Let $\mathrm{SNR}_{\mathrm{opt}}$ be the minimum detectable ratio $\sigma/\eta = b - 1$ using the power law detector with its exponent optimally chosen to match the target burstiness. In addition, let SNR_{lh} be the minimum detectable σ/η for a given burstiness λ using the power law detector with its exponent specified by likelihood theory, assuming the target burstiness is $\lambda = 1$. These signal-to-noise ratios are compared in Fig. 6, where it is shown that while there is little gain in detector performance to be had for $\lambda > 0.5$, significant improvement is possible for $\lambda < 0.5$ by choosing the detector exponent to match the target burstiness. Interestingly, the condition for and degree of improvement is fairly consistent across the whole $\mathrm{GW}(\beta, \gamma)$ noise family.

IV. Concluding Remarks

This correspondence identifies (8) as a useful relation for the study of power law detection of bursty targets in multiplicative noise, generalizing the results obtained by Fawcett and Maranda for Gaussian noise to a wide range of noise distributions including the $\mathrm{GW}(\beta, \gamma)$ family treated here. The present analytic approach using (8) usefully complements the other known approaches of numerically inverting characteristic functions and Monte Carlo estimation of detector probabilities.

Relation (8) shows generally that the more bursty the target, the greater the optimal detector exponent. Since a power law with a larger exponent more greatly amplifies strong sensor signals relative to weaker signals, this observation supports [13], wherein it is speculated that a good detector of bursty targets assumes a form that, to some degree, suppresses weak (presumably noise-only) sensor signals relative to strong signals.

Finally, it is noted that (8) readily generalizes to detection probabilities other than 0.5 and to multiplicative noises W_k with dependencies [2].

References

[1] D. Andescavage, "An analytic theory of power law detection in multiplicative noise," Honors thesis, Bucknell Univ., Lewisburg, PA, May 1997.

[2] D. Andescavage and M. R. Frey, "An analytic theory of bursty signal detection in multiplicative noise," in *Proc. Phys. Eng. Sci. Sect., Joint Stat. Meet.*, Anaheim, CA, Aug. 10–14, 1997.

[3] J. Bae and I. Song, "Rank-based detection of weak random signals in a multiplicative noise model," *Signal Process.*, vol. 63, pp. 121–131, 1997.

[4] J. Bae, I. Song, H. Morikawa, and T. Aoyama, "Nonparametric detection of known signals based on ranks in multiplicative noise," *Signal Process.*, vol. 60, pp. 255–261, 1997.

[5] R. S. Blum and S. A. Kassam, "Approximate analysis of the convergence of relative efficiency to ARE for known signal detection," *IEEE Trans. Inform. Theory*, vol. 37, pp. 199–206, Jan. 1991.

[6] R. A. Brooks and A. C. Bovik, "Robust techniques for edge detection in multiplicative Weibull image noise," *Pattern Recogn.*, vol. 23, pp. 1047–1057, 1990.

[7] P. J. Crepeau, "Uncoded and coded performance of MFSK and DPSK in Nakagami fading channels," *IEEE Trans. Commun.*, vol. 40, pp. 487–493, Mar. 1992.

[8] N. R. Draper and D. E. Tierney, "Exact formulas for additional terms in some important series expansions," *Commun. Statist.*, vol. 1, pp. 495–524, 1973.

[9] J. Fawcett and B. Maranda, "The optimal power law for the detection of a Gaussian burst in a background of Gaussian noise," *IEEE Trans. Inform. Theory*, vol. 37, pp. 209–214, Jan. 1991.

[10] M. R. Frey, "Power law detection of Poisson bursts," in *Proc. Twenty-Fifth Annu. Conf. Inform. Sci. Syst.*, Johns Hopkins Univ., Baltimore, MD, Mar. 20–22, 1991.

[11] W. A. Gardner, "A unifying view of second-order measures of quality for signal classification," *IEEE Trans. Commun.*, vol. COMM-28, 807–816, June 1980.

[12] L. M. Garth and Y. Bresler, "On the use of asymptotics in detection and estimation," *IEEE Trans. Signal Processing*, vol. 44, pp. 1304–1307, May 1996.

[13] K. Gerlach, M. Steiner, and F. C. Lin, "Detection of a spatially distributed target in white noise," *IEEE Signal Processing Lett.*, vol. 4, pp. 198–200, July 1997.

[14] Y.-S. Lee and T.-K. Lin, "High order Cornish-Fisher expansion," *Appl. Statist.*, vol. 41, pp. 233–240, 1992.

[15] E. W. Stacy, "A generalization of the gamma distribution," *Ann. Math. Stat.*, vol. 33, pp. 1187–1192, Sept. 1962.

A Classification-Driven Partially Occluded Object Segmentation (CPOOS) Method with Application to Chromosome Analysis

Boaz Lerner, Hugo Guterman, and Its'hak Dinstein

***Abstract*—Classification of segment images created by connecting points of high concavity along curvatures is used to resolve partial occlusion in images. Modeling of shape or curvature is not necessary nor is the traditional excessive use of heuristics. Applied to human cell images, 82.6% of the analyzed clusters of chromosomes are correctly separated, rising to 90.5% following rejection of 8.7% of the images.**

***Index Terms*— Chromosome analysis, image classification, image segmentation, neural networks, partial occlusion.**

I. Introduction

To generate a description of an image, it is necessary to segment the image into regions of interest or objects, each having a high level of uniformity in some parameter such as brightness, color, or texture. However, objects in images in real-world applications do very often partially occlude each other; hence, their segmentation is not trivial and requires the application of a dedicated procedure. A failure to recognize partially occluded objects or partially occluded object boundaries using global features (geometrical features, moments, Fourier descriptors, etc.) has led researchers to seek new approaches. These were usually based on matching model and scene features (landmarks) such as holes, points, line or curve segments, or on matching the relative positions of critical points in the model and the scene (see e.g., [2]). These methods, however, are sometimes either sensitive to scale variation, computationally expensive, or depend highly on the curvature and, hence, are sensitive to noise.

