Computer Assisted Detection of Polycystic Ovary Morphology in Ultrasound Images

CONFERENCE PAPER · JUNE 2007

DOI: 10.1109/CRV.2007.18 · Source: IEEE Xplore

4 AUTHORS, INCLUDING:

Roger A Pierson
University of Saskatchewan
286 PUBLICATIONS **6,372** CITATIONS
SEE PROFILE

Eric Neufeld
University of Saskatchewan
95 PUBLICATIONS **522** CITATIONS
SEE PROFILE

Computer Assisted Detection of Polycystic Ovary Morphology in Ultrasound Images

Maryruth J. Lawrence
Department of Biomedical Engineering
57 Campus Drive
University of Saskatchewan
Saskatoon, SK, CANADA, S7N 5A9
pradeepa_a@yahoo.com

Mark G. Eramian
Department of Computer Science
110 Science Place
University of Saskatchewan
Saskatoon, SK, CANADA, S7N 5C9
eramian@cs.usask.ca

Roger A. Pierson
Obstetrics, Gynecology, and
Reproductive Sciences
103 Hospital Drive
University of Saskatchewan
Saskatoon, SK, CANADA, S7N 0W8
pierson@erato.usask.ca

Eric Neufeld
Department of Computer Science
110 Science Place
University of Saskatchewan
Saskatoon, SK, S7N 5C9, CANADA
eric@cs.usask.ca

Abstract

Polycystic ovary syndrome (PCOS) is an endocrine abnormality with multiple diagnostic criteria due to its heterogenic manifestations. One of the diagnostic criteria includes analysis of ultrasound images of ovaries for the detection of number, size, and distribution of follicles within the ovary. This involves manual tracing and counting of follicles on the ultrasound images to determine the presence of a polycystic ovary (PCO). We describe a novel method that automates PCO detection. Our algorithm involves segmentation of follicles from ultrasound images, quantifying the attributes of the automatically segmented follicles using stereology, storing follicle attributes as feature vectors, and finally classification of the feature vector into two categories. The classification categories are: PCO present and PCO absent. An automatic PCO diagnostic tool would save considerable time spent on manual tracing of follicles and measuring the length and width of every follicle. Our procedure was able to achieve classification accuracy of 92.86% using a linear discriminant classifier. Our classifier will improve the rapidity and accuracy of PCOS diagnosis, reducing the risk of the severe complications that can arise from delayed diagnosis.

1. Introduction

Polycystic ovary syndrome (PCOS) is an ovarian abnormality that affects 5 - 10 % of women of reproductive age [6]. The exact cause of PCOS is not well established, but insulin resistance and androgen excess play an important role on its onset. Insulin resistance leads to excess insulin levels causing over-production of male hormones, elevated levels of blood fats, and subsequent ovarian dysfunction. Women with PCOS are at greatly increased risk of cardiovascular disease, diabetes, and obesity; however, most are first diagnosed by a reproductive endocrinologist upon referral for infertility.

Symptoms of PCOS include irregular menses, infertility, obesity, excessive production of male hormones exhibited by male-pattern facial and bodily hair growth, acne, and male-pattern balding [1]. It is therefore important to screen for polycystic ovaries in order to improve the rapidity with which this condition can be diagnosed. An automatic detection algorithm, such as that described herein, could be applied to routine scans and facilitate early, accurate diagnosis by family physicians and general radiologists, and allow intervention that may abate or obviate the severe consequences of the disease.

The current criteria for the diagnosis of PCOS was established jointly by the American Society for Reproductive Medicine (ASRM) and the European Society of Human Reproduction and Embryology (ESHRE). There are three cri-

Fourth Canadian Conference on Computer and Robot Vision(CRV'07)
0-7695-2786-8/07 $20.00

teria for the diagnosis of PCOS: Oligo- and/or anovulation (failure to ovulate), clinical and/or biochemical signs of excessive production of male hormones, and the presence of polycystic ovaries in at least one of the ovaries (in an ultrasound examination). The ASRM/ESHRE recommend that PCOS should be diagnosed if two of these three criteria are met [2].

