Segmentation of Mosaicism in Cervicographic Images using Support Vector Machines

Zhiyun Xue [a], L. Rodney Long [a], Sameer Antani [a], Jose Jeronimo [b], George R. Thoma [a]
[a]National Library of Medicine, NIH, Bethesda, MD
[b]Program for Appropriate Technology in Healthcare (PATH), Seattle, WA

ABSTRACT

The National Library of Medicine (NLM), in collaboration with the National Cancer Institute (NCI), is creating a large digital repository of cervicographic images for the study of uterine cervix cancer prevention. One of the research goals is to automatically detect diagnostic bio-markers in these images. Reliable bio-marker segmentation in large biomedical image collections is a challenging task due to the large variation in image appearance. Methods described in this paper focus on segmenting mosaicism, which is an important vascular feature used to visually assess the degree of cervical intraepithelial neoplasia. The proposed approach uses support vector machines (SVM) trained on a ground truth dataset annotated by medical experts (which circumvents the need for vascular structure extraction). We have evaluated the performance of the proposed algorithm and experimentally demonstrated its feasibility.

1. INTRODUCTION

Cervical cancer has high incidence rate worldwide, especially in under-developed countries. Invasive cervical cancers develop slowly and are preceded by a long stage of precursor conditions referred to as cervical intraepithelial neoplasia (CIN). CIN is detectable through screening techniques and treatable if detected early. Among screening methods used, *cervicography* is a cost-effective alternative to colposcopy in resource-limited areas. Similar to colposcopy, cervicography is based on the visual examination of certain changes on the cervix surface following the application of 5% acetic acid solution. The epithelium region, which changes color from pink to white with a certain degree of opacity, is of special clinical interest due to its suggestion of abnormality. This region is referred to as an *acetowhite* region. The color pictures taken during cervicography using a specially designed optical camera are called cervicographic images or cervigrams. Two examples of cervigrams having acetowhite regions are given in Figure 1.

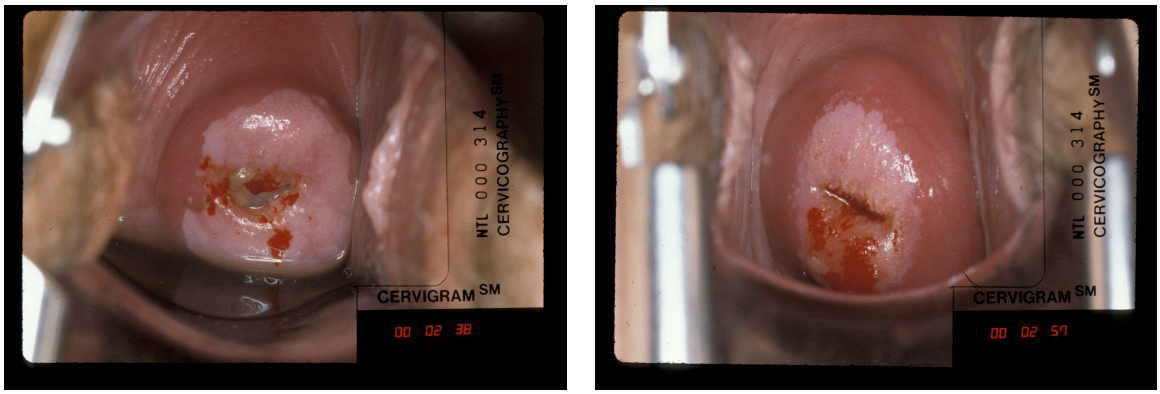

(a) Example 1 (b) Example 2

Figure 1 Examples of cervigrams

Color, opacity, margin shape, and surface contour of acetowhite areas are important features considered by physicians in diagnosing disease and severity. In addition vascular abnormality inside acetowhite areas is an important feature. Abnormal vascular patterns within acetowhite areas consist of three types: *punctation*,

In Proceedings of SPIE Medical Imaging 2009.

mosaicism, and *atypical vessels* [1]. Punctations refer to the appearance of capillaries shown as red / black points in a stippling pattern in an end-on view. Mosaicism refers to the appearance of inter-connecting blood vessels running parallel to the surface shown as cobbled areas of mosaic pattern. Unlike punctation and mosaicism, atypical vessels are irregular vessels having no specific patterns. In our work, we focus on identifying mosaicism. Mosaic areas may be classified as *fine* or *coarse*, based on vessel caliber and intercapillary distance. Coarse mosaic areas ("coarse mosaics") are formed by vessels having larger caliber and larger intercapillary distances, while fine mosaic areas are a network of fine-caliber blood vessels that are located close to one another. Coarse mosaics tend to be associated with more severe degrees of abnormality, such as high grade CIN lesions and early preclinical invasive cancer. A schematic representation of mosaics is shown in Figure 2 and an example of coarse mosaicism in a cervicographic image is shown in Figure 3.