In this work, we use a new approach called a classification-driven partially occluded object segmentation (CPOOS) method to resolve partial occlusion in images. Classification-driven segmentation methods have been applied before to separate two touching,

Manuscript received July 21, 1997; revised April 6, 1998. This work was supported in part by the Paul Ivanier Center for Robotics and Production Managment, Ben-Gurion University, Beer-Sheva, Israel. The associate editor coordinating the review of this paper and approving it for publication was Prof. Yu-Hen Hu.

B. Lerner is with the University of Cambridge Computer Laboratory, New Museums Site, Cambridge, U.K. (e-mail: boaz.lerner@cl.cam.ac.uk).

H. Guterman and I. Dinstein are with the Department Electrical and Computer Engineering, Ben-Gurion University of the Negev, Beer-Sheva, Israel.

Publisher Item Identifier S 1053-587X(98)07090-1.

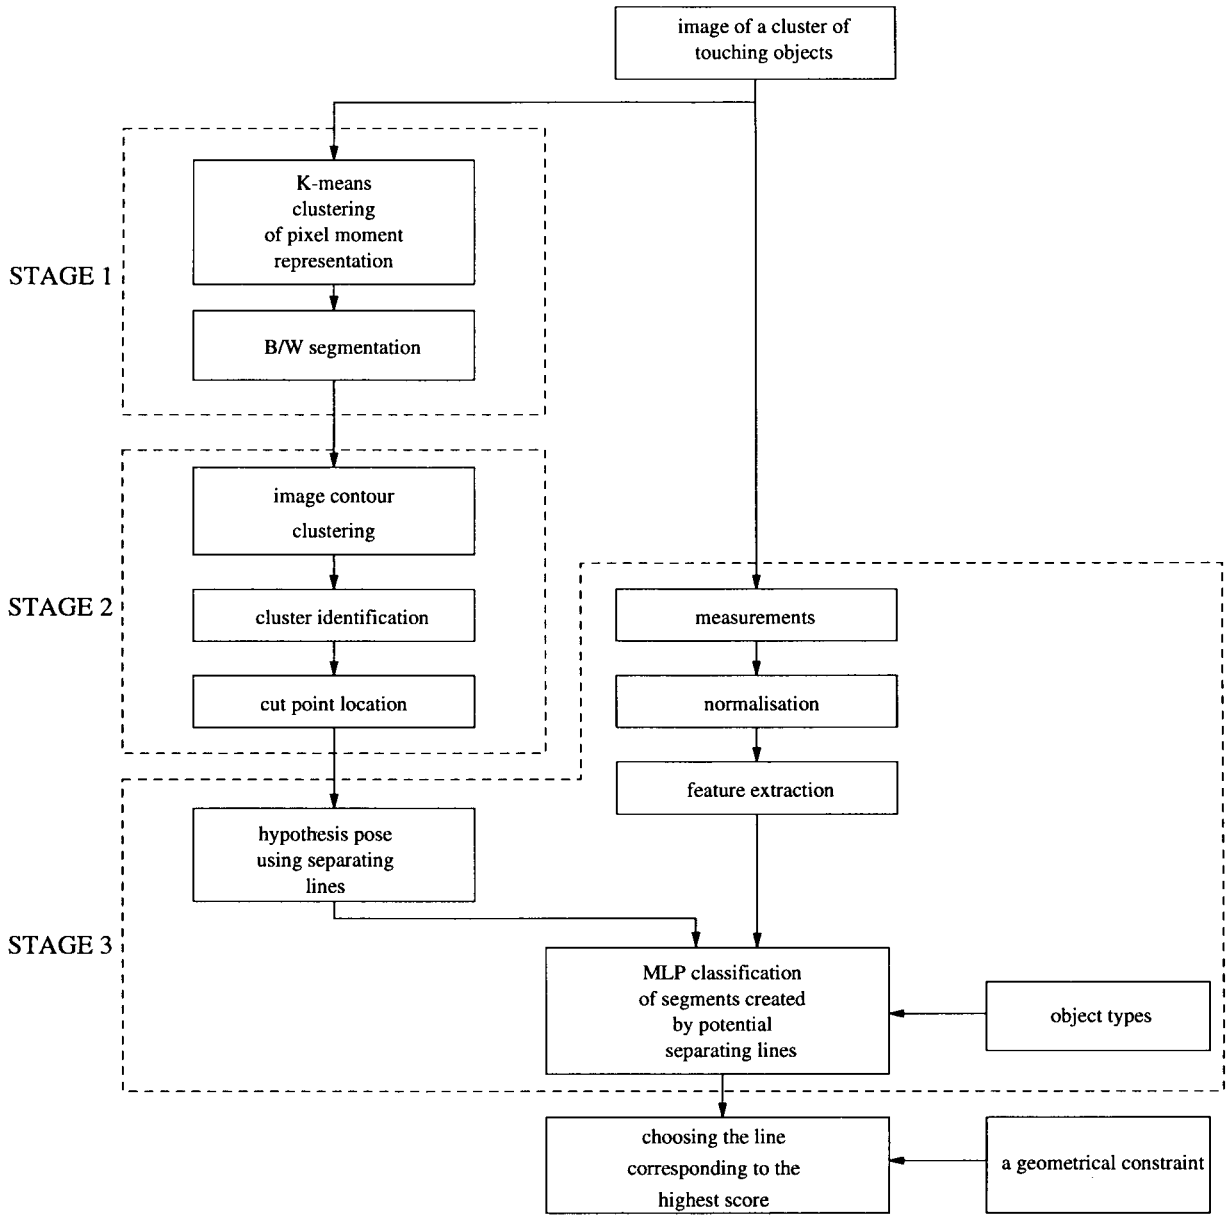

Fig. 1. Flow chart of the CPOOS method.

