

A Data Driven Approach to Cervigram Image Analysis and Classification

Edward Kim and Xiaolei Huang

Abstract Cervical cancer is one of the leading causes of death for women worldwide. Early detection of cervical cancer is possible through regular screening; however, in developing countries, screening and treatment options are limited due to poor (or lack of) resources. Fortunately, low cost screening procedures utilizing visual inspection after the application of acetic acid in combination with low cost DNA tests to detect HPV infections have been shown to reduce cervical cancer by nearly 30%. To assist in this procedure, we developed an automatic, data centric system for cervigram (photographs of the cervix) image analysis. In the first step of our algorithm, our system utilizes nearly a thousand annotated cervigram images to automatically locate a cervix region of interest. Next, by utilizing both color and texture features extracted from the cervix region of interest on several thousand cervigrams, we show that our system is able to perform a binary classification on cervigram images with comparable accuracy to a trained expert. Finally, we analyze and report the effect that the color and texture features have on our end classification result.

1 Introduction

Cervical cancer afflicts an estimated 12,200 women in the US [1] and 529,800 women worldwide [2] every year. Fortunately, it can be cured, if it is detected during its early stages and treated appropriately. However, among the new cervical cancer cases found worldwide each year, 85% of them are in developing countries [2]. This disproportionate burden in low-resource world areas with medically underserved populations is mainly due to the lack of screening. Screening can prevent cervical

Edward Kim
Lehigh University, Bethlehem, PA, e-mail: edk208@lehigh.edu

Xiaolei Huang
Lehigh University, Bethlehem, PA e-mail: xih206@lehigh.edu

cancer by detecting Cervical Intraepithelial Neoplasia (CIN), also known as cervical dysplasia. The CIN classification is specified in several grades: CIN1 (mild), CIN2 (moderate), and CIN3 (severe). In a clinical setting, one of the most important goals of screening for cervical cancer is the differentiation of normal/CIN1 from CIN2/3+. If a lesion is classified as CIN2/3+, it will require treatment whereas mild dysplasia in CIN1 typically will be cleared by immune response in a year or so, and thus can be observed or treated more conservatively.

To address the problem of CIN classification, we utilize a low cost, photographic screening test, called Cervicography. The photographs of the cervix, or cervigrams, can provide valuable and insightful information to assist in diagnosis and disease grading. One of the most important observations in cervigrams is the acetowhite region which is caused by the whitening of potentially malignant regions of the cervix epithelium after applying dilute (3-5%) acetic acid. All forms of precancerous tissue exhibit some degree of opacity, or acetowhiteness, after contact with acetic acid. Thus, accurately interpreting the severity of this tissue region is critically important to cervigram image analysis. Additionally, other visual features or observations can assist with disease classification. These features include the identification or presence of mosaicism, punctation, atypical vessels or vasculature, blood, polyps, cyst, etc. However, as cervigram regions have very high variability in color, shape, and size, it is difficult to identify and characterize these regions individually for both trained medical professionals and computer algorithms.

2 Related Work

In recent years, there have been several automatic or semi-automatic image analysis algorithms applied to cervigram images. A common process in many of these previous works was the automatic detection of the cervix region. This region of interest (ROI) contains the relevant information necessary for accurate tissue and disease classification. In Li et al. [13], the region of interest is found by the analysis of local color features and optimized through expectation maximization. Zimmerman et al. [18] developed a two-stage segmentation process utilizing image intensity, saturation, and gradient information and reported their results on 120 images. Gordon et al. [9] uses a Gaussian mixture model to automatically find the cervix region of interest, and then separates the cervix tissue region into three types: the columnar epithelium, the squamous epithelium, and the acetowhite region. Gordon et al. also tested on a set of 120 cervigram images. Xue et al. [17] focuses on the removal of specular regions and the identification of the acetowhite region in the ROI. Similarly, Xue et al. tested on 120 cervigram images, and used $L^*a^*b^*$ color features, Gaussian mixture models, and k-means clustering to achieve their results.

Further image classification tasks in the region of interest can be performed as exhibited in several previous works. In Ji et al. [11], the authors use texture features to recognize important vascular patterns found in cervix images. They collected 5 images per vascular pattern class (network, hairpin, punctation1, punctation2, mo-

saic1, mosaic2) for a total of 30 images. Similarly, Srinivasan et al. [16] uses a filter bank of texture models for recognizing punctation and mosaicism on ten images.