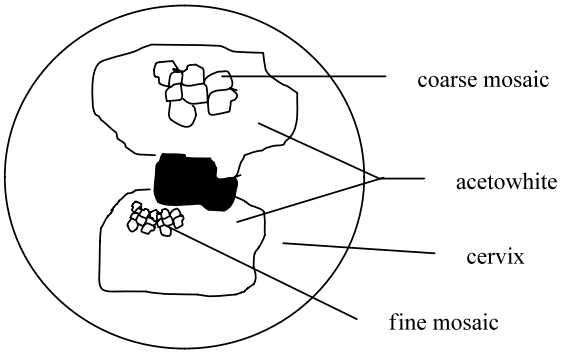

Figure 2. A schematic representation of mosaics

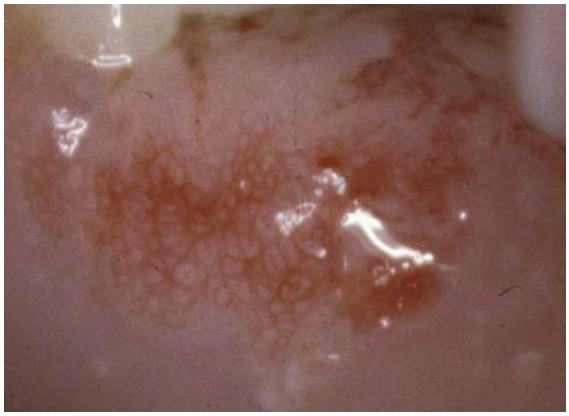

Figure 3. An example of mosaicism

For automated detection of mosaic areas in uterine cervix images, there are very few studies reported in the literature. Most of them [2,3,4] are based on extracting vascular structures using a series of mathematical morphological operations first, then extracting features as texture measures from those skeletonized vascular structures, and classifying them using unsupervised algorithms or minimum-distance classifier. The performance of these methods greatly relies on the accuracy of the first step, vascular structure extraction, which depends on hand-crafted heuristics and may suffer from over-segmentation. Therefore, a more robust method is needed, especially if the repertoire of images under consideration is very large and has large variations across images.

The cervigrams under consideration were taken during two major NCI-funded cohort studies in cervical cancer, the *Guanacaste* and *ALTS* projects. The *Guanacaste* project is a population-based natural history study of human papillomavirus (HPV) and cervical neoplasia in a rural area of Costa Rica [5]. The *ALTS* project was developed in four geographical areas of the United States, enrolling thousands of volunteer patients with abnormal Pap smear of ASCUS or LSIL [6]. During these projects, approximately 100,000 cervigrams were acquired; these have been digitized and archived by the National Library of Medicine (NLM) in order to manage, evaluate, and collect information from them, in collaboration with the National Cancer Institute (NCI). The cervigrams in the archive have large variations in appearance due to variations of illumination and intrinsic content variability, which makes the reliable segmentation of mosaic areas a challenging problem.

In this paper, we propose a new method of automatic segmentation and classification of mosaic patterns in cervigrams in which a support vector machine (SVM) classifier is applied, using learning from a "ground truth" dataset annotated by medical experts in oncology and gynecology. In our approach, the acetowhite region is split into tiles, and texture features are extracted from each tile. The SVM classifier is trained

using the texture features of tiles obtained from ground truth images. Given a new test image, the trained SVM classifier is applied to classify each tile in the test image, and the classified tiles are combined to generate the final segmentation map. The beauty of this approach is that it circumvents the step of vascular structure extraction used in [2,3,4], while taking full advantage of experts' knowledge to tackle the problem of large variations in appearance of target regions across images.

The rest of the paper is organized as follows. In Section 2, a brief description on the collection of ground truth data is provided. Section 3 gives the detailed explanation of the proposed algorithm. The results of experimental evaluation of the proposed algorithm and some discussions are presented in Section 4, followed by conclusions and future work given in Section 5.