printed characters in different optical character recognition (OCR) applications (see e.g., [4]). However, methods for separating printed characters are critically dependent on a post-processing stage, usually using a lexicon or a graph algorithm, to complete the segmentation. The CPOOS method output, however, is robust enough to provide the ultimate classification without any post-processing. Moreover, to the best of our knowledge, the CPOOS method is the first classification-driven segmentation method for problems in applications other than character recognition. Section II of the correspondence describes chromosome analysis, which is used here to evaluate the CPOOS method. Sections III and IV present the CPOOS method and the results of applying it to chromosome analysis, respectively, whereas Section V summarizes the benefits of the method for chromosome and generic image analyses.

II. Chromosome Analysis as an Application for Partially Occluded Object Segmentation

Chromosome images are best discriminated by a combination of specific geometrical and intensity-based features [5]. The geometrical features are the length and centromeric index (the ratio of the short arm length to the total length) of a chromosome. The intensity-based features, which are called density profile (d.p.) features, are averaged intensities along sections perpendicular to the chromosome medial axis. Considering our previous experience [5], we represent chromosomes in this work by those two geometrical features and 64 d.p. features.

The inability to automatically segment clusters of touching and overlapping chromosomes is the main reason that prevents human chromosome analysis from being a *completely* automatic procedure [5], [6]. Current methods to segment partially occluded chromosome images are based on finding paths between chromosomes but yield a relatively poor success rate (50–70%) [6]. Liang [6] analyzes concavities along the chromosome shape and performs a heuristic search for a minimum density path between touching chromosomes with a success rate of 90–95%. However, the method involves threshold selection that dictates an immediate rethresholding, and it uses a large number of empirical thresholds that might suit images taken in specific conditions but not necessarily in others. The

TABLE I
COMPARISON OF THREE METHODS FOR SEGMENTING TOUCHING CHROMOSOMES

	Liang's method [6]	Agam & Dinstein's method [1]	The CPOOS method
based on	intensity image	B/W image	B/W and intensity images
curvature representation	*k-curvature*	*k-curvature*	*k-curvature*
constraints on cut points	on the first cut points (strict)	none	none
use of pre & post-processing	large	(very) large	small
use of heuristics	large	moderate	small
selection of a separating line by	finding the first cut point and a heuristic search for a minimum density path	hypothesis verification using the convex-hull of a segment	(MLP) classification of pairs of separated segments
accumulated experience	large (~700 images)	moderate (tens of images)	moderate (tens of images)

procedure is time consuming, and its success depends critically on the determination of the first cut point, as this is the origin of the split line, and on the existence of a minimum density path between chromosomes. In addition, it is expected [6] that when the method is applied to banded chromosomes (which are those usually used in routine chromosome analysis), the cytogeneticist will have to modify the segmentation results. In Agam and Dinstein [1], a separating line between touching chromosomes is chosen as the line that best verifies one of a series of hypotheses. Any two concave points along the cluster contour that obey some geometrical constraints pose a hypothesis to be verified by evaluating the fit of the separated segments to some prototype shapes. Although the method is applied to clusters of more than two chromosomes, it is a complicated method composed of a relatively large number of stages, almost every one of which requires pre and/or postprocessing.

To summarize, rejecting images containing clusters of touching or partially occluded chromosomes would require the analysis of many more images and will make chromosome analysis more time consuming, especially when searching for infrequent abnormalities. Current separation methods are complex, heuristic-based methods; hence, a *simple* yet reliable method for completely automatic separation of partially occluded chromosome images from any source is needed.

III. THE CPOOS METHOD

The CPOOS method is a classification-driven segmentation method that is applied here to chromosome analysis. First, and before applying the method, isolated objects in the image are classified into their classes. Thereafter, the CPOOS method is applied to those clusters that could not be separated and therefore classified. The first stage of the CPOOS method (Fig. 1) is K-means clustering[1] of an algebraic moment representation of the image pixels to create a binary image. Algebraic moments provide a more comprehensive characterization than the pixel gray level does since they depend on the neighborhood of the pixel as well. Since at this stage we are looking for a rapid implementation of a two-level (B/W) segmentation, only two moments (m_{01} and m_{02}), which are calculated in a 3×3 neighborhood around the central pixel, are found to provide a robust B/W segmentation. Moreover, the common problem of threshold selection is avoided using this approach.

[1]The term "cluster" is usually used to describe both the aggregation of partially occluded objects [6] and the operation of grouping similar patterns (pixels) into groups (clustering) [3].

In the second stage, clusters of touching objects are identified by their size and their failure to be classified as one of the object classes. The r most concave points along the cluster curvature are selected and suggested as potential cut points to draw possible separating lines between touching objects. Points among these r potential cut points, which are closer than a distance s from more acutely concave points, are eliminated. The values used for r and s are 8 and 10, respectively, and they remain unchanged throughout the experiments.