As demonstrated by the previous work, there are many complex visual features that contribute to the problem of cervigram image analysis. The isolation of the cervix region of interest is an important first step used to remove the unwanted effects of the background image noise. Then, analysis of the region of interest can be performed by looking at color and texture features. Color features play a key role in the cervix and tissue classification task, whereas texture also plays an important role in the identification of mosaicism and vessel pattern analysis. Further, given the high variability exhibited by cervigrams, testing on larger datasets is necessary to validate the effectiveness of cervigram image analysis on real world datasets.

3 Methodology

In our work, we develop a unique approach to cervigram image analysis. In contrast to many previous works that utilize a more generative model towards cervigram image analysis, we developed a discriminative, data centric system that would be able to utilize similar cervigram cases in a collection of annotated cervigrams to perform a binary classification, e.g. normal/CIN1 and CIN2/3+.

To be more explicit, we do not attempt to directly characterize the visual properties present in cervigram images. Instead, we utilize thousands of training images collected by the National Cancer Institute (NCI) and National Library of Medicine (NLM) to classify a new cervigram image. For this process, we will be utilizing two distinct databases. The first database consists of 939 expertly labeled cervigrams. There are detailed annotations linking these images to expert markings including the delineation of the cervix region of interest. We will refer to this database as D^1 . Our second database, D^2 , is of a larger scale and contains tens of thousands of patient records and cervigram images [10]. Each record also has been labeled with a final outcome, which we can utilize in the classification stage of our system. The final outcome is determined by expert practitioners and has been given a final diagnosis (normal/CIN1, CIN2/3+). These expert annotations of final diagnosis are based on the analyzing the histology of the patient images, a commonly used gold standard to define the ground truth diagnosis.

In summary, our system takes as input a new test cervigram image and uses large amounts of training data to reach a final disease classification. This classification result involves several steps. The first step is the translation of raw image data into a compact color and texture feature representation. Using our representation, we can then attempt to leverage our first database of annotated cervigrams to isolate the cervix region of interest. Finally, we can again utilize our second database and the specific visual features located in the ROI of a given cervix region to obtain a final disease classification. By using the databases in our classification task, we are indirectly utilizing the variables that went into the diagnosis of a patient cervigram image, without having to individually model the complex visual characteristics.

3.1 Visual Feature Extraction and Representation

Through our research, and as exhibited in many previous works, we have found that both color and texture features are necessary to represent the visual cues present in cervigram tissue regions (acetowhite regions, mosaicism, punctation, etc.). Additionally, the relative size and position of abnormal characteristics are also important to capture in our feature representation. Thus, we utilize a spatial pyramid of color and texture features as described in Lazebnik et al. [12]. This spatial pyramid representation is able to preserve the geometric correspondence of visual features.

Color features - Color plays an important role in cervical lesion identification and classification. One of the most important visual features on the cervix that have relevant diagnostic properties is the presence of Acetowhite regions, or the whitening of potentially malignant cervical regions with the application of dilute acetic acid. The perceived color and thickness of an Acetowhite region is also relevant to cervical lesion grading. Thus, we extract pyramid color histogram features, PLAB, from a cervigram image to represent various color regions. We convert the pixel colors in a cervigram into the perceptually uniform $L^*a^*b^*$ color space. A property of this color space is that a small change in the color value corresponds to about the same small change in visual appearance. Our PLAB descriptor is also able to represent local image color and its spatial layout. For each channel (L^* , a^* , or b^*) of the color space, we extract 3 pyramid levels, with a 16 bin histogram from each region. A pyramid is constructed by splitting the image into rectangular regions, increasing the number of regions at each level. Thus, a single channel histogram consists of 336 bins, and our complete PLAB descriptor consists of 1008 bins.