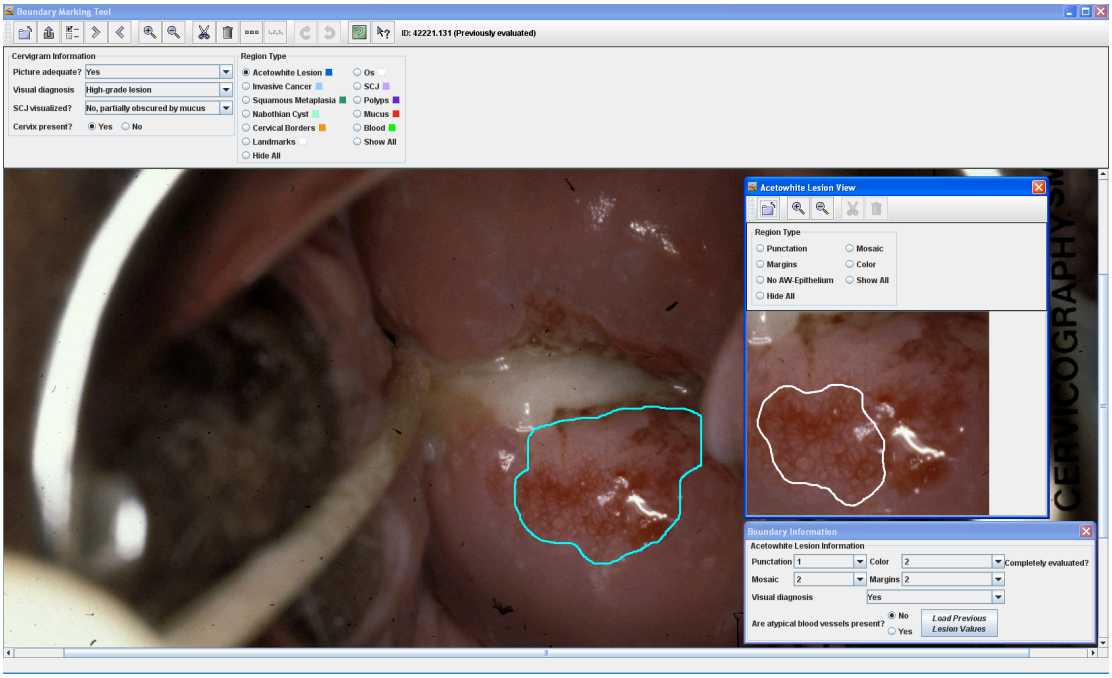

Figure 4. BMT with acetowhite and mosaic areas marked

2. GROUND TRUTH DATA COLLECTION

To collect the ground truth data used in this study, we used our in-house developed software: the Boundary Marking Tool (BMT). The BMT is a Web-accessible Java application which allows gynecologist experts in cervical cancer working at geographically dispersed sites to draw boundaries of anatomical regions on the image as well as entering a detailed set of data relevant to the evaluation of uterine cervix images [7]. For particular regions, the expert may display a detailed view that allows sub-classification of region contents. Relating to our study, for example, for acetowhite regions, the presence or absence of mosaicism is recorded, along with a classification of the mosaicism as coarse or fine. As shown in Figure 4, the acetowhite region (in blue) and the mosaic area inside the acetowhite region (in white) were both marked, and the mosaic was classified as coarse by the expert. All these BMT outputs including the spatial boundary data are saved as records in a central MySQL database that resides on a server at NLM.

3. MOSAIC DETECTION USING SVM

3.1. Pre-processing

Since abnormal vascular features such as mosaicism are significant only if they are within acetowhite areas [1], the first step of mosaicism detection is to extract acetowhite areas, so that searching for mosaicism may be confined to these areas. This approach reduces computation costs and increases detection accuracy. Our

collaborators have been conducting several studies on automatic segmentation and recognition of acetowhite epithelium. For example, Greenspan et al. [8] proposed an unsupervised segmentation algorithm based on Gaussian mixture models to identify squamous epithelium, columnar epithelium and acetowhite lesions. Yang et al. [9] developed a technique based on deterministic annealing and K-means clustering to segment and recognize acetowhite regions. Huang et al. [10] proposed a method for automated detection of acetowhite lesions using mean-shift clustering and SVM classification. For the work reported in this paper, we used cervigrams with acetowhite areas manually marked by medical experts using the BMT, in order to de-couple segmentation issues from the problem of mosaicism detection. In our pre-processing, the original image is first cropped according to the bounding box of the acetowhite region marked by experts. As shown in Figure 1, the area of acetowhite epithelium is only a fraction of the entire image. By this method the search space is greatly reduced.

The mosaicism could appear anywhere inside acetowhite regions, with variable shapes and sizes of mosaic sub-regions. To handle this variability and uncertainty, the cropped image is split into tiles with fifty percent of overlap, texture features are extracted from each tile, and each tile is independently classified as mosaic or non-mosaic using the SVM classifier. In implementation, the boundary of the image might be expanded based on the size of tile so that the whole acetowhite area will be taken into consideration, as illustrated by Figure 5. The index/location of each tile is recorded so that the SVM-classified tiles can be recombined into a segmentation map image in the post-processing stage.