By current convention, the ultrasonographic morphology of a polycystic ovary (PCO) is characterized by the presence of 12 or more ovarian follicles (roughly spherical fluid-filled cavities in which eggs develop) which are 2-9mm in size, and/or a total ovarian volume of more than 10cm^3 [2]. In contrast, a normal ovary usually contains fewer than 10 follicles with one dominant follicle reaching a maximum diameter of 20-25mm prior to ovulation. Conventional detection of PCOS involves analyzing ovarian ultrasound images for PCO morphology, and testing for biochemical/clinical signs for excess male hormones. Ovarian ultrasound images are analysed for PCO by counting the number of follicles, identifying follicle size and distribution, and evaluating the ratio of number of follicles to ovarian volume.

Accurate analysis of the pelvic ultrasound images is important for the detection of PCO. Presently, most ultrasound images are analysed manually for diagnostic purposes. Hence there is a high variance in the detection of PCO among different gynecologists/radiologists. Our work aims to automate the analysis of pelvic ultrasound images for detecting PCO by examining the number, size, and distribution of follicles within the ovary. This was achieved using a three step process involving segmentation of follicles from ultrasound images using image processing methods, application of a mathematical stereology to quantify the attributes of the segmented follicles, and pattern recognition techniques to classify the feature vector obtained from the previous steps into one of the two categories: PCO present or PCO absent.

The novelty of our method lies in the amalgamation of the follicle segmentation technique and the stereological methodology which results in the creation of a feature vector quantifying ovarian morphology.

The problem of detecting the overall spatial distribution of follicles in the ovary was also studied. Two different general distributions are possible: peripheral distribution and random distribution. In peripheral distribution, follicles are distributed along the periphery of the ovary, and in random distribution follicles are distributed more or less uniformly within the ovary. Normal ovaries usually present a random distribution and may exhibit one or two "dominant" follicles which, by definition, are at least 2-3mm larger than the rest. Polycystic ovaries typically exhibit a more peripheral distribution pattern, although there are notable exceptions. A method of automatically detecting the mode of follicle distribution is explained in Section 6.3.

Figure 1. Examples of normal and polycystic ovaries. The normal ovary (left) contains a dominant follicle (large dark region) and exhibits a random follicle distribution. The polycystic ovary (right) has no dominant follicle and exhibits peripheral distribution.

2. Follicle Segmentation

Ovarian follicles are roughly spherical, fluid-filled structures in which oocytes (eggs) develop. Follicles imaged in two dimensions appear as dark, roughly circular regions in ultrasound images since fluid does not reflect ultrasonographic pulses. Figure 1 shows examples of a normal and a polycystic ovary. PCO ovaries typically exhibit a larger number of smaller, possibly irregularly shaped follicles, and a peripheral distribution of follicles.

Various follicle segmentation techniques have been proposed and they can be categorized into grey-level thresholding and graph searching techniques [8, 12], region growing methods [10], texture-based methods [7, 9] , and object recognition algorithms [11].

For the automatic PCO diagnosis algorithm described herein, the follicles whose structural and geometric characteristics have to be determined were segmented using a region growing algorithm based largely on that of Potočnik and Zazula [10]. This framework was chosen because, of the known follicle segmentation algorithms, it has the highest follicle recognition rate of 78% and relatively few false positives compared to its competitors (specificity 0.71) [10]. It is also fully automatic. The algorithm operates in three phases: Identification of homogeneous regions, region growing, and follicle extraction. Since this algorithm was designed to segment follicles in normal ovaries, modifications were made in order to accomodate the different properties of polycystic ovaries. Subsections 2.1 through 2.3 detail the algoirthm; modifications from the original version [10] are noted.

