaussian mixture models. If k is known, optimization algorithms, such as the EM algorithm [11], can estimate efficiently the parameters of these distributions. The esti-

mated density can then be easily segmented to classify the original data. The main drawback of this approach is that histograms obtained from real data cannot always be modeled as mixtures of Gaussians, for example, luminance histograms of natural images, as we shall see in the experimental section. Nonparametric approaches give up any assumption on the underlying data density. Among them, bi-level or multilevel thresholding methods, such as [1], [7], [8], and [22], divide the histogram into several segments by minimizing some energy criterion (variance, entropy, etc.).

In all cases, the number of modes in the final segmentation must be estimated. This number can be specified *a priori*, becoming a method parameter. It can also be estimated if its *a priori* distribution is hypothesized. The selection of this parameter is crucial since a wrong choice leads to an over or under segmentation of the data. Generally, *ad hoc* procedures are used to estimate the number of modes.

Other nonparametric approaches (for instance, mean shift [9]) find peaks (local maxima) of the histogram without estimating the underlying density. These methods tend to detect too many peaks in histograms coming from real noisy data. Some criterion is, therefore, needed to decide which peaks from the detected ones correspond to true modes [25]. Indeed, one of the main challenges of histogram analysis is the detection of small modes among big ones (see, for example, Fig. 3).

A different approach has been recently proposed in [14]. The authors propose to fit the simplest density function compatible with the data. Such a method is globally convincing, but the choice of the data-compatibility threshold is not formalized, only justified by experiments.

The limitations observed in the previous methods have motivated the development of a new nonparametric approach, robust to small variations in the histogram due to the limited number of samples, and local enough to detect isolated small modes. This new method is based on the automatic detection of unimodal intervals in the histogram, which allow us to segment it.

In Section II, the theoretical framework of the proposed approach is described in detail. Several tests are presented in Section III, with applications to document segmentation.

II. NEW APPROACH TO HISTOGRAM ANALYSIS

A density function f is said to be **unimodal** on some interval $[a, b]$ if f is increasing on some $[a, c]$ and decreasing on $[c, b]$. It seems appropriate to segment a histogram by looking for segments on which it is “likely” that the histogram is the realization of a unimodal law. On such intervals, we will say that the histogram is “statistically unimodal” (this expression will be precisely defined later). Obviously, such a segmentation is gener-

Manuscript received January 12, 2006; revised July 4, 2006. This work was supported in part by the Office of Naval research under Grant N00014-97-1-0839, in part by the Direction Générale des Armements (DGA), in part by the Ministère de la Recherche et de la Technologie (projet RNRT ISII), and in part by the Ministerio de Ciencia y Tecnología under Grant TIC2002-02172. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Giovanni Poggi.

J. Delon is with the CNRS, Télécom, Paris 75015, France (e-mail: julie.delon@enst.fr).

A. Desolneux is with the Laboratory of Applied Mathematics, University Paris 5, Paris 75270, France (e-mail: desolneux@math-info.univ-paris5.fr).

J.-L. Lisani and A. B. Petro are with the University of Illes Balears, Palma de Mallorca 07122, Spain (e-mail: joseluis.lisani@uib.es; anabelen.petro@uib.es).

Color version of Fig. 7 is available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIP.2006.884951

ally not unique. In particular, the segmentation defined by all the local minima of the histogram has this property. However, small variations due to the sampling procedure should clearly not be detected as modes. In order to get a “minimal” division of the histogram, these fluctuations should be neglected. We arrive at two requirements for an admissible segmentation:

- in each segment, the histogram is “statistically unimodal;”
- there is no union of several consecutive segments on which the histogram is “statistically unimodal.”

What are the right tests to decide whether a histogram is “statistically unimodal” on an interval or not? In a nonparametric setting, any unimodal density on the considered interval should be hypothesized and the compatibility between this density and the observed histogram distribution should be tested. Unfortunately, this leads to a huge number of tests and this is, therefore, impossible. There is, however, a way to address this question by testing a small number of adequate unimodal laws. In [16], this problem was solved for the case of decreasing laws. Our purpose here is to extend this method to the segmentation of any histogram into meaningful modes. We shall treat the problem in three stages in Sections II-A–C:

- step A: testing a histogram against a fixed hypothesized density;
- step B: testing a histogram against a qualitative assumption (decreasing, increasing);
- step C: segmenting a histogram and generating an estimate of its underlying density.

A. Distribution Hypothesis Testing

Consider a discrete histogram $h = (h_i)_{i=1\dots L}$, with N samples on L bins $\{1, \dots, L\}$. The number h_i is the value of h in the bin i . It follows that:

$$\sum_{i=1}^L h_i = N. \quad (1)$$

For each discrete interval $[a, b]$ of $\{1, \dots, L\}$, let $r(a, b)$ be the proportion of points in $[a, b]$

$$r(a, b) = \frac{1}{N} \left(\sum_{i=a}^b h_i \right). \quad (2)$$

Assume that an underlying discrete probability law $p = (p_i)_{i=1\dots L}$ is hypothesized for h . One would like to test the adequacy of the histogram h to this given density. For

TABLE I
NUMBER OF SEGMENTS FOUND BY THE FTC ALGORITHM
FOR 100 SAMPLES OF SIZE 2000 OF DIFFERENT LAWS

	unif.	Gauss.	mix. of 2 Gaussian laws		
			$d = 2\sigma$	$d = 3\sigma$	$d = 4\sigma$
1 segment	99	100	100	24	0
2 seg.	1	0	0	76	100
3 seg.	0	0	0	0	0

each interval $[a, b]$ of $\{1, \dots, L\}$, let $p(a, b)$ be the probability for a point to fall into the interval $[a, b]$

$$p(a, b) = \sum_{i=a}^b p_i. \quad (3)$$

Consider the hypothesis \mathcal{H}_0 that h originates from p . In other words, the N samples of the histogram h have been sampled independently on $\{1, \dots, L\}$ with law p . A simple way to accept or to reject \mathcal{H}_0 is to test for each interval $[a, b]$ the similarity between $r(a, b)$ and $p(a, b)$. Under the hypothesis \mathcal{H}_0 , the probability that $[a, b]$ contains at least $Nr(a, b)$ samples among N is given by the binomial tail $\mathcal{B}(N, Nr(a, b), p(a, b))$, where

$$\mathcal{B}(n, k, p) = \sum_{j=k}^n \binom{n}{j} p^j (1-p)^{n-j}. \quad (4)$$

In the same way, the probability that $[a, b]$ contains less than $Nr(a, b)$ samples is $\mathcal{B}(N, N(1-r(a, b)), 1-p(a, b))$. If one of these probabilities is too small, the hypothesis \mathcal{H}_0 can be rejected. We define for each interval $[a, b]$ its number of false alarms [see (5), shown at the bottom of the page].