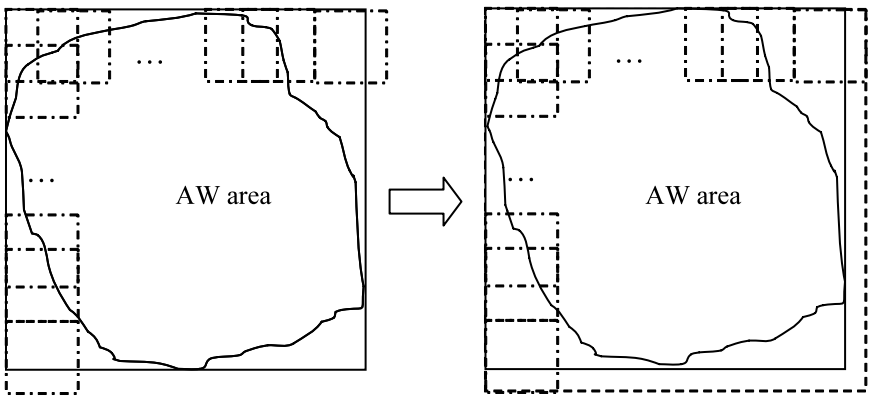

Figure 5. Illustration of image expansion

3.2. Texture Feature Extraction

The vascular patterns in cervigrams can be represented and characterized by texture information. In our work, we extracted four sets of features from each image tile after applying transforms based on the Gabor filter and Log-Gabor filter, discrete wavelet transform, and gray level co-occurrence matrix, respectively. The Gabor filter has been widely used for extracting texture features because of its optimal joint localization properties in space and frequency. To extract texture features from an image, we transformed the image by filtering with a set of Gabor filters of different orientations and spatial frequencies that appropriately cover the spatial frequency domain; then we extract certain features from the coefficients of the transform. In this study, we applied a Gabor filter bank with K orientations and S scales as described in [11]. We used the mean μ_{mn} and standard deviation σ_{mn} ($m = 1,...,S$; $n = 1,...,K$) of the magnitude of the transform coefficients to represent the image. We also used the Log-Gabor filter [12], which compensates for some limitations of the Gabor filter. Similarly as for the Gabor filter, we used a filter bank with K orientations and S scales, and calculated the mean μ_{mn} and standard deviation σ_{mn} of the magnitude of the transform coefficients in the region. To obtain additional, multiscale texture features we used the discrete wavelet transform [13]. For each sub-image obtained by applying the discrete wavelet

transform to the entire image, we calculated the mean and standard deviation of the magnitude of its wavelet coefficients as features. Finally, we applied the gray level co-occurrence matrix (GLCM) to capture the spatial dependence of gray-level values by using second-order statistics. Haralick [14] proposed a number of statistical features which were derived from the GLCM. We used four of them in our application: contrast, correlation, homogeneity, and energy.

3.3. SVM Training and Classification

The Support Vector Machine (SVM) is an effective and powerful classification algorithm [15]. It has been successfully used in various classification studies including medical applications, such as classifying lesion tissues in cervigrams [10] and detecting microcalcifications in mammograms [16]. In this paper, we investigated its use on mosaic detection in cervigrams. Specifically, we treated mosaic detection of each tile as a two-class classification problem. For each tile, we applied the SVM classifier to determine whether it was mosaic or non-mosaic.

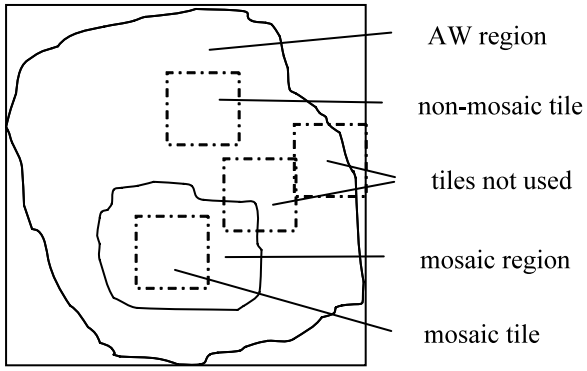

Figure 6. Selection of training tiles

The training tiles were obtained as follows. As described in the Section 3.1, each cropped input ground truth image was split into tiles with fifty percent overlap. Then, as illustrated in Figure 6, a tile which was completely within the mosaic region marked by experts was labeled as a mosaic tile and was treated as a positive training example. A tile which was totally inside the acetowhite region, but totally outside the mosaic region was labeled as a non-mosaic tile and was treated as a negative training example. Those tiles which did not fit in either category were not used as training examples. Let vector $\mathbf{x}_i$ denote the feature vector extracted from tile i $(i = 1,...,n)$. Let scalar $y_i \in \{+1, -1\}$ denote the class label of tile i. If tile i is a mosaic tile, then $y_i = +1$. If tile i is a non-mosaic tile, then $y_i = -1$. Given the set of $\{(\mathbf{x}_i, y_i),\ i = 1,2,...,n\}$ training examples, the SVM classifier finds the linear hyperplane that maximizes the separating margin between the positive and negative training examples in a high-dimensional feature space, and obtains the decision function with the form:

$$f(\mathbf{x}) = \sum_{i=1}^{n} \alpha_i y_i K(\mathbf{x}, \mathbf{x}_i) + b \tag{1}$$

where K is the kernel function, the α_i are the model coefficients, and b is the offset of the decision boundary from the origin. The coefficients α_i are obtained by solving the optimization problem:

Max

$$W(\alpha_1,...\alpha_n) = \sum_{i=1}^{n}\alpha_i - \frac{1}{2}\sum_{i=1}^{n}\sum_{j=1}^{n}\alpha_i\alpha_j y_i y_j K(\mathbf{x}_i,\mathbf{x}_j) \tag{2}$$

subj ect to

$$\sum_{i=1}^{n}\alpha_i y_i = 0,\ 0 \le \alpha_i \le C\ \ for\ i = 1,...,n$$

The parameter C is the regularization parameter and is selected by the user. A larger C corresponds to assigning a higher penalty to the training errors. In this study, the kernel function K which implicitly maps the input feature vector into a high dimensional feature space is set to be a radial basis function (RBF) defined as:

$$K(\mathbf{x},\mathbf{y}) = \exp\left(-\gamma\|x-y\|^2\right)\gamma > 0 \tag{3}$$

Therefore, given the training examples $\{(\mathbf{x}_i, y_i),\ i = 1,2,...,n\}$, only two parameters C and γ need to be determined in order to get the decision function $f(\mathbf{x})$ of the SVM. In implementation, the *C*-support vector classification (*C*-SVC) package of LIBSVM [18], an open source library for SVM, is used. As suggested in [17], in order to obtain good results fast and easily, the data (features) is linearly scaled to [0, 1] and the procedure of v-fold cross-validation is applied to find the best parameter values for C and γ; then these parameters C and γ are used to train the whole training set.

Given a new testing image, we first cropped it and split it into tiles as described in the preprocessing step. For each testing tile j $(j = 1,2...m)$, a feature vector $\mathbf{x}_j$ was computed and scaled first, and then classified by determining on which side of the decision boundary (the optimal separating hyperplane) it falls; it then was assigned one of the class labels +1 or –1, representing a mosaic tile or non-mosaic tile:

$$g(\mathbf{x}_j) = sign\left(f(\mathbf{x}_j)\right) = sign\left(\sum_{i=1}^{n}\alpha_i y_i K(\mathbf{x}_j,\mathbf{x}_i) + b\right) \tag{4}$$

3.4. Post-processing

In this step, the final segmentation map of the mosaic regions was obtained by combining the SVM-classified tiles. As illustrated in Figure 7, for each sub-tile (the patch with diagonal lines) whose size was a quarter of the size of the tile (because the tiles are fifty-percent overlapped with each other), all the tiles that enclose it (there are four tiles for the sub-tile shown in the Figure 7) will be checked. If at least one of them is labeled as +1, then the sub-tile is set to be a white (all ones) tile. If all of them are labeled as -1, then the sub-tile is set to be a black (all zeros) tile. After obtaining this segmentation map with white area indicating the mosaic regions of interest (ROI), the following procedure is applied to further refine the map. The area outside the marked acetowhite regions is set to be non-mosaic region since the mosaics are significant only when they are confined inside acetowhite regions. The isolated ROI whose area is too small (the size of the sub-tile) is removed. Finally, the segmentation map is cropped back to the original size if the image had been expanded during the stage of pre-processing.

Figure 7. The generation of segmentation map

4. EXPERIMENTS AND DISCUSSION

We tested the proposed algorithm using a set of cervigrams collected in the NCI *Guanacaste* and *ALTS* projects and annotated by medical experts, as described in Section 2. This dataset consists of 39 cervigrams in which both the boundary of acetowhite regions and the boundary of mosaic regions were expert-marked. In the experiment, the cervigrams were divided into two groups, with one group exclusively used as training images and the other group exclusively used as testing images.