In the final stage, lines connecting pairs of cut points are hypothesized as separating lines. The two segments created by such a line are suspected of being the touching objects, and a classifier is employed to verify the corresponding hypothesis. For chromosome analysis, the chromosomal features of Section II and a rectangular fit factor (the ratio of the object's area to the area of its minimum enclosing rectangle) are measured to yield a 67-dimensional feature vector. The d.p. features are projected onto their principal axes, which enables the use of approximately 90% of the variance in the d.p. features with only 1/16 of the number of original features and with almost no performance degradation [5]. An MLP trained by the backpropagation (BP) learning algorithm [3] classifies the two reduced vectors (four eigenfeatures plus the three geometrical features) to provide hypothesis verification. The MLP enables a simple yet reliable method by which to select the "correct" separation line. The classifier is trained beforehand based on *a priori* knowledge of the identity of the two types of chromosomes that compose the cluster. This knowledge is gained following the completion of the classification of the isolated chromosomes and the application of a simple elimination criterion.

The classification is based on three classes. The first two (the "wholes") are the two expected types of chromosomes that compose the cluster [see Fig. 4(b) for a demonstration] and the third (the "brokens") is of those images created by any arbitrary "wrong" separating line [Fig. 4(a), (c), (d)]. The expected output vectors for the two segments are (100) and (010) for a "correct" separating line or (001) and (001) for a "wrong" one. Since the presentation of the segments to the network is arbitrary, both orders of output vectors [(100) followed by (010) and *vice versa*] are checked. During the test, the first output value of the first output vector and the second output value of the second output vector are averaged to yield a score, which is an estimation of the average maximum *a posteriori* probability [3] of the two input vectors. The highest score assigned by the classifier indicates the "correct" separating line and