Texture features - Texture features play an important role in representing various vasculature patterns, punctation, mosaicism, and tissue thickness characteristics. Similar to our color features, we represent texture as a pyramid histogram of oriented gradients, or PHOG feature [4]. The PHOG descriptor represents local image shape and its spatial layout. The shape correspondence between two images can be measured by the distance between their PHOG descriptors using a spatial pyramid kernel. To extract the PHOG descriptors from a cervigram image, we first compute the gradient response using a sobel edge filter. If we use an 8 bin orientation histogram over 4 levels, the total vector size of our PHOG descriptor for each image is 680 bins. For an illustration of our PLAB and PHOG feature, see Figure 1.

Image similarity measurement - To compute the image similarity, we use a weighted sum of the similarities between the two images' color and texture features. The cost function that measures this dis-similarity, or distance, is defined as,

$$C_s(X, Y) = \lambda(d(X^c, Y^c)) + (1 - \lambda)(d(X^t, Y^t)) \quad (1)$$

Where X, Y represent two distinct images, X^c, Y^c are the PLAB color feature vectors of X and Y , and X^t, Y^t are the PHOG texture features of X and Y respectively. The distance between feature vectors is computed by d , and the λ term weights the influence of the two features on the final similarity computation.

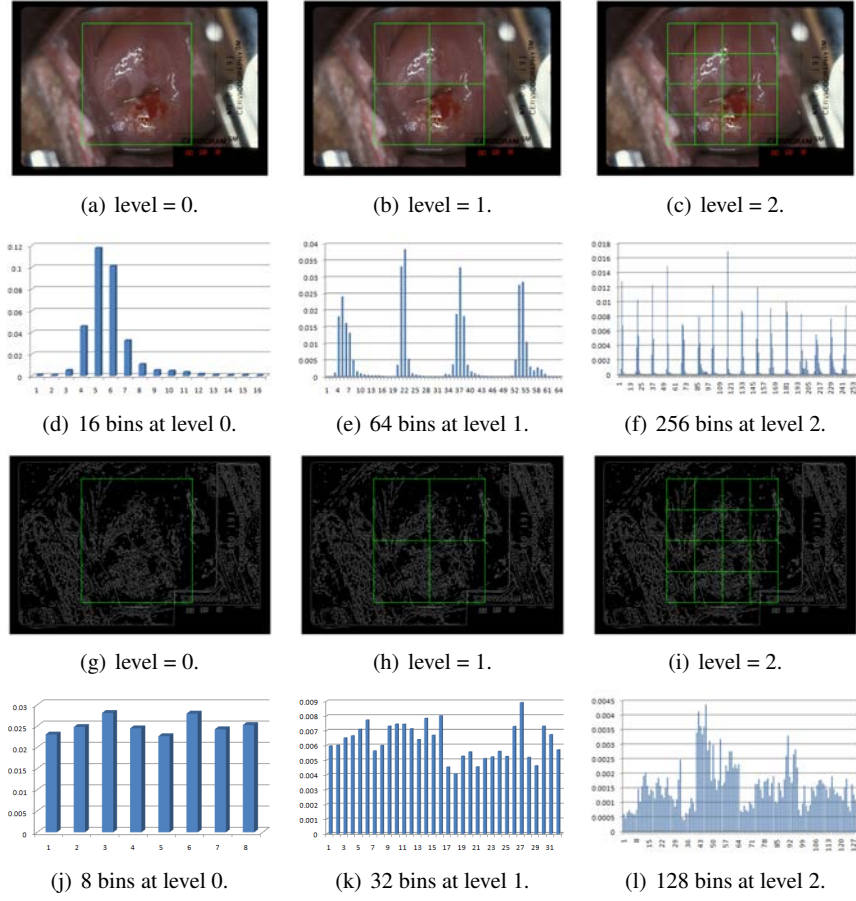


Fig. 1 Example of the PLAB and PHOG features extracted at multiple levels in a rectangular region of interest. In (a)-(c) the $L^*a^*b^*$ color space is sampled into 16 bins per region (per channel). The L^* channel in the PLAB feature vector is represented in (d)-(f). The edges of the input image are computed by a sobel edge filter and partitioned into a pyramid of regions (g)-(i). 8 orientation bins are extracted from each rectangle and concatenated into the PHOG feature vector represented in (j)-(l).