The size of these color cervigrams is 2400 × 1640 pixels. The tile size was set to 64 × 64 pixels, which was found to adequately represent mosaic characteristics and result in good segmentation in experimental testing. The training group consisted of 19 cervigrams, which corresponds to 674 SVM positive training examples and 1789 SVM negative training examples. To extract features from each training example/tile, the parameters of texture descriptors were set as follows: for both Gabor filter-based texture descriptors and Log-Gabor filter-based texture descriptors, a filter bank with $K = 6$ orientations and $S = 4$ scales was used, which resulted in a feature vector of length 48 $\bar{f} = [\mu_{00}, \sigma_{00}, \mu_{01}, ...\mu_{S-1,K-1}, \sigma_{S-1,K-1}]$, respectively. For the DWT-based texture descriptor, a two-level wavelet transform was used, resulting in 7 sub-images and a feature vector of length 14 $\bar{f} = [\mu_0, \sigma_0, ..., \mu_8, \sigma_8]$. In the experiment, the GLCM features were calculated with 4 directions (0, 45, 90, and 135 degrees) and 4 distances (2, 4, 6, and 8 pixels). Therefore, 16 displacement vectors were used, and the feature vector contained 64 elements. The intensity of the image was quantized into 8 levels to reduce the computation cost. This resulted in a final feature vector length of 174. To train the SVM classifier with the RBF kernel, a fivefold cross-validation procedure was applied. The best cross-validation accuracy of 93 % was achieved when the regularization parameter, $C = 8$, and the kernel parameter, $\gamma = 0.5$. (The cross-validation accuracy is defined as the total number of correctly classified examples divided by the total number of examples classified in the procedure of cross-validation during training.)

The segmentation performance was initially evaluated by visual inspection. Overall, the proposed algorithm appears very promising in identifying the areas where mosaics are located. Eight pairs of example testing images are shown in Figure 8, in which all (a) images show the mosaic region marked by experts and (b) images show the region generated by the automatic segmentation method. The yellow boundary indicates the AW region and blue boundary identifies the mosaic region. The proposed method segments the majority of the mosaics, even though the appearance of the repetitive pattern of mosaics is variable across images. It bears mentioning that, with careful visual examination and comparison of both the ground truth marking and the segmentation results of the proposed approach, we found, for several cases, that the algorithm identified some areas which we believe are mosaic areas but which experts did not mark. An example is shown in Figure 8.5 (b). Further discussion of the results with experts will be conducted in the future to resolve this observation. Another aspect of the results which indicates need for further improvement is that, as shown in the results, the boundary of the segmented map is zig-zag due to the tile splitting and "piling" (placing tiles adjacently to form a connected region). The performance of segmentation was also quantitatively evaluated by two measures: true positive fraction (TPF) and false positive fraction (FPF) as defined below

$$TPF = \frac{S \cap R}{R}, \; FPF = \frac{|S - R|}{\overline{R}} \tag{5}$$

where R denotes the mosaic region marked by experts, $\overline{R}$ denotes its complement, S denotes the mosaic ROI segmented by the automatic approach. The TPF was 0.69, and the FPF was 0.16.