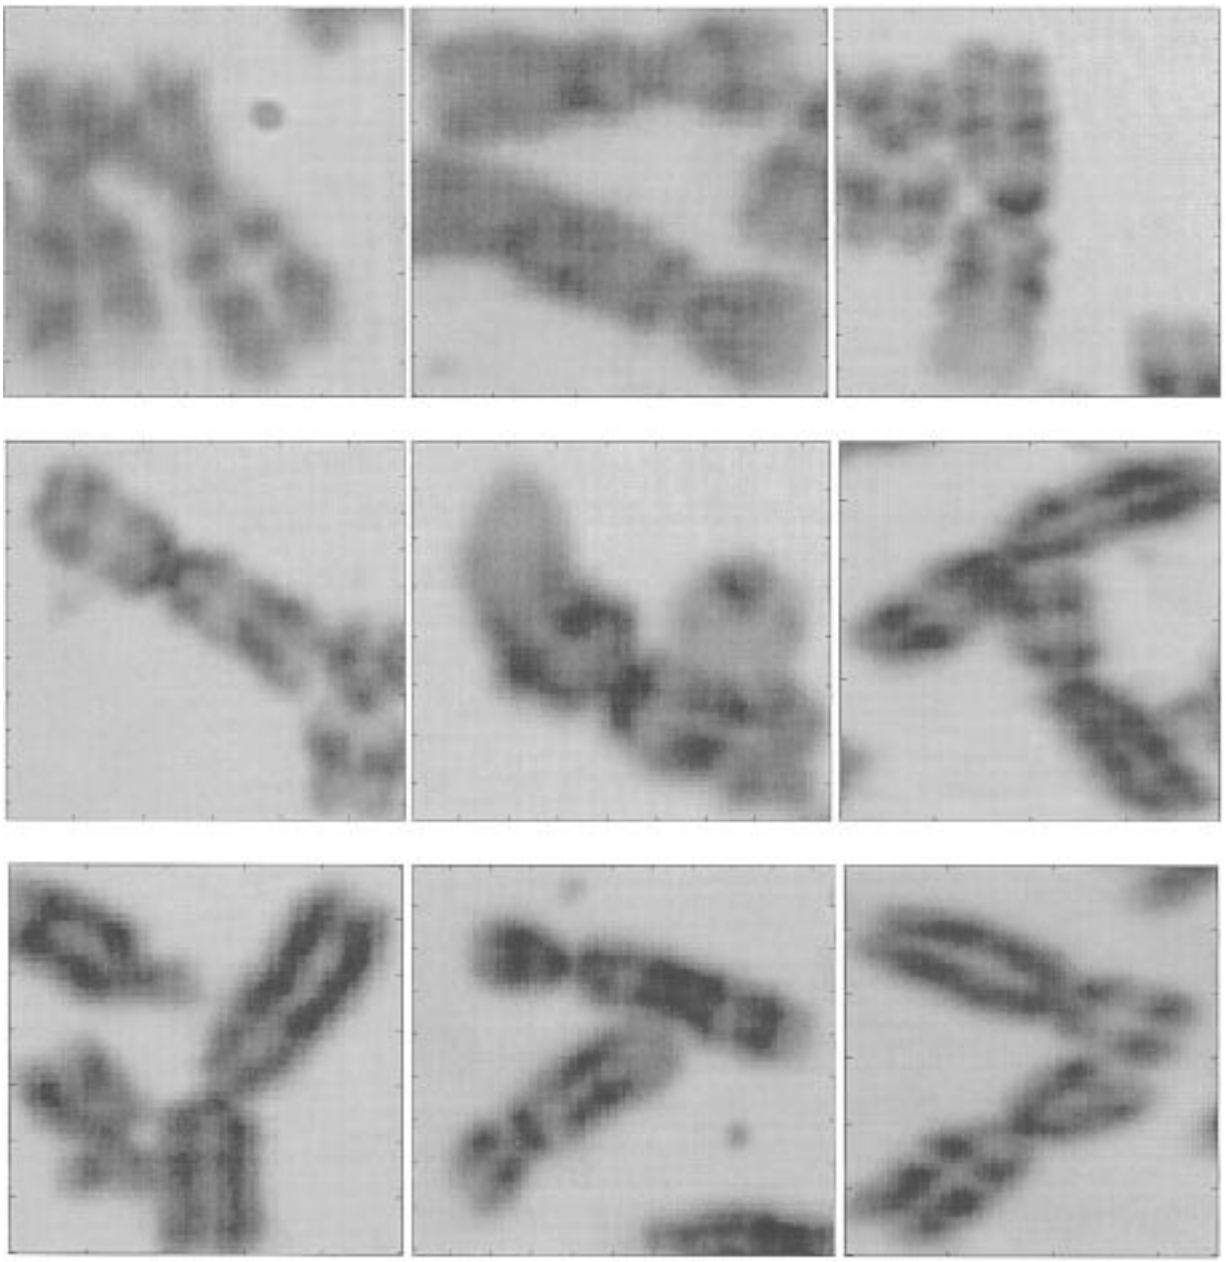

Fig. 2. Examples of subimages (used once for training and once for testing) of combinations of different touching or partially occluded chromosomes.

thereby verifies the corresponding hypothesis. Being a classification-driven segmentation procedure, the CPOOS method is dependent on achieving accurate chromosome classification capability, which has been recently demonstrated [5]. In addition, a "geometrical constraint" is obtained during the training using the maximum and minimum values of the length and centromeric index. A separating line in a test image that creates two segments such that at least one of them yields a length or a centromeric index that is either larger than one of the above maxima or smaller than one of the above minima is rejected and is not checked by the classifier.

Compared with other partially occluded object segmentation methods, the CPOOS is a classification-driven segmentation method that combines a simple preliminary segmentation stage with a recognition phase. It uses both the object curvature and the intensity image. Table I summarizes few of the characteristics of the CPOOS method compared with those of the segmentation methods of Liang [6] and Agam and Dinstein [1] in separating touching chromosomes.

In using classification-driven segmentation methods in OCR applications [4] to separate any two touching characters, we require as many classifier outputs as the number of digits or characters. Furthermore, huge training sets are required to represent all the characters and all the combinations of two touching characters. These classifiers are therefore complex and difficult to train. However, by incorporating, in the CPOOS method, *a priori* information regarding the identity of the objects that create the occlusion, the number of classifier outputs is set to be relatively low (3), which can therefore allow a lower complexity network and shorter training sessions.

IV. Segmentation Results in Chromosome Analysis

Only subimages of normal cell images that include clusters of touching or partially occluded chromosomes are selected and analyzed here (Fig. 2). Training and test images are gathered for different combinations of two touching chromosomes. Since touching among