Definition 1: An interval $[a, b]$ is said to be an ε -**meaningful rejection** of \mathcal{H}_0 if

$$\text{NFA}_p([a, b]) \leq \frac{\varepsilon}{2}. \quad (6)$$

Proposition 1: Under the hypothesis \mathcal{H}_0 , the expectation of the number of ε -meaningful rejections among all the intervals of $\{1, \dots, L\}$ is smaller than ε .

The proof of proposition 1 is obvious [12] and uses a Bonferroni argument, taking into account the number of tests $(L(L+1)/2)$ (the number of different intervals in $\{1, \dots, L\}$). This means that testing a histogram h following a law p will lead on the average to less than ε wrong rejections. It may be asked how reliable this estimate is. In [7], Grompone and Jakubowicz have shown that the expectation of ε -meaningful events could be approximated by $\varepsilon/100$. This will be confirmed in Section III (see Table I). Thus, in practice, we fix $\varepsilon = 1$, and just talk about meaningful rejections.

$$\text{NFA}_p([a, b]) = \begin{cases} \frac{L(L+1)}{2} \mathcal{B}(N, Nr(a, b), p(a, b)), & \text{if } r(a, b) \geq p(a, b) \\ \frac{L(L+1)}{2} \mathcal{B}(N, N(1-r(a, b)), 1-p(a, b)), & \text{if } r(a, b) < p(a, b). \end{cases} \quad (5)$$

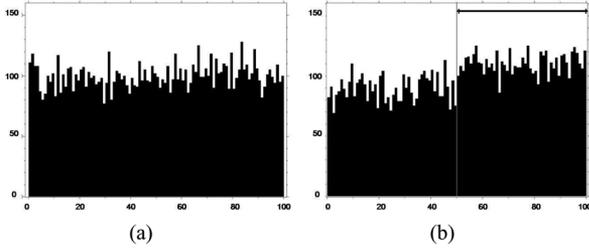


Fig. 1. Histograms of $N = 10000$ samples distributed on $L = 100$ bins, tested against the uniform law on $[1, 100]$. (a) Realization of the uniform law on $[1, 100]$. (b) Realization of a mixture of two uniform laws: $[1, 50]$ with a weight 0.45, and $[50, 100]$ with weight 0.55.

Definition 2: We say that a histogram h follows the law p on $[1, L]$ if h contains no meaningful rejection for \mathcal{H}_0 .

Fig. 1 shows two histograms which have been tested against the uniform law on $[1, 100]$. The first one is a realization of this law, and no rejection is found. The second is a mixture of two uniform laws on different intervals. In this case, several rejections of the uniform law on $[1, 100]$ are found. The rejection with the lowest NFA_p (the interval $[50, 100]$) is shown in Fig. 1(b).

B. Testing the Monotone Hypothesis

Next, we test if a histogram h follows a decreasing hypothesis on $[1, L]$ (the increasing case can be deduced by symmetry). This test will be useful later to give a suitable meaning to the expression “being statistically unimodal on an interval.” The aim of a ideal test is to examine the adaptation of h to any decreasing density on $[1, L]$. This operation is obviously impossible but can be circumvented by using an estimate of the most likely decreasing law that fits h .

Let $\mathcal{P}(L)$ be the space of discrete probability distributions on $\{1, \dots, L\}$, i.e., the vectors $r = (r_i)_{i=1, \dots, L}$ such that

$$\forall i \in \{1, 2, \dots, L\}, \quad r_i \geq 0 \text{ and } \sum_{i=1}^L r_i = 1. \quad (7)$$

Let $\mathcal{D}(L) \subset \mathcal{P}(L)$ be the space of all decreasing densities on $\{1, \dots, L\}$. If $r = (1/N)h \in \mathcal{P}(L)$ is the normalized histogram of our observations, let \bar{r} be the Grenander estimator of r . Introduced by Grenander in 1956 [16], this estimator is defined as the nonparametric maximum likelihood estimator restricted to decreasing densities on the line.

Definition 3: The distribution \bar{r} is the only distribution of $\mathcal{D}(L)$ which achieves the minimal Kullback–Leibler distance from r to $\mathcal{D}(L)$, i.e.,

$$KL(r||\bar{r}) = \min_{p \in \mathcal{D}(L)} KL(r||p) \quad (8)$$

where $\forall p \in \mathcal{D}(L)$, $KL(r||p) = \sum_{i=1}^L r_i \log(r_i/p_i)$.

Grenander shows in [16] (see also [3]) that \bar{r} is merely “the slope of the smallest concave majorant function of the empirical repartition function of r .” \bar{r} also achieves the minimal L^2 -distance from r to $\mathcal{D}(L)$. It can easily be derived from r by an algorithm called “Pool Adjacent Violators” (see [2] and [4]).

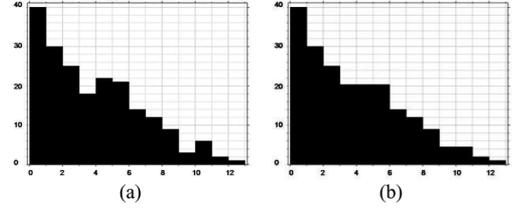


Fig. 2. (a) Original histogram. (b) Grenander estimator obtained from the “Pool Adjacent Violators” algorithm.

Pool Adjacent Violators: Consider the operator $D : \mathcal{P}(L) \rightarrow \mathcal{P}(L)$ defined by: for $r = (r_i)_{i=1, \dots, L} \in \mathcal{P}(L)$, and for each interval $[i, j]$ on which r is increasing, i.e., $r_i \leq r_{i+1} \leq \dots \leq r_j$ and $r_{i-1} > r_i$ and $r_{j+1} < r_j$, set

$$D(r)_k = \frac{r_i + \dots + r_j}{j - i + 1}, \text{ for } k \in [i, j] \\ \text{and } D(r)_k = r_k, \text{ otherwise.} \quad (9)$$

This operator D replaces each increasing part of r by a constant value (equal to the mean value on the interval). A finite number (less than the size L of r) of iterations of D yields a decreasing distribution denoted \bar{r}

$$\bar{r} = D^L(r). \quad (10)$$

An example of discrete histogram and its Grenander estimator are shown in Fig. 2.