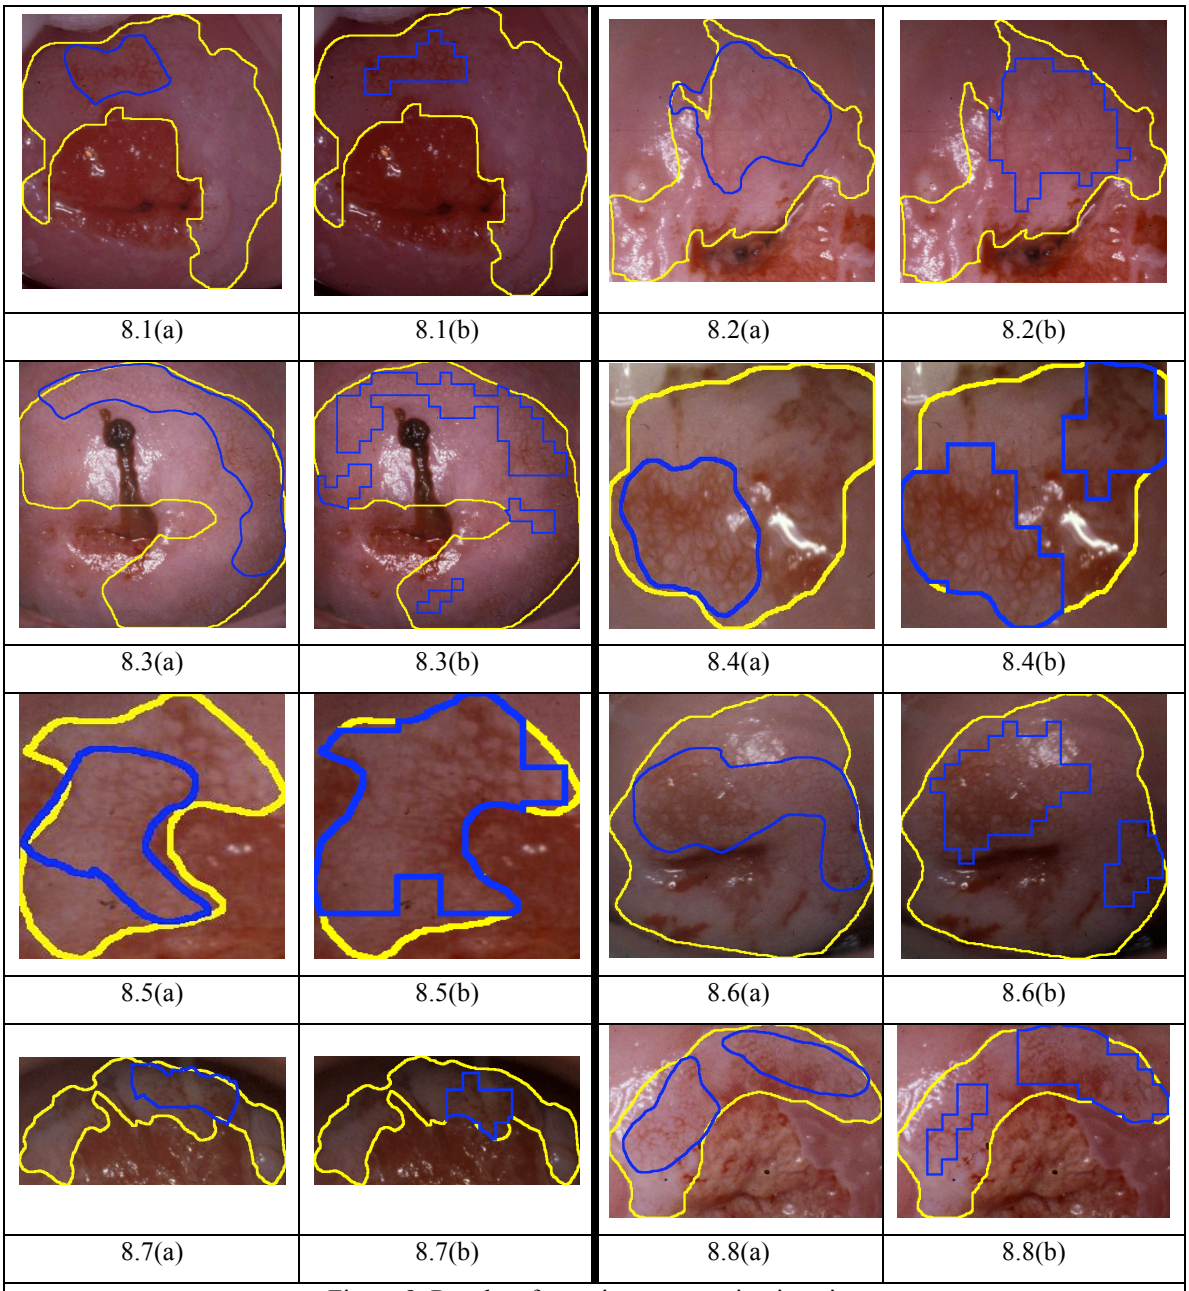

8.1(a)	8.1(b)	8.2(a)	8.2(b)
8.3(a)	8.3(b)	8.4(a)	8.4(b)
8.5(a)	8.5(b)	8.6(a)	8.6(b)
8.7(a)	8.7(b)	8.8(a)	8.8(b)

Figure 8. Results of mosaic segmentation in pairs.

(Left, (a): mosaic areas marked by experts. Right, (b): mosaic areas segmented by the proposed method)

5. CONCLUSIONS

Mosaicism is one of the important vascular abnormalities in cervicography which employs visual inspection of the cervix using acetic acid, and is indicative of a lesion with high CIN. Automatic detection of mosaicism in cervicographic images is a challenging task especially in a large image collection with high variation in appearance of the acetowhite and mosaic regions. In this paper, we have proposed a novel segmentation method based on SVM, which circumvents the step of vascular structure extraction, while taking advantage of expert knowledge. This work is part of our on-going project of developing a content-based image retrieval system for the NLM/NCI archive of 100,000 cervigrams [18, 19]. The performance

of the proposed algorithm was evaluated and the feasibility of the approach was demonstrated by a five-fold cross-validation score of 93%. Several tasks on further improving the proposed algorithm for mosaic segmentation, such as refinement of the segmentation boundary, testing on a larger dataset, and investigating more features, have been identified.

ACKNOWLEDGEMENT

This research was supported by the Intramural Research Program of the National Institutes of Health (NIH), National Library of Medicine (NLM), and Lister Hill National Center for Biomedical Communications (LHNCBC).