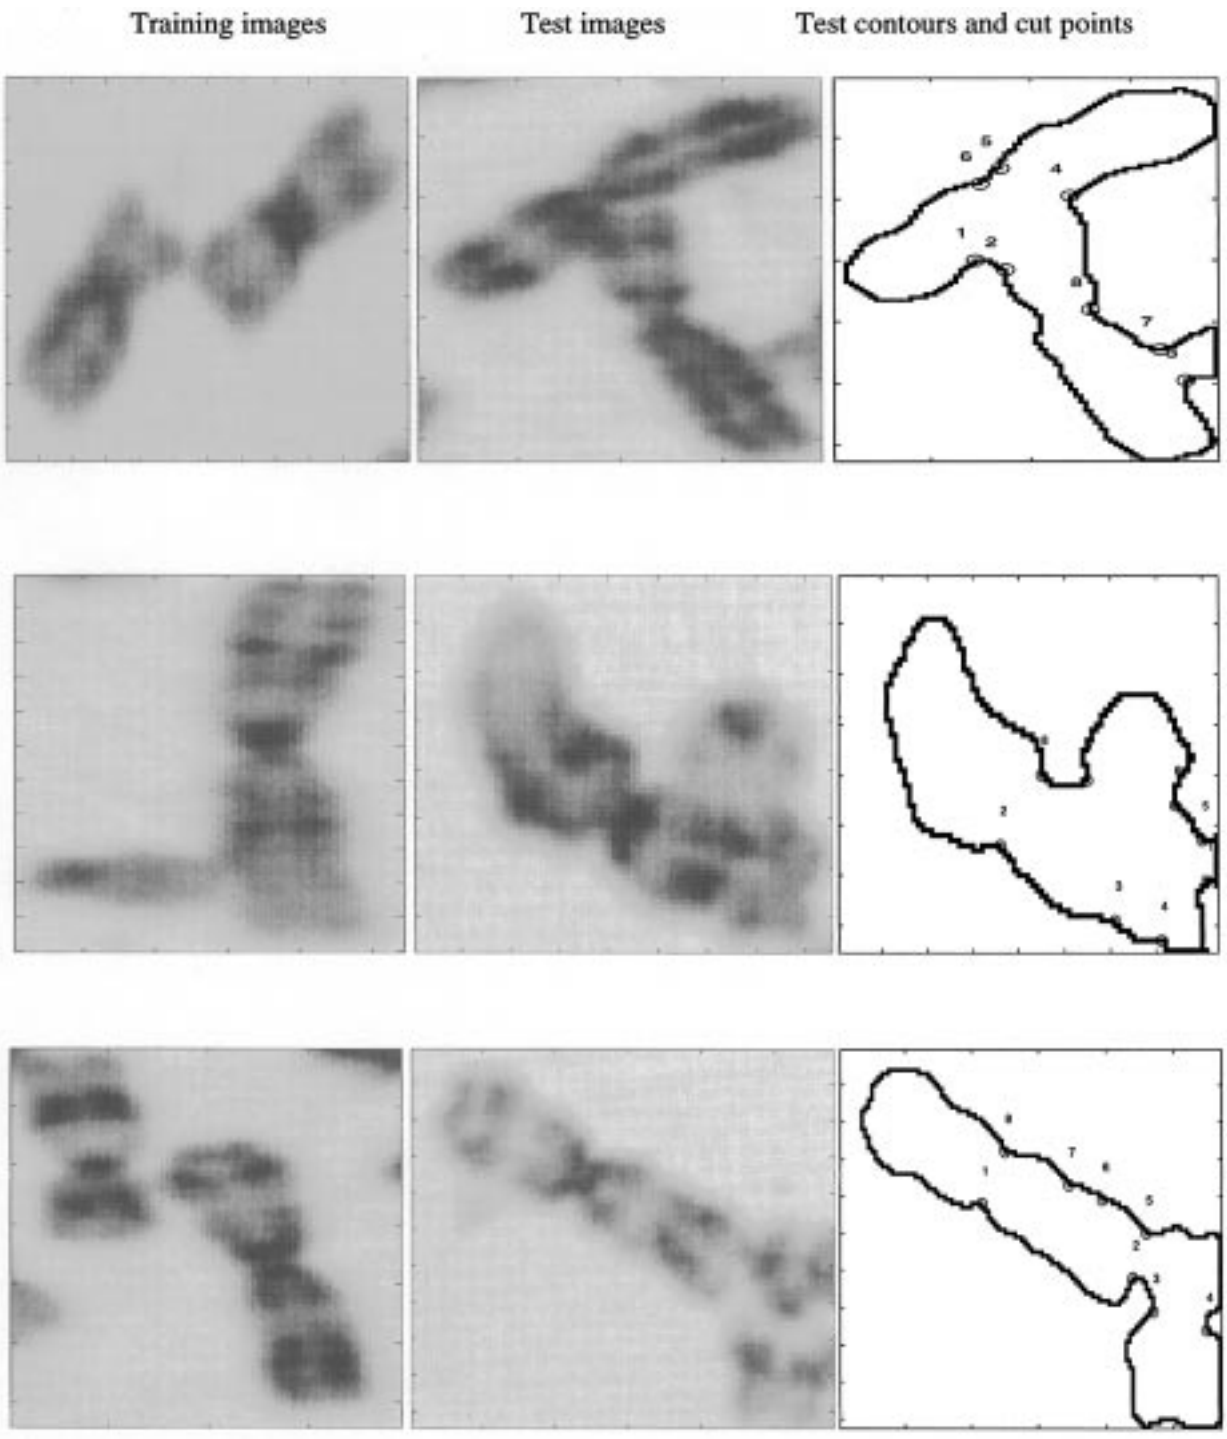

Fig. 3. Examples of clusters of touching and partially occluded chromosomes used for training and test with the contours and cut points (circles) of the test images.

chromosomes is arbitrary and there are many as $p(p-1)/2$ such combinations ($p = 24$), it is possible to collect only a few images with the same combination of touching chromosomes. Moreover, to our best knowledge, there is no standard data base of cell images; hence, we use our self-extracted base, which allows us to use only one pair of images for each combination in most of the experiments. Therefore, training and testing in each pair is alternated. For each combination and training image, we create segment images of arbitrary cluster separation using all the cluster's potential separating lines. These images define the data set of the third class (the "brokens"). The images of the first two classes (the "wholes") are collected randomly from a manually segmented chromosome data base [5]. The two sets ("brokens" and "wholes") are combined to one training image set of around 300 images ($\sim$100 from each class). The d.p. feature vectors—one from each image of the set—are measured, normalized, and projected into the first four principal axes to define, together with the geometrical features, the training set.

To cover all the possible combinations of two touching chromosomes, $p(p-1)/2$ networks are required. However, each one of these networks is of a modest configuration (7:2:3) (two hiddens)