For the distance measure between two histogram-like feature vectors, h_X and h_Y , we use the χ^2 measure defined as,

$$d(h_X, h_Y) = \chi^2(h_X, h_Y) = \frac{1}{2} \sum_{k=1}^K \frac{[h_X(k) - h_Y(k)]^2}{h_X(k) + h_Y(k)} \quad (2)$$

Where K is the total number of bins present in the feature vectors.

3.2 Finding the Region of Interest

In the first step of our algorithm, we isolate the (cervix) region of interest in a new cervigram. Given a new cervigram image our goal is to annotate a tight bounding box around the cervix. Some previous works have used the local color and position features in order to isolate the cervix region [9, 8, 17]. However, due to the high variability in color, size, and position of the cervix in cervigrams, these approaches based solely on local image features suffer from low specificity. In contrast, we take a different approach to the region of interest detection problem. Our approach is data driven; we rely on an expertly labeled database of 939 cervigram images with their delineated rectangular regions of interest in order to find a suitable bounding box for the region of interest in a new cervigram image.

We will refer to this step in our algorithm as our **optimized bounding box** method, and is defined as the following. Given a new cervigram test image, we extract the color and texture features from the whole image and compute the image similarity between this image and every other image in our 939 database (D^1) by Equation 1. We sort the list of images in the database by decreasing similarity and extract the top M matching cervigram images. These images should *globally* resemble the test image; however, there is no guarantee that the cervix region of the top images match the location and size of the test image. Thus, we only use the top M matching cervigrams for their annotated ROI and use these ROIs as candidate bounding boxes. We denote the top M matching cervigrams' ROIs as B_m where $m = [1..M]$, and our ground truth bounding boxes in D^1 as D_n^1 , where $n = [1..N]$. Then, for every candidate bounding box, we recompute the color and texture features of the test image inside the candidate ROI. We then compute the similarity between each candidate bounding box and every ground truth cervix ROI in our first database of $N(=939)$ images. Among the MXN comparisons, we find the pair of ROIs that gives the smallest distance, and the candidate ROI in this pair will be our final ROI for the test image.

Mathematically speaking, we choose the minimum distance bounding box pair to obtain our final ROI, \hat{B}_m ,

$$\langle \hat{B}_m, \hat{D}_n^1 \rangle = \arg \min_{\langle B_m, D_n^1 \rangle} C_s(D_n^1, B_m) \quad n \in [1..N], m \in [1..M] \quad (3)$$

3.3 Cervigram classification

Given the cervix region of interest, we can now more accurately match a test cervix region to our database of cervigram ROIs. Our second database, D^2 , consists of thousands of cervigrams that have been analyzed by experts and given a disease diagnosis of CIN1 through CIN3. This data can be utilized to train a classifier which will compute a disease classification for the given test image.

We build two classification methods into our system, a support vector machine classification and a majority vote classification.

3.4 Support Vector Machine classification

In D^2 , we compute the ROIs for every image, and extract their color and texture features. We then build a binary linear classifier that can discriminate between normal/CIN1 and CIN2/3+ using a Support Vector Machine (SVM) [5]. We extract our color and texture features from the cervix ROI and concatenate these vectors into a single 1688 bin vector (1008 bin color vector + 680 texture vector). We can build a SVM model based upon a subset of D^2 , which will attempt to separate the two classes (normal/CIN1 and CIN2/3+) in a high dimensional space. Given a new cervigram ROI, the SVM model can predict a CIN classification based upon the decision boundary obtained by our training data. We describe the size of our training and test set, as well as the parameters of our SVM model, and how we obtained them in our results section.

3.5 Majority vote classification

For our majority vote classifier, we again compute the extracted color and texture features in the cervix ROI. We can then compute a matching score between the test cervix ROI and every other ROI in our second database. Using the matching score, we can sort the similarity of the images to the new image and find the top Q most similar cases. These top cases vote on a classification, where the majority vote label is selected as the final output of our system. Mathematically speaking, given Q , the top cluster of $|q| = Q$ cases can be obtained by minimizing,

$$C(X) = \sum_{q=1}^Q C_s(\hat{D}_q^2, \hat{B}_X) \quad q \subset D^2 \quad (4)$$

Where \hat{D}^2 are the computed ROIs from our second database, and X is our test cervigram, and \hat{B}_X is its corresponding ROI. When minimized, the set of q consists of the top Q matches from our second database. Each $q \in Q$ has a corresponding binary label i.e. less than CIN2 or CIN2/3+, and will cast a vote for the final classification. The majority vote label is assigned to the test image, X .