REFERENCES

1. J.W. Sellors and R. Sankaranarayanan, “Colposcopy and Treatment of Cervical Intraepithelial Neoplasia - A Beginner’s Manual”, Edited by J.W. Sellors and R. Sankaranarayanan, Published by the International Agency for Research on Cancer, France, 2003.
2. Q. Ji, J. Engel, and E. Craine, “Texture analysis for classification of cervix lesions”, IEEE Transactions on Medical Imaging, Vol. 19, No. 11, pp. 1144-1149, 2000.
3. B. Tulpule, D. L. Hernes, Y. Srinivasan, S. Mitra, Y. Sriraja, B. S. Nutter, B. Phillips, R. L. Long, D. G. Ferris, “A probabilistic approach to segmentation and classification of neoplasia in uterine cervix images using color and geometric features”, Proc. of SPIE Medical Imaging, Vol. 5748, pp. 995-1003, February 2005.
4. W. Li and A. Poisson, "Detection and characterization of abnormal vascular patterns in automated cervical image analysis," Lecture Notes in Computer Science: Advances in Visual Computing, Vol. 4292, pp. 627-636, 2006.
5. R. Herrero, M. H. Schiffman, C. Bratti, et al., “Design and methods of a population-based natural history study of cervical neoplasia in a rural province of Costa Rica: the Guanacaste project”, Rev Panam Salud Publica, No. 1, pp. 362-375, 1997.
6. M. Schiffman, M. E. Adrianza, “ASCUS-LSIL triage study: design, methods and characteristics of trial participants”, Acta Cytol, Vol. 44, No. 5, pp. 726-742, 2000.
7. J. Jeronimo, R. Long, L. Neve, et al. Digital tools for collecting data from cervigrams for research and training in colposcopy. Journal of Lower Genital Tract Disease, 10(1):16-25, January 2006.
8. S. Gordon, G. Zimmerman and H. Greenspan, “Image segmentation of uterine cervix images for indexing in PACS”, Proc. of the 17th IEEE Symposium on Computer-Based Medical Systems, pp. 298-303, CBMS 2004, Bethesda, MD, 2004.
9. S. Yang, J. Guo, P. King, Y. Sriraja, S. Mitra, B. Nutter, D. Ferris, M. Schiffman, J. Jeronimo, L. R. Long, “A multispectral digital cervigram analyzer in the wavelet domain for early detection of cervical cancer”, Proc. of SPIE Medical Imaging, Vol.5370, pp. 1833-1844, May 2004.
10. X. Huang, W. Wang, Z. Xue, S. Antani, L. R. Long, J. Jeronimo, "Tissue classification using cluster features for lesion detection in digital cervigrams", Proc. of SPIE Medical Imaging, Vol. 6914, pp. 69141Z-1-8, February 2008
11. B. S. Manjunath, W. Y. Ma, “Texture features for browsing and retrieval of image data”, IEEE Trans Pattern Analysis and Machine Intelligence, Vol. 18, No. 8, pp. 837-842, 1996.
12. D. J. Field. “Relations between the statistics of natural images and the response properties of cortical cells”, Journal of Optical Society of America, pp. 2379-2394, 1987.
13. M. N. Do, M. Vetterli, “Wavelet-based texture retrieval using generalized Gaussian density and Kullback-Leibler distance”, IEEE Transactions on Image Processing, Vol. 11, No. 2, pp. 146-158, 2002
14. R. Haralick, K. Shanmugam, I. Dinstein, “Texture features for image classification”, IEEE Trans Systems, Man and Cybernetics, Vol. 3, pp. 610-621, 1973.
15. V. Vapnik, Statistical Learning Theory. New York: Wiley, 1998.

16. I. El-Naqa, Y. Yang, M. N. Wernick, N. P. Galatsanos, and R. M. Nishikawa, “A support vector machine approach for detection of microcalcifications," IEEE Trans. on Medical Imaging, Vol. 21, NO. 12, pp. 1552-1563, 2002.
17. C. Chang and C. Lin, “LIBSVM: a library for support vector machines”, 2001. Software available at http://www.csie.ntu.edu.tw/~cjlin/libsvm.
18. L. R. Long, S. Antani, G. R. Thoma, “Image Informatics at a National Research Center”, Computerized Medical Imaging and Graphics, Vol. 29, pp. 171-193, February 2005.
19. Z. Xue, S. Antani, L. R. Long, J. Jeronimo, G. Thoma, “A Web-accessible content-based cervicographic image retrieval system”, Proc. of SPIE Medical Imaging, Vol. 6919, pp. 691907-1-9 February, 2008.