The previous definitions of meaningful rejections can obviously be applied to this case by taking $p = \bar{r}$ in the hypothesis \mathcal{H}_0 , with \bar{r} the Grenander estimator of $r = (1/N)h$.

Definition 4: Let h be a histogram of N samples and \bar{r} the Grenander estimator of $r = (1/N)h$. An interval $[a, b]$ is said to be a **meaningful rejection for the decreasing hypothesis** if

$$NFA_{\bar{r}}([a, b]) \leq \frac{1}{2} \quad (11)$$

where $NFA_p([a, b])$ is defined for any density law p in (5).

Definition 5: We say that a histogram h follows the **decreasing hypothesis** (respectively, the increasing hypothesis) on an interval $[a, b]$ if the restriction of the histogram to $[a, b]$ [i.e., $h|_{[a, b]} = (h_a, h_{a+1}, \dots, h_b)$] contains no meaningful rejection for the decreasing (resp. increasing) hypothesis.

C. Piecewise Unimodal Segmentation of a Histogram

Definition 6: We say that a histogram h follows the **unimodal hypothesis** on the interval $[a, b]$ if there exists $c \in [a, b]$ such that h follows the increasing hypothesis on $[a, c]$ and h follows the decreasing hypothesis on $[c, b]$.

We call segmentation of h a sequence $1 = s_0 < s_1 < \dots < s_n = L$. The number n is termed length of the segmentation. Our aim is to find an “optimal” segmentation S of h , such that h follows the unimodal hypothesis on each interval $[s_i, s_{i+1}]$ of S . If S is the segmentation defined by all the local minima of h , h follows obviously the unimodal hypothesis on each of its segments, but this segmentation is not reasonable in general [see

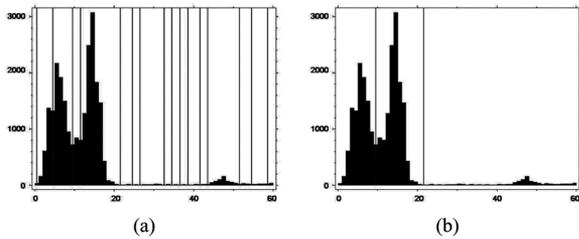


Fig. 3. (a) Initialization of the algorithm (all the local minima of the histogram). The histogram presents small oscillations, which create several local minima. (b) Final segmentation after FTC algorithm. Three modes are detected in this histogram, one is very small.

Fig. 3(a)]. A segmentation following the unimodal hypothesis on each segment is generally not unique. In order to be sure to build a minimal (in terms of number of separators) segmentation, we introduce the notion of “admissible segmentation.”

Definition 7: Let h be a histogram on $\{1, \dots, L\}$. A segmentation S of h is **admissible** if it satisfies the following properties:

- h follows the unimodal hypothesis on each interval $[s_i, s_{i+1}]$;
- there is no interval $[s_i, s_j]$ with $j > i + 1$, on which h follows the unimodal hypothesis.

The first requirement avoids under segmentations, and the second one avoids over segmentations. It is clear that such a segmentation exists. Starting from the segmentation defined by all the local minima of h , merge recursively the consecutive intervals until both properties are satisfied.

Fine to Coarse (FTC) Segmentation Algorithm

- 1) Initialize $S = \{s_0, \dots, s_n\}$ as the finest segmentation of the histogram, i.e., the list of all the local minima, plus the endpoints $s_0 = 1$ and $s_n = L$.
- 2) Repeat:
 - Choose i randomly in $[1, \text{length}(S) - 1]$. If the pair of segments on both sides of s_i can be merged into a single interval $[s_{i-1}, s_{i+1}]$ following the unimodal hypothesis, group them. Update S .
 - Stop when no more pair of successive intervals in S follows the unimodal hypothesis. $\text{length}(S)$ has decreased by one with each merging. Now $\text{length}(S) = l_0$.
- 3) For j from 3 to l_0 , repeat step 2 with the unions of j segments.

The necessity of step 3) comes from the fact that a union of j successive segments can follow the unimodal hypothesis whereas no more union of k successive segments for $k < j$ does.

The result of this algorithm on a histogram is shown in Fig. 3.

In the histogram of Fig. 3, an energy-minimizing algorithm (for example, the one presented in [1]) gives similar results if it is specified that three segments are required. The separator between the second and third modes is not located exactly at the same place, but this variation has a negligible effect on the classification, since very few points are represented in this zone of the histogram. If only two segments are required, the second and third modes are united. It is interesting to note that for the energy defined in [1], the bimodal segmentation has almost the

same energy as the three-modal segmentation. This implies that with a term penalizing the number of segments in the energy, the bimodal segmentation would certainly be chosen instead of the three-modal one. Therefore, the small mode cannot be found by this kind of method.

Fig. 6 shows the result of the FTC algorithm on a more oscillating histogram. Popular techniques of histogram analysis, such as mean shift [9], would over segment this histogram, as noted in [25], since many of the observed small oscillations would be detected as peaks.

III. EXPERIMENTS

The experimental section is organized as follows. First, some experiments on synthetic data are performed to test the ability of the method to segment mixtures of laws without *a priori* assumption. Then, some experiments on image segmentation are displayed, and the validity of modeling real data histograms by Gaussian mixtures is discussed. The section ends with experiments on document segmentation, and the robustness of the method is tested.

A. Some Results On Synthetic Data

From a given probability law, 100 histograms of $N = 2000$ samples each were generated and quantized on 50 bins. For each distribution, the number of segments found by the FTC algorithm was noted. Table I shows for different classical laws the number of distributions among the 100 leading to one, two, or three segments. The laws used here are the uniform law, a Gaussian distribution of standard deviation ten and mixtures of two Gaussian functions $(1/2)\mathcal{N}(\mu, \sigma) + (1/2)\mathcal{N}(\mu + d, \sigma)$, with $\sigma = 5$ and $d = 2\sigma, 3\sigma$ or 4σ .