2.1. Identification of Homogeneous Regions

Since follicles appear nearly black in ultrasonographic images, a homogeneous region was considered to be one in which the pixels had similar greyscale values. Identification of homogeneous regions was achieved by filtering the image with an adaptive neighborhood median filter using a threshold T_1, which was set to the mean grey level of the original image. Pixels with intensity below T_1 were filtered using an 11×11 neighborhood, and pixels with intensity above this threshold were filtered using a 5×5 neighborhood. This causes a greater amount of smoothing in darker regions which are more likely to be part of a follicle and better preserves edges in brighter regions. The smoothing step was then repeated to ensure elegant smoothing of follicle regions. An 11×11 neighborhood size was chosen because a distance of 50 - 60 pixels in an ovarian ultrasound image corresponds to approximately 10mm in reality. This size ensured that follicle regions were smoothed more thoroughly than the follicle edges or the high intensity regions. Since this step was part of the coarse estimation of follicle regions, highly accurate determination of the threshold T_1 was not paramount. The filtered image was then thresholded using a new rough threshold, T_2, which was set to the mean intensity of the smoothed image minus one standard deviation of pixel intensities in the smoothed image. Structures which were incorrectly merged using this procedure were coarsely separated using binary watershed segmentation. The above procedure resulted in some undesired homogeneous regions that were too small to be follicle regions. Such regions were removed by deleting regions whose area in pixels fell below a threshold of T_3. T_3 was set to 50, which is approximately a quarter of the area of the minimum detectable follicle size. The identified homogeneous regions are then sorted in descending order by area and are passed on to the region growing step.

2.2. Region growing

The homogeneous regions from the previous phase are intial approximations of the follicles. The approximation typically underestimates follicle area and results in regions roughly centered within the actual follicle region. It is therefore appropriate to use region growing to expand the regions to the actual follicle boundary.

Each homogeneous region was grown using an iterative process in which an individual pixel is marked as a potential candidate for merging with a homogeneous region if it satisfied two merging criteria. The first criterion (Equation 1) is based on the intensity of an individual pixel; the second criterion (Equation 3) is based on weighted gradients.

Let R_0 be the initial homogeneous region prior to being grown. Region growing proceeds iteratively. Let R_i denote the resulting region after i iterations of region growing. Let $p_0 ... p_n$ be the pixels from the outer boundary of region R_i. For each iteration, Equations 1 and 3 were evaluated for all pixels $p_0 ... p_n$ in the outer boundary. Pixels that satisfied both criteria were marked as potential candidates for merging with R_i to form R_{i+1}. The first criterion was

$$|I(p_i) - \mu(R_i)| \leq \alpha\sigma(R_i), \tag{1}$$

where, $I(p_i)$ is the intensity of pixel p_i, $\mu(R_i)$ is the mean grey-level of region R_i, and $\sigma(R_i)$ is the standard deviation of grey-levels for pixels in R_i. The scaling parameter, α, was chosen to be 1.

The second criterion used edge and texture information. Edges can be detected by computing the gradient of the image. However, since region boundaries are not well expressed in ultrasound images, a weighted gradient was used. The weighted gradient magnitude is:

$$grad(p_i) = ||\nabla I_k(p_i)||(e^{G/tex(p_i)} - 1), \tag{2}$$

where $||\nabla I_k(p_i)||$ is the gradient magnitude of I_k, $tex(p_i)$ is the ratio of the mean grey-level and one standard deviation of grey-levels in the 11×11 neighborhood about pixel p_i, and $G = 2\ln(2)$. As $tex(p_i) \rightarrow \infty$, the exponential approaches 0, and when $tex(p_i) = 1.91$, exponential quantity is 1. The value 1.91 is the average signal-to-noise ratio (SNR) in regions with ultrasound speckle. Hence, $grad(p_i)$ is small for anechogenic follicle regions in which there is no speckle and large for edges and noisy regions. The second merging criterion was given as:

$$|grad(p_i) - \mu(grad(R_i))| \leq \alpha\sigma(grad(R_i)), \tag{3}$$

where, $\mu(grad(R_i))$ is the mean weighted gradient and $\sigma(grad(R_i))$ is the standard deviation of the weighted gradient in region R_i, and α was set as 2 [10].