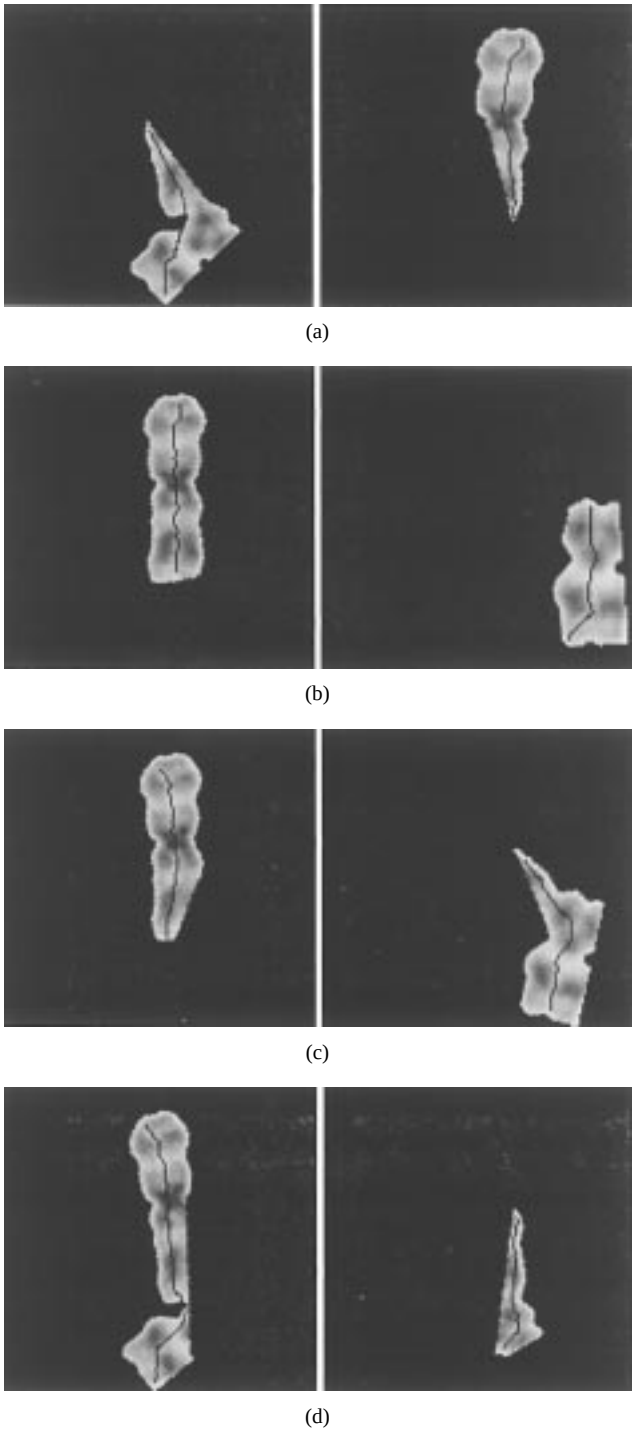

(a) (b) (c) (d)

Fig. 4. Separation of the cluster of the last test image of Fig. 3 by lines connecting points (a) 1 and 5, (b) 2 and 5 (the "correct" line), (c) 2 and 7, and (d) 4 and 8 (the black lines along the segment images are the medial axes along which the profiles are measured).

and only classifies three classes; therefore, training is very short (100 epochs). Moreover, since training is done beforehand, the verification procedure (test) is very rapid. Alternatively, since the identity of the objects that compose the cluster is known, we could, instead of training $p(p-1)/2$ networks, train only the network required to separate the specific cluster.

Fig. 3 shows three examples, with different levels of complexity, of training and test images of clusters of touching or partially occluded chromosomes along with the cluster contours and cut points of the test

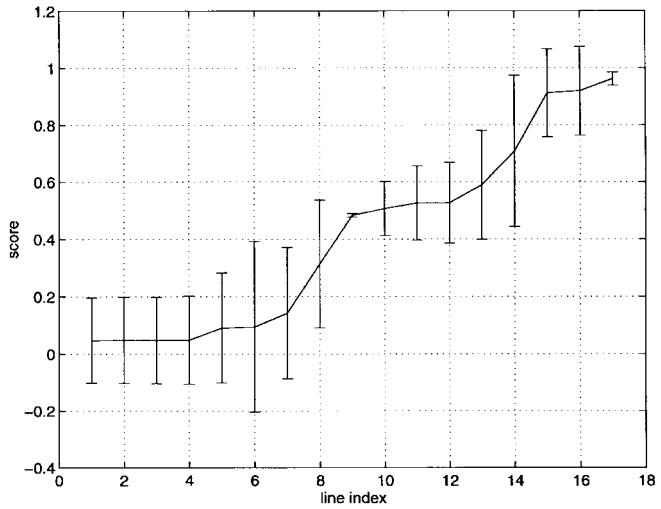

Fig. 5. Scores of different separating lines applied to the last test image of Fig. 3 (averages and standard deviations for ten experiments). Lines are arranged by their scores from the poorest to the best.

images. A few examples of pairs of segments of the last test image cut by different separating lines are given in Fig. 4. Only those segments of a potential separating line that satisfy the "geometrical constraint" are classified, and that line obtains a score by the classifier. Lines that create two segments composed of the expected chromosomes will have high scores (close to one). Those lines that create arbitrary ("wrong") separation and hence segments of nonchromosomes will have low scores (approximately zero). Finally, lines responsible for segments where only one of them can be classified correctly (with some probability) as a chromosome receive medium scores (around 0.5).

Fig. 5 shows scores obtained by different lines applied to the last test image of Fig. 3. All the results reported here are averages of ten experiments with random training sets derived from each training image and ten randomly initialized classifiers; therefore, the results are averages over 100 simulations. The lines in Fig. 5 are arranged by their scores so the "correct" line (2–5 in this case), which is line number 17, receives an average score of 0.963 with a standard deviation of 0.02.

Following the limitations described before, we are able to analyze 46 images, each one of which is used once for training and once for testing. These images represent both the different types of chromosomes and a variety of situations of touching and partial occlusion among chromosomes. In four of the images, the CPOOS method fails, and a "wrong" line is selected. This mainly occurs when a "correct" line "cuts" a segment excessively so that the segment is not consistent with the "geometrical constraint" and, hence, not classified, and the corresponding line is disqualified for separation. In addition, since each image is tested based on a limited number of training patterns, a rejection option has been added for cases where all the lines receive low average scores (below 0.4). Different definitions of this option are responsible for different points on the error-rejection curve of the problem.