For a uniform or a Gaussian law, the number of segments is almost always found to be 1. For Gaussian mixtures, the results are of course closely related to the distance d between the means of the Gaussian distributions. When $d = 2\sigma$, the FTC algorithm always finds a single segment. It begins to find two segments when $d \simeq 2.5\sigma$, and finds two segments in 99% of the cases as soon as $d \geq 3.4\sigma$. Fig. 4 shows that these results correspond to intuition. When $d = 2\sigma$, the two Gaussian functions cannot be distinguished, whereas the mixture clearly shows two modes when $d \geq 3\sigma$. These results also obviously depend on the number N of points. The larger N is, the more each histogram looks like the real mixture law, and the sooner the algorithm finds two segments. Of course, segmenting Gaussian mixtures can be made more efficiently by dedicated algorithms if we really know they are Gaussian. In practice, observed mixtures are seldom Gaussian mixtures.

B. Some Experiments on Image Segmentation

Fig. 5 displays an image that contains a uniform background and a set of small objects of different intensities. The intensity histogram shows a large peak corresponding to the background and very small groups of values corresponding to the objects in the foreground. The FTC algorithm segments the histogram into four modes. The associated regions are shown in Fig. 5, each one of them corresponding to a different object in the scene.

Modes of a gray level histogram do not necessarily correspond to semantical visual objects. In general, modes simply

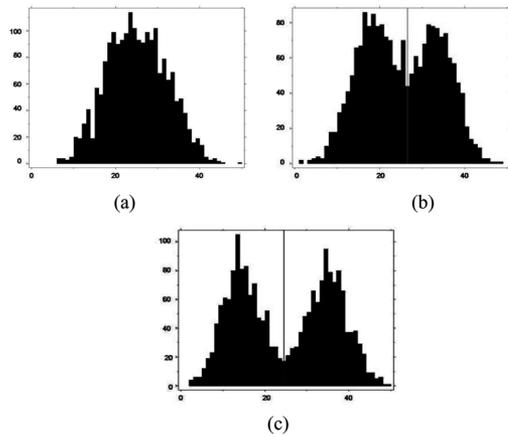


Fig. 4. Examples of Gaussian mixtures of the form $(1/2)\mathcal{N}(\mu, \sigma) + (1/2)\mathcal{N}(\mu + d, \sigma)$, where $\sigma = 5$. (a) Case $d = 2\sigma$. (b) $d = 3\sigma$. (c) $d = 4\sigma$. For $d = 2\sigma$, the FTC algorithm finds no separator. For the other mixtures, the vertical line indicates the segmentation found.

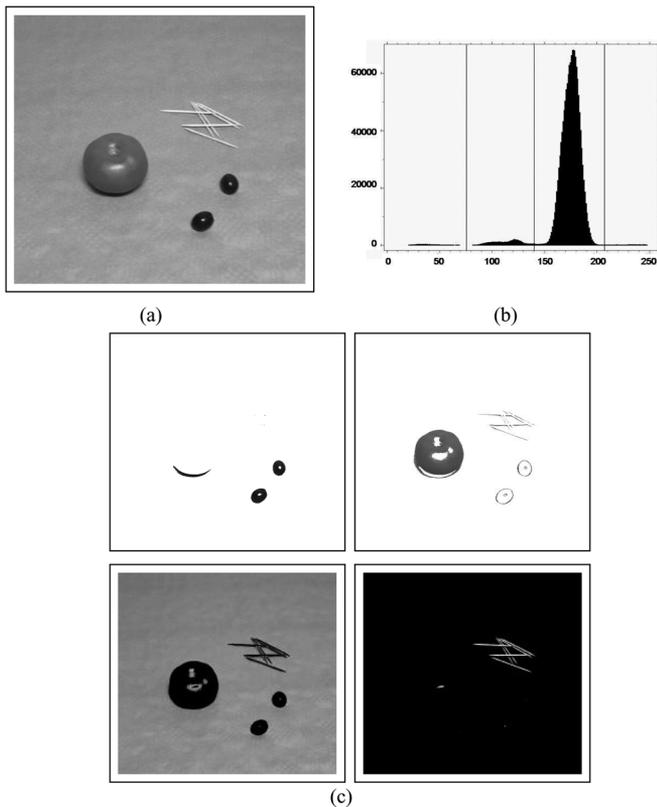


Fig. 5. (a) Original image (399×374 pixels). (b) Its intensity histogram segmented into four modes. (c) Regions of the image corresponding to the four obtained histogram modes (in decreasing level of intensity, the background is either white or black depending on the mean intensity of the represented mode).

correspond to regions with uniform intensity and segmenting the histogram boils down to quantizing the image on the so-defined levels. This is the case in Fig. 6. The histogram of “Lena” is automatically divided into seven modes, and the corresponding image quantization is shown in Fig. 6(c). Remember that no information about the spatial relations between image pixels is used to obtain the segmentation. Some authors (e.g., [23] and

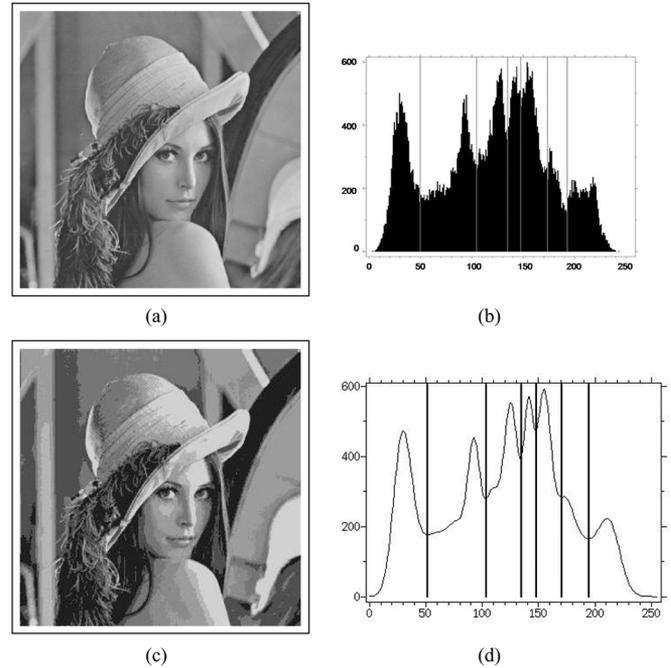


Fig. 6. (a) Image (256×256 pixels) Lena. (b) Its intensity histogram segmented into seven modes by the FTC algorithm. Observe that this histogram presents strong oscillations. In the initialization, the histogram presented 60 local minima among 256 bins. The segments have been merged until they follow Definition 7 of an admissible segmentation. (c) Image Lena quantized on the 7 levels defined by the histogram segmentation shown in (b). (d) Best mixture of 14 Gaussian laws for the histogram (b), found by an EM algorithm. The local minima of this mixture, indicated by the vertical lines, correspond almost exactly to the separators found in (b).