4 Results

We perform several experiments to evaluate our system. Our first experiment measures how accurately we are able to isolate the cervix region of interest. Our second experiment measures the ability of our system to correctly classify a new cervigram image. Finally, our third experiment measures the effect that our color and texture features have on the final outcome.

4.1 Isolating the region of interest

In this experiment, we analyze how accurately we are able to detect the region of interest in a new cervigram image. This experiment tests on 450 cervigrams in D^1 , but utilizes all of the 939 (minus the test cervigram, e.g. leave-one-out) expertly annotated bounding boxes to obtain the final result as described in Equation 3. The majority of the images are of the same size and resolution; however, there still is a small amount of variability. Our features representations are normalized to account for this variability.

To measure the accuracy of our region of interest calculation, we use the Jaccard similarity coefficient defined as, $J(A, \hat{B}) = \frac{|A \cap \hat{B}|}{|A \cup \hat{B}|}$ which measures the similarity of our bounding box calculation to the ground truth region of interest specified by a trained physician. In this equation we can view A as the ground truth cervix region bounding box, and \hat{B} as the minimum bounding box found in Equation 3. In essence, the Jaccard coefficient can be viewed as the area of intersection of bounding boxes, divided by the total area covered by both bounding boxes. A Jaccard similarity coefficient closer to one has a greater similarity to the ground truth; whereas, a coefficient value close to zero has nearly no overlap area similar to the ground truth. In this experiment, we compare three different methods.

1. **Image Bounding Box (IBB)** - In this method, we choose the most similar image to the test image, based upon a global image similarity. The bounding box is transferred from the top match to the test cervigram. This method does not eliminate the noise present outside the region of interest.
2. **Average Bounding Box (ABB)** - This method again uses global image similarity; however, we take the top K matches and average the position of their bounding boxes to achieve our final bounding box. We used a gradient descent on the Jaccard similarity coefficient (with respect to D^1) to obtain the best K ($= 15$).
3. **Optimized Bounding Box (OBB)** - This is our method described in section 3.2 that utilizes candidate bounding boxes from a global image match. We compute the final bounding box by finding the *argmin* in equation 3. For our experiments, we set M to be 100 i.e. we have 100 candidate bounding boxes to choose from.

We show several sample images, their corresponding ground truth bounding boxes, and the results of these three methods in Figure 2.




Fig. 2 Sample region of interest calculations on five (a)-(e) images. The ground truth (1) and expert annotations are shown here. The image bounding box (IBB) (2), average bounding box (ABB) (3), and optimized bounding box (OBB) (4) and their corresponding Jaccard (jsc) coefficients are also displayed. In column (c) row(4), we see a low Jaccard index, but visually close ROI calculation.

However, the Jaccard index does not describe the entire story. As seen in Figure 2(4)(c), the Jaccard index can be quite low, yet the cervix region may be more accurately located than indicated. To further describe the accuracy of our cervix ROI calculation, we compute the Euclidean distance between the ground truth centroid position, and our computed ROI centroid position. Additionally, we record the difference of the aspect ratio between the ground truth bounding box and our computed rectangle as further evidence. The aspect ratio is computed by the absolute difference between the width divided by height of the computed bounding box and the width divided by height of the ground truth bounding box. In Table 1 we report the average Jaccard index values, centroid difference, and aspect ratio difference. And as seen from our results, our optimized bounding box method outperforms the other two methods, and has the desired effect of maintaining a low centroid difference and consistent aspect ratio.

4.2 Accuracy of our disease classification

In our second experiment, we utilize a subset of 2,000 cervigram images obtained from the NIH/NCI database D^2 , consisting of 1,000 normal/CIN2 grade images and

Table 1 Comparison of three different region of interest methods. We report the average Jaccard index, centroid distance, and aspect ratio difference of the computed bounding box and ground truth bounding box on 450 test images.