Fig. 6 shows, for the 38 correctly classified images, a histogram of the number of times (out of ten experiments) the "correct" line receives the highest score. For example, in nine of the images, the "correct" line is selected in six out of the ten experiments held for each one of these nine images. In 84% of the images, the "correct" line achieves the highest score in at least four out of ten experiments. In Fig. 7, the results of the CPOOS method are compared with those of an expert for the 46 images. Cases where the

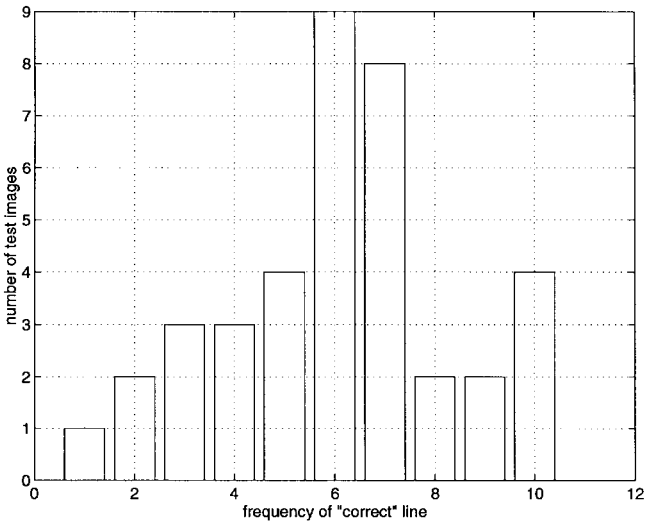

Fig. 6. Histogram for the number of times the "correct" separating line is selected (out of ten experiments) for the 38 correctly segmented test images.

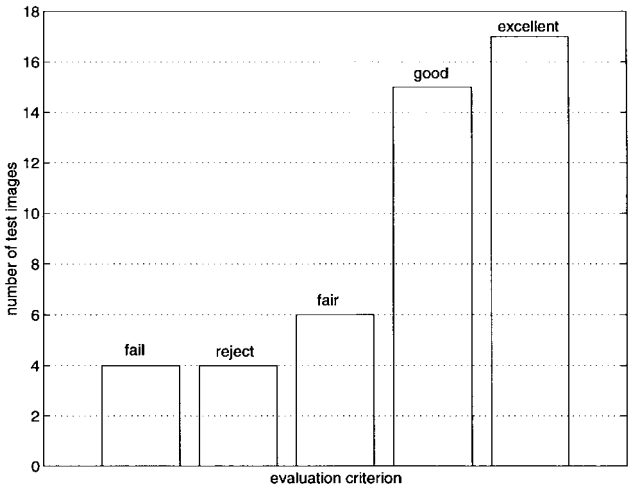

Fig. 7. Histogram for the frequency of evaluation criteria.

CPOOS method chooses the "wrong" separating line are evaluated as "fail" and cases of rejection as "reject." Cases where the chosen line is considered as the correct one by the expert and the score obtained by the CPOOS method is between 0.4–0.7 are evaluated as "fair." Those cases, as above but with scores between 0.7–0.9, are assessed to be "good," and those with scores higher than 0.9 are assesed to be "excellent." Avoiding the rejection option, the CPOOS method achieves a probability of correct segmentation of 82.6% ("fair" + "good" + "excellent") and a probability of error segmentation of 8.7%. However, taking advantage of this option and rejecting, according to the previous criterion, 8.7% of the images, the probability of correct segmentation reaches 90.5%.

V. Conclusion

To resolve partial occlusion in images, the CPOOS method incorporates object classification during the segmentation, along with a combination of object and boundary information, to provide

a) a flexible segmentation method that uses a minimum of heuristics and, therefore, can be adapted easily to a new environment and data (only retraining of the classifier is needed);
b) the possibility for a designer to choose his/her own classifier according to their experience with the application and the data;
c) a *simple* and *practical* solution to automatic human cell image segmentation.

When applied to a relatively small database of human cell images, the CPOOS method achieves over 90% probability of correct segmentation when allowed to reject 8.7% of the images. Enlarging the data base to enable training using more than one image for each test image will probably improve accuracy significantly while reducing the proportion of data rejected. Thus, efforts should be made to expand our database or to combine it with other small databases to establish an extended, publicly available standard database for such research and for an objective comparison of separation methods.

A further step in the research is to modify the method to handle more than one cluster in an image and clusters of more than two objects through hierarchical application of the method combined with suitable context-based heuristics. Consequently, and for chromosome inspection, the method may lead to a *completely* automatic human chromosome analysis.

Finally, we believe that, as is demonstrated for chromosome analysis, the CPOOS method leads toward a *generic* classification-driven segmentation method of partially occluded objects in images.

Acknowledgment

The authors thank the anonymous referees whose comments helped improve this manuscript.

References

[1] G. Agam and I. Dinstein, "Geometric separation of partially overlapping nonrigid objects applied to automatic chromosome classification," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 19, pp. 1212–1222, 1997.
[2] N. Ansari and E. J. Delp, "Partial shape recognition: A landmark-based approach," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 12, pp. 470–483, 1990.
[3] C. M. Bishop, *Neural Networks for Pattern Recognition.* Oxford, U.K.: Oxford Univ. Press, 1995.
[4] S. Kahan, T. Pavlidis, and H. S. Baird, "On the recognition of printed characters of any font and size," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-9, pp. 274–288, 1987.
[5] B. Lerner, "Toward a completely automatic neural network based human chromosome analysis," *IEEE Trans. Syst. Man Cybern. (Special Issue on Artificial Neural Networks)*, vol. 28, pp. 544–552, 1998.
[6] J. Liang, "Fully automatic chromosome segmentation," *Cytometry*, vol. 17, pp. 196–208, 1994.