[8]) propose the use of such information to improve the results of histogram thresholding techniques.

As mentioned in the beginning of Section II, segmenting a histogram consists of two steps: 1) choosing a set of possible densities and 2) looking for the simplest of these densities which better adapts itself to the histogram for some statistical test. In the FTC algorithm, the densities proposed are a set of mixtures of unimodal laws, constructed from local Grenander estimators of the histogram. The test consists in looking for meaningful rejections. Another option is to use an EM algorithm to look for the best mixture of k Gaussian laws fitting the histogram. For each k , the adequacy of the mixture to the histogram is measured by a Kolmogorov–Smirnov test. The final segmentation is then defined by all the local minima of the selected mixture. This can be tested on the “Lena” histogram. The EM algorithm is initialized by a k -means algorithm. For a significance level of 5%, the first value of k leading to an accepted mixture is $k = 14$ (the p -value for this mixture is 0.053). The adaptation between this mixture and the histogram is confirmed by a Cramer von Mises test at a significance level of 5% (p -value = 0.0659). Fig. 6(d) shows this best mixture of 14 Gaussian laws and indicates its local minima. Observe that this density is constituted of seven modes that correspond exactly to the modes found previously.

With more demanding tests (a Chi-square test, or the search of meaningful rejections presented in Section II-A), all mixtures are rejected until $k = 20$ (the modes found in this case are still the seven same modes). This illustrates the discrepancy between the number of Gaussians needed to correctly represent the law

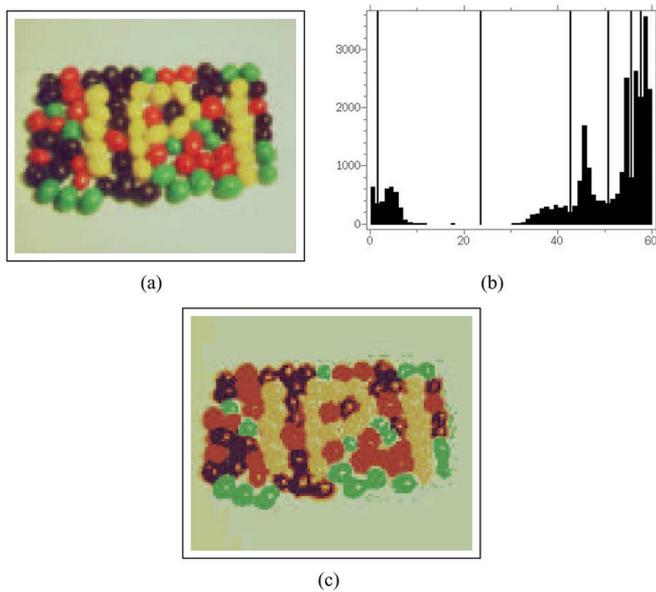


Fig. 7. (a) “Beans” image. (b) Hue histogram of the image and corresponding segmentation in six modes (remember that the hue histogram is circular). (c) Corresponding segmentation of the image.

and the actual number of modes. This discrepancy can be explained by the following observation: when a digital picture is taken, the sensors of the camera and the postprocessing that is used to store the picture into a file are nonlinear. Even if the real intensity values of a given object followed a Gaussian law, the intensity distribution of this object on the picture would not be well represented by a Gaussian function. In particular, the corresponding mode on the histogram can be highly nonsymmetric (see, e.g., Fig. 6). Such a mode needs several Gaussian laws to be well represented, whereas a unique unimodal law fits it. As a consequence, looking for a mixture of unimodal laws is more adapted in this case than looking for Gaussian mixtures.

Fig. 7 shows an example of image segmentation using the hues instead of the gray levels. Remember that hue histograms are circular. The FTC algorithm is perfectly adapted to this case.

C. Some Experiments in Document Image Analysis

Histogram thresholding is widely used as a preprocessing step for document understanding and character recognition. Its main use in this domain is to sort out the background and the characters in scanned documents. In this kind of document, the intensity histogram generally presents two different modes: one large mode that represents the background, and another one, much smaller, corresponding to the text. Many different binarization methods have been proposed (see [15] and [21]) to find the best histogram thresholds for grayscale images. Nowadays, different methods are still studied, using simple spatial features [10], texture features [20], or mathematical morphology information [6].

However, binarization methods present two drawbacks. First, when the foreground region (the text here) is too small in comparison to the background (see Fig. 8), the position of

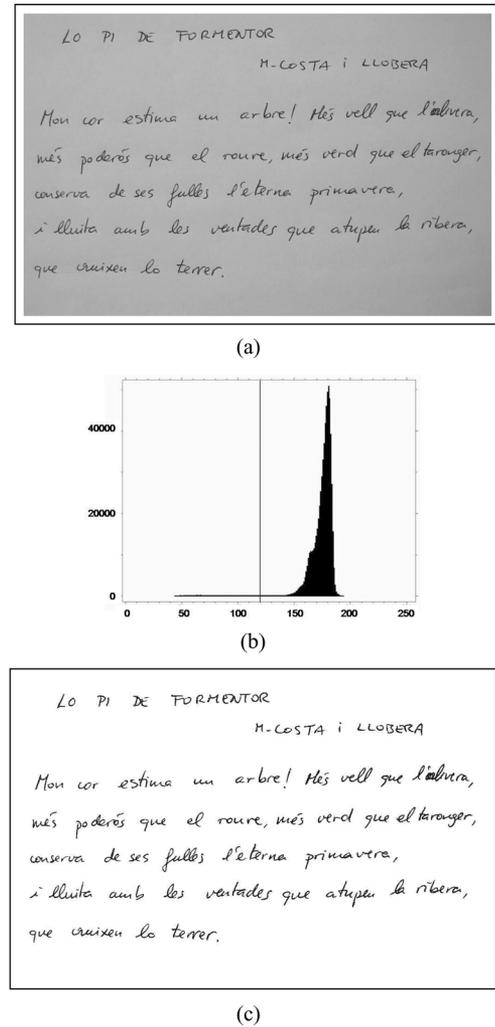


Fig. 8. (a) Original image (1010×661 pixels). (b) Intensity histogram and the threshold obtained. (c) Pixels corresponding to the left segment of the histogram.