Method	Jaccard Index (std)	Centroid Dist.	Aspect Ratio
Image Bounding Box (IBB)	0.611 (0.012)	32.37	0.148
Average Bounding Box (ABB)	0.699 (0.013)	23.91	0.109
Optimized Bounding Box (OBB)	0.736 (0.014)	25.72	0.096

std = Standard Deviation

1,000 CIN2/3+ cases. We perform a ten fold cross validation, binary classification on this dataset to evaluate how well our system is able to differentiate between these classes. We test using two classifiers, our majority vote classifier described in section 3.5 and a linear Support Vector Machine described in section 3.4.

There are several parameters that need to be set for both classification methods. For our majority vote technique, we train the number of voting cases, Q , by computing the Dice similarity coefficient, $DSC = \frac{2 \cdot TP}{(2 \cdot TP + FP + FN)}$ over our training set while varying Q between the top matching case ($Q = 1$) to the top fifty ($Q = 50$) matches. In the Dice similarity coefficient, TP denotes true positive cases, FP denotes false positive cases, and FN are false negative cases. As we increase Q , the DSC score steadily increases from 0.70 at $Q = 1$ and asymptotically approaches 0.75 when $Q > 30$. Therefore, we set Q to be the top 33 most similar cervigram images to the input image that will vote to obtain the final classification output of our system. For the weight parameter of our color versus texture features (λ value), we chose $\lambda = 0.7$. The analysis behind this value can be seen in the next section, Section 4.3.

For our linear SVM, we train the parameters of our model using a five-fold cross validation on the training images, and use this model to classify the new input image. The results of both our methods can be seen in Table 2. Additionally, in this table, we also report comparative results from multiple studies around the world as reported by Sankaranarayanan et al. [15]. In these studies, direct visual inspection of the cervix was conducted after the application of acetic acid, and each patient is given a result of positive or negative for CIN2/3+.

4.3 Weighted effect of Color and Texture

In our final experiment, we view the effects of the color and texture features on our final classification result. Using our majority vote classifier, and altering the λ in equation 1, we can evaluate the influence of these features. As a baseline data point, a value of $\lambda = 0.5$, gives both color and texture features equal weight. On a set of 1,000 cervigram images from D^2 (500 normal/CIN1, 500 CIN2/3+), we

Table 2 Comparison of Sensitivity, and Specificity for different classification methods for detecting CIN2/3+. Our automatic classification method is comparable to manual inspection by experts.

Method	Samples ^a	Sensitivity,% (95% CI)	Specificity,% (95% CI)
Our automatic method ^b , Majority Vote	2,000	73 (65-81)	77 (67-87)
Our automatic method ^b , L-SVM	2,000	75 (69-82)	76 (66-86)
Denny et al. [6], 2000, South Africa	2,885	67 (56-77)	84 (82-85)
Belinson et al. [3], 2001, China	1,997	71 (60-80)	74 (71-76)
Denny et al. [7], 2002, South Africa	2,754	70 (59-79)	79 (77-81)
Sankaranarayanan et al. [14], 2004 India and Africa	54,981	79 (77-81)	86 (85-86)

^a Samples for our method are number of distinct cervigram images. Samples for the comparative studies correspond to number of patients.

^b Confidence intervals calculated by ten-fold cross validation
 CI = Confidence Interval

vary the parameter from 0 - 1, where a λ of less than 0.5 gives more weight to our texture feature and a λ of greater than 0.5 gives more weight to our color feature. We analyze how this parameter affects our sensitivity and specificity of our majority vote classifier on dataset, D^2 and display the results in Figure 3.




Fig. 3 The effect of changing our λ value between 1 - 0 when using a majority vote classifier. A balanced weight of 0.5 has equal contribution of color and texture, whereas a value closer to 1 favors the color feature. Empirically, a λ value of 0.7 has the best specificity and sensitivity on D^2 .

From these results it is clear that the color feature plays a vital role in the CIN designation for cervigram images. By using the color feature alone ($\lambda = 1$), we are still able to achieve fairly good results, but if we only use texture features ($\lambda = 0$), our sensitivity and specificity drop dramatically.