the threshold becomes arbitrary, and the foreground may not be well detected. Second, binarization methods are not adapted to any kind of written documents (see Fig. 10). When the background pattern is complicated, or when different inks are used in the text, segmenting the histogram in only two modes is not a good solution. Finding automatically the number of modes in the histogram allows one to get more than two modes when necessary.

In simple written documents, where only one ink has been used, the segmentation found by the FTC algorithm is bimodal. This is the case of the example shown in Fig. 8. The histogram is segmented into two modes, one of them corresponding to the text characters [see Fig. 8(c)]. This example shows that the FTC algorithm is able to find very small modes when they are isolated enough. In Fig. 9, although the image presents several different gray shades, the FTC algorithm also segments the histogram into two modes Fig. 9(b), separating clearly the characters in the check from the background, as we can see in Fig. 9(c). In these experiments, it must be underlined that the size of the images can interfere in the results. In a text image, the larger the proportion of the white pixels, the more difficult it becomes to

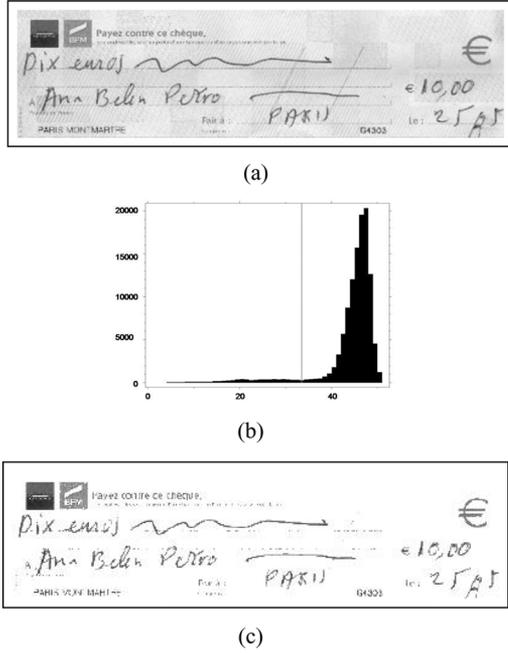


Fig. 9. (a) Original image (755 × 201 pixels). (b) Intensity histogram of the original image and corresponding segmentation. This histogram presents several local minima, but the final segmentation is bimodal. (c) Pixels corresponding to the first segment of the histogram.

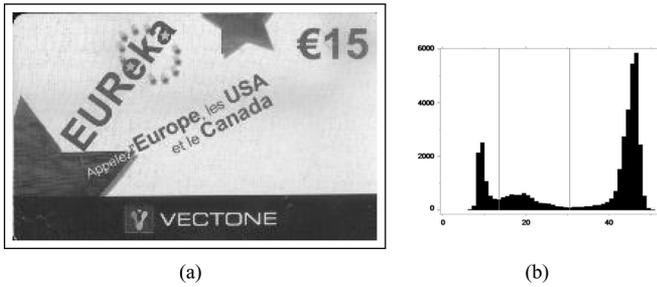


Fig. 10. (a) Original image (246 × 156 pixels). (b) Intensity histogram with the three modes obtained. The first mode on the left corresponds to the lower and darker band of the image, the middle mode corresponds to the text and the stars, and the last mode is the background one.

extract a black mode from the histogram, since it tends to become negligible in front of the white one. In these cases, it is interesting to narrow the image around the text.

In the case of the histogram of the image shown in Fig. 10, the algorithm finds three different modes, corresponding to three intensity regions in the image. The first mode represents the band in the bottom of the image, the second mode corresponds to the text and the stars of the image, and the third one is the background. A bi-level histogram thresholding method could not yield this separation (e.g., [20] and [22])

D. Sensitivity of the Method

Generally, two factors can influence the segmentation: the noise in images and the histogram quantization noise.

Theoretically, if we add a noise b to an image u , its intensity distribution becomes $h_u * h_b$, where h_u is the gray level histogram of u , and h_b the noise histogram. This results in a blur in the histogram. If the image has N pixels, and if the noise is an impulse noise, added to $p\%$ of the pixels, then $h_u * h_b =$

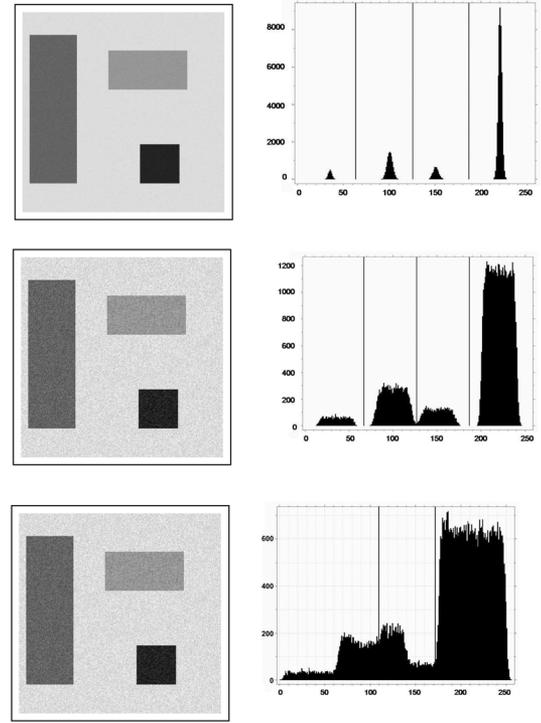


Fig. 11. Performance evaluation of the FTC algorithm in the presence of additive noise. Top, reference image (256 × 256) and its histogram. Middle and bottom, images corrupted with uniform noise (SNR = 24 dB and SNR = 18 dB, respectively), and their corresponding histograms. The segmentation results are marked on the histograms.