5 Discussion and Conclusion

We present an automatic cervigram image analysis algorithm that is able to isolate the region of interest in a new cervigram image, and ultimately classify the image as less than CIN2 or CIN2/3+. Our system is data centric, meaning we do not attempt to directly model the complex features present in the cervix anatomy, but rather utilize thousands of training images to perform our analysis. Furthermore, our system performs well, and is shown to be comparable to human observers. To improve our results, we have been exploring multi-modal classification methods by utilizing patient data such as patient age, HPV types (16/18/31), and health behavior (history of smoking). Another possible improvement would be the use of vector weights to find the optimal balance of color and texture. Because a color change of the acetowhite region has a high correlation to the severity of a CIN classification, this would suggest that a color feature would be important in cervigram image analysis. Through our experiments, we were able to confirm that color does in fact play a vital role in the CIN classification task.

Acknowledgements This research was supported in part 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 (LH- NCBC). The image and clinical data for this work comes from the National Cancer Institute (NCI) Guanacaste/ALTS projects.

References

1. Atlanta: Cancer facts and figures. American Cancer Society (2010)
2. Atlanta: Global Cancer facts and figures, 2nd edition. American Cancer Society (2011)
3. Belinson, J., Pretorius, R.: Cervical cancer screening by simple visual inspection after acetic acid. *Obstetrics & Gynecology* **99**(3), 518 (2002)
4. Bosch, A., Zisserman, A., Munoz, X.: Representing shape with a spatial pyramid kernel. pp. 401–408 (2007)
5. Chang, C.C., Lin, C.J.: LIBSVM: a library for support vector machines (2001). Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>
6. Denny, L., Kuhn, L., Pollack, A., Wainwright, H., Wright Jr, T.: Evaluation of alternative methods of cervical cancer screening for resource-poor settings. *Cancer* **89**(4), 826–833 (2000)
7. Denny, L., Kuhn, L., Pollack, A., Wright Jr, T.: Direct visual inspection for cervical cancer screening. *Cancer* **94**(6), 1699–1707 (2002)
8. Gordon, S., Zimmerman, G., Greenspan, H.: Image Segmentation of Uterine Cervix Images for Indexing in PACS. In: *Proceedings of the 17th IEEE Symposium on Computer-Based Medical Systems* (2004)
9. Gordon, S., Zimmerman, G., Long, R., Antani, S., Jeronimo, J., Greenspan, H.: Content analysis of uterine cervix images: initial steps towards content based indexing and retrieval of cervigrams. In: *Proc. of SPIE medical imaging*, vol. 6144, pp. 1549–1556 (2006)
10. Jeronimo, J., Long, L., Neve, L., Michael, B., Antani, S., Schiffman, M.: Digital tools for collecting data from cervigrams for research and training in colposcopy. *Journal of Lower Genital Tract Disease* **10**(1), 16 (2006)
11. Ji, Q., Engel, J., Craine, E.: Classifying cervix tissue patterns with texture analysis. *Pattern Recognition* **33**(9), 1561–1574 (2000)

12. Lazebnik, S., Schmid, C., Ponce, J.: Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In: *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*, vol. 2, pp. 2169–2178 (2006)
13. Li, W., Gu, J., Ferris, D., Poirson, A.: Automated image analysis of uterine cervical images (2007)
14. Sankaranarayanan, R., Basu, P., Wesley, R., Mahe, C., Keita, N., Mbalawa, C., Sharma, R., Dolo, A., Shastri, S., Nacoulma, M., et al.: Accuracy of visual screening for cervical neoplasia: results from an iarc multicentre study in india and africa. *International Journal of Cancer* **110**(6), 907–913 (2004)
15. Sankaranarayanan, R., Gaffikin, L., Jacob, M., Sellors, J., Robles, S.: A critical assessment of screening methods for cervical neoplasia. *International Journal of Gynecology & Obstetrics* **89**, S4–S12 (2005)
16. Srinivasan, Y., Nutter, B., Mitra, S., Phillips, B., Sinzinger, E.: Classification of cervix lesions using filter bank-based texture mode (2006)
17. Xue, Z., Antani, S., Long, R., Thoma, G.: Comparative performance analysis of cervix roi extraction and specular reflection removal algorithms for uterine cervix image analysis. In: *Proc SPIE*, vol. 6512, pp. 411–419 (2007)
18. Zimmerman-Moreno, G., Greenspan, H.: Automatic detection of specular reflections in uterine cervix images. In: *SPIE Medical Imaging Symposium* (2006)