$(1 - p)h_u + p(\frac{N}{256})\mathbf{1}_{[0,255]}$, which is not really disturbing for the FTC algorithm (adding a uniform noise on a histogram does not change its unimodality property on a given interval). Moreover, this kind of noise can be easily removed by a median filter on the image.

In the case of additive noise (Gaussian or uniform noise), the operation smooths the shape of the histogram h_u . As a consequence, the number of modes found can decrease when the standard deviation of the noise increases too much. However, this kind of image noise can be efficiently handled by NL-means algorithms [5] before computing the histogram.

The performance of the FTC algorithm in the presence of additive noise can be evaluated as follows [19], [24]: 1) create a synthetic image (Fig. 11, top) and segment it manually; 2) add increasing quantities of uniform noise (Fig. 11, middle and bottom); and 3) segment the histograms using FTC and evaluate the probability of error

$$P(\text{error}) = \sum_{j=1}^N \sum_{i=1, i \neq j}^N P(R_i | R_j) P(R_j) \quad (12)$$

where N is the number of regions in the manually segmented image ($N = 4$ in our example), $P(R_j)$ is the proportion of pixels in the j th region and $P(R_i | R_j)$ is the proportion of pixels in the j th region assigned to the i th region by the FTC algorithm.

The results of this evaluation are shown in Table II. Since the histogram of the original synthetic image was composed of four Gaussians, the EM algorithm was also used to estimate this

TABLE II
PERFORMANCE EVALUATION. P(ERROR) FOR DIFFERENT LEVELS OF NOISE

SNR (dB)	Inf.	36	30	27	24	22	18
FTC	0	0	0	0	0	0.08	0.19
EM	0	0	0	0.08	0.44	0.72	0.71
denoise+FTC	0	0	0	0	0	0	0
denoise+EM	0	0	0	0	0	0	0.70

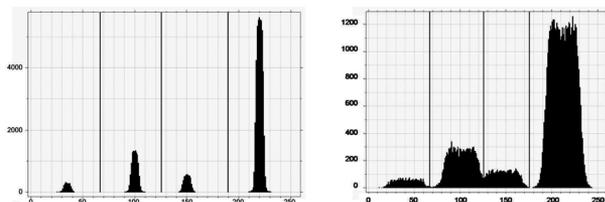


Fig. 12. Gray level histograms of images in Fig. 11 after denoising with NL-means algorithm. This algorithm improves the SNR of the images from (left histogram) 24 to 34 dB and from (right histogram) 18 to 23 dB. The segmentation results of the FTC algorithm are marked on the histograms. Compare them to the ones shown in Fig. 11 (middle and bottom rows).

mixture. Remember that when the SNR decreases, the Gaussian mixture hypothesis no longer holds and the EM algorithm gives poor results. The FTC method, however, is able to cope with this distortion down to lower SNR values.

The performance of the histogram segmentation method after denoising with NL-means algorithm has also been tested. Fig. 12 displays the histograms of some of the denoised images. The denoising method improves the SNR of the images, thus reducing the probability of error, both for the EM and the FTC algorithms, but, again, FTC gives better results than EM up to higher levels of noise (no segmentation errors are observed for SNR bigger than 17 dB; see Table II).

The real noise in histograms is the quantization noise, coming from the fact that the histograms have a finite number N of samples. If a histogram h originates from an underlying density p , the larger N is, the more h looks like p . When $N \rightarrow \infty$, a segmentation algorithm should segment the histogram at each local minima of its limit p . Consider the example of Fig. 13. An image of size 1600×1200 is subsampled several times by a factor of 2. Each intermediate image yields a histogram. These histograms can all be considered as realizations of the density given by the histogram of the original image. The smaller the number of samples, the less information we have, and the less the histogram can be segmented with certainty. Fig. 13 shows that the number of segments found by the FTC algorithm increases with N . The segmentation tends towards the separators of the deterministic histogram of the continuous underlying image.

IV. CONCLUSION

This paper presents a new approach to segment a histogram without *a priori* assumptions about the number or shape of its modes. The central idea is to test the simplest multimodal law fitting the data. The proposed adequacy test, called “meaningful rejections,” is a multiple test which presents the advantage of being simultaneously local and global. This method is more generic than looking for Gaussian mixtures and avoids overestimating the number of modes. The corresponding algorithm is able to detect very small modes when they are isolated, which

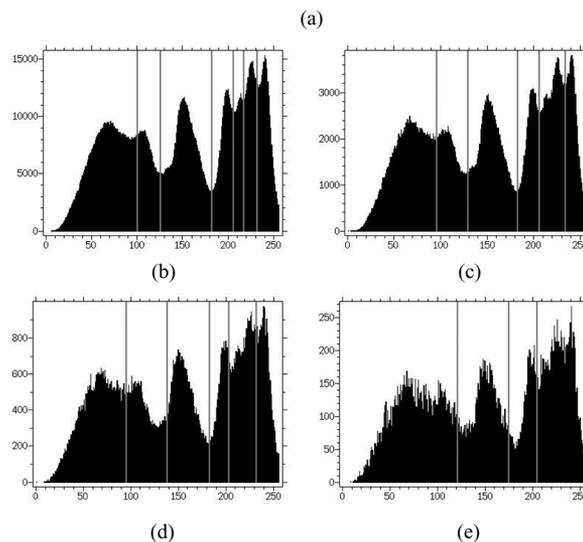


Fig. 13. Sensitivity of the method to quantization. The larger the number of samples, the more certain the segmentation. It follows that the histogram is more and more segmented when N increases. The segmentation tends towards the segmentation of the deterministic histogram of the continuous underlying image. (a) Original image, 1600×1200 wide. (b) Histogram of the original image (1920000 samples). (c) Histogram of the image subsampled by a factor of 2 (480000 samples). (d) Histogram of the image subsampled by a factor of 4 (120000 samples). (e) Histogram of the image subsampled by a factor of 8 (30000 samples).

makes it well adapted to document analysis. The statistical aspect of the approach makes it robust to quantization noise: the larger the number of samples is, the more the histogram can be considered as deterministic and the more it is segmented. Several tests on histograms computed from real or synthetic data endorse the efficiency of the method. Now, it is clear that such a method should be extended to higher dimension in order to segment color histograms. First results have been obtained by segmenting hierarchically color histograms in the HSV space. A direct adaptation of the method to any dimension is currently studied.

REFERENCES

- [1] L. Alvarez and J. Esclarín, “Image quantization using reaction-diffusion equations,” *J. SIAM*, vol. 57, no. 1, pp. 153–175, Feb. 1997.
- [2] M. Ayer, H. Brunk, G. Ewing, W. Reid, and E. Silverman, “An empirical distribution function for sampling with incomplete information,” *Ann. Math. Statist.*, vol. 26, no. 4, pp. 641–647, 1955.
- [3] R. Barlow, D. Bartholomew, J. Bremner, and H. Brunk, *Statistical Inference Under Order restrictions*. New York: Wiley, 1972.
- [4] L. Birgé, “The Grenander estimator: A nonasymptotic approach,” *Ann. Statist.*, vol. 17, no. 4, pp. 1532–1549, 1989.
- [5] A. Buades, B. Coll, and J. Morel, “A review of image denoising algorithms, with a new one,” *SIAM Multiscale Model. Simul.*, vol. 4, no. 2, pp. 490–530, 2005.

- [6] J. Caldas, L. Bandeira, J. d. Costa, and P. Pina, "Combining fuzzy clustering and morphological methods for old documents recovery," *Proc. IbPRIA*, no. 2, pp. 387–394, 2005.
- [7] C. Chang, K. Chen, J. Wang, and M. L. G. Althouse, "A relative entropy based approach to image thresholding," *Pattern Recognit.*, vol. 27, no. 9, pp. 1275–1289, 1994.
- [8] H. Cheng and Y. Sun, "A hierarchical approach to color image segmentation using homogeneity," *IEEE Trans. Image Process.*, vol. 9, no. 12, pp. 2071–2082, Dec. 2000.
- [9] D. Comaniciu and P. Meer, "Mean shift: A robust approach toward feature space analysis," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 5, pp. 603–619, May 2002.
- [10] A. Dawoud and M. Kamel, "Iterative multimodel subimage binarization for handwritten character segmentation," *IEEE Trans. Image Process.*, vol. 13, no. 9, pp. 1223–1230, Sep. 2004.
- [11] A. Dempster, N. Laird, and D. Rubin, "Maximum likelihood from incomplete data via EM algorithm," *J. Roy. Statist. Soc. B*, vol. 39, no. 1, pp. 1–38, 1977.
- [12] A. Desolneux, L. Moisan, and J.-M. Morel, "A grouping principle and four applications," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 25, no. 4, pp. 508–513, Apr. 2003.
- [13] R. Duda, P. Hart, and D. Stork, *Pattern Classification*. New York: Wiley, 2000.
- [14] G. Frederix and E. Pauwels, A Statistically Principled Approach to Histogram Segmentation CWI, 2004, Tech. Rep.
- [15] C. Glasbey, "An analysis of histogram-based thresholding algorithms," *CVGIP: Graph. Models Image Process.*, vol. 55, no. 6, pp. 532–537, 1993.
- [16] U. Grenander, *Abstract Inference*. New York: Wiley, 1981.
- [17] R. Grompone and J. Jakubowicz, *Forthcoming*. . .
- [18] J. N. Kapur, P. K. Sahoo, and A. K. C. Wong, "A new method for gray-level picture thresholding using the entropy of the histogram," *Comput. Vis., Graph., Image Process.*, vol. 29, no. 3, pp. 273–285, 1985.
- [19] Y. Lim and S. Lee, "On the color image segmentation algorithm based on the threshold and the fuzzy *c*-means technique," *Pattern Recognit.*, vol. 23, no. 9, pp. 935–952, 1990.
- [20] Y. Liu and N. Srihari, "Document image binarization based on texture features," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 19, no. 5, pp. 540–544, May 1997.
- [21] N. Otsu, "A threshold selection method from grey-level histograms," *IEEE Trans. Syst., Man, Cybern.*, vol. SMC-9, no. 1, pp. 62–66, Mar. 1979.
- [22] T. Pun, "A new method for grey-level picture thresholding using the entropy of the histogram," *Signal Process.*, vol. 2, pp. 223–237, 1980.
- [23] P. Saha and K. Udupa, "Optimum image thresholding via class uncertainty and region homogeneity," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, no. 7, pp. 689–706, Jul. 2001.
- [24] O. Tobias and R. Seara, "Image segmentation by histogram thresholding using fuzzy sets," *IEEE Trans. Image Process.*, vol. 11, no. 12, pp. 1457–1465, Dec. 2002.
- [25] H. Wang and D. Suter, "False-peaks-avoiding mean shift method for unsupervised peak-valley sliding image segmentation," in *Proc. 7th Digital Image Computing: Techniques and Applications*, 2003, pp. 581–590.



Julie Delon was born in France in 1978. From 1997 to 2001, she studied applied mathematics at the Ecole Normale Supérieure de Cachan, Cachan, France. She received the Ph.D. degree with a dissertation on image analysis from the CMLA, Cachan, in December 2004.

She is currently a Research Scientist with the CNRS, Télécom, Paris, France.



Agnès Desolneux received the Ph.D. degree with a dissertation on meaningful events in images from the Department of Applied Mathematics, Ecole Normale Supérieure de Cachan, Cachan, France, in 2000.

She is a CNRS Researcher with the Laboratory of Applied Mathematics, University Paris 5, Paris, France. She is currently working on statistical methods in image analysis.



José-Luis Lisani was born in Mallorca, Spain, in 1970. He received the Ph.D. degree in applied mathematics from the University of Paris-Dauphine, France, and the University of Illes Balears, Palma de Mallorca, Spain, respectively, in 2001.

He is currently an Assistant Professor at the University of Illes Balears, and he collaborates with the CMLA Laboratory, ENS Cachan, France. His research interests include image and video processing with applications to video surveillance, database indexing, and satellite images analysis.



Ana Belén Petro was born in Mallorca, Spain, in 1977. She received the Ph.D. degree with a thesis in color image analysis from the University of Illes Balears, Palma de Mallorca, Spain, in 2006.

She is currently an Assistant Professor at the University of Illes Balears, and she collaborates with the CMLA Laboratory, ENS Cachan, France.