In the original method, the marked potential candidates were merged with the homogeneous region if at least four of their neighbors were either in R_i already, or had also been marked as potential candidates. In our method, two new criteria have been added in addition to the above two merging criteria. Marked potential candidates are merged with the homogeneous regions based on the values of two region-based scalar descriptors: Solidity and eccentricity. Solidity is the proportion of the pixels in the convex hull of a region that are also in the region, and eccentricity is the ratio of the lengths of the major and minor axes of a region. If the solidity of the merged region (original region merged with a potential candidate), is less than the original region, then the potential candidate is unmarked. Also, if eccentricity of the merged region is greater than that of the original region, or if it is greater than a threshold T_4 which is set as 0.72, then the potential candidate is unmarked. An eccentricity of 1 corresponds to a circular shape, and it was found that an

eccentricity of 0.72 gave optimum results for segmentation of follicular structures. These shape descriptors are an important addition to the merging criteria as they ensure that regions retain a compact shape during growth. The growing was halted when the final region R_n was equal to that of the previous step R_{n-1}.

2.3. Follicle extraction

The identified regions of the previous step were further analyzed in an attempt to remove those that did not correspond to an actual follicle. Identified regions with an area less than 220 pixels were removed since this corresponds to the approximate area of the smallest visible follicles (2-3mm diameter). Also, if the ratio of a follicle's area to that of the area of its bounding box was less than 0.5, it was removed from consideration. All the regions satisfying these two measures were labelled and holes inside them filled.

Regions touching the image borders are removed in [10], but were retained in our method. This was because follicles in a polycystic ovary tend to be located along its periphery and we wished to retain peripheral follicles for the follicle distribution analysis described in Section 6.3.

3. Stereology and Feature Extraction

The second phase of the polycystic ovary detection algorithm is the generation of a feature vector for the image which describes the segmented follicles. The feature vectors are used in the third phase of the algorithm (see Section 4) which classifies feature vectors as arising from either polycystic or normal ovaries.

Features were derived using a mathematical methodology called stereology [13], originally developed to understand the 3D geologic composition of the earth from core samples. Stereology is now routinely used in histology (study of tissues or cells using a microscope) [5] to infer 3D structure from small samples or biopsies. In stereology, two-dimensional images are viewed as projections of three-dimensional objects. Stereology relates three-dimensional parameters of structures to two-dimensional measurements that are obtained from 2D slices through the structures [13]. A variety of geometric attributes of follicles can be calculated using stereology, such as the follicle count, distribution of follicles within the ovary, and follicle size.

Stereology defines a *structure* as the space containing the components of interest. *Phase* and *particles* make up the components of a structure, where the particles are discrete elements and a phase is the aggregate of all particles of the same kind. The fundamental quantitative descriptor of the structural entities is the *density* of various components (for our application, follicles, blood vessels, corpus luteum) within the structure (the ovary). The component of interest in the detection of PCO is the follicle. The basic quantities that describe these components are their volume, surface area, count, and diameter. Hence, the quantitative properties of the ovary can be described by volume density, surface density, numerical density, and mean follicle diameter. *Volume density* is defined as volume of the phase within the unit volume of the structure, *surface density* is surface area of the phase within the unit volume of the structure, and *numerical density* is the number of follicles in the structure [13].

Table 1. Example feature vectors extracted from polycystic ovaries.

SD	VD	Profiles	meanD	maxD
0.031009	0.13876	15	23.288	46.07
0.026463	0.12936	21	21.296	55.497
0.029709	0.14693	12	25.452	41.796

In histological applications, volume density measurements are made using techniques such as the principle of Delesse, linear integration, and point counting [13]. Methods such as intersection of test lines in space, intersection with profile boundaries and the buffon principle are used for surface density measurements [13]. The methods used in histological applications can be replaced by computer vision/image processing techniques on the digitized ultrasound images.

Volume density of a component was calculated as the ratio of the sum of follicle profile areas to the sum of the section area, where the area is calculated using a stereological method called the point counting method. The point counting method performed on histological slices was replaced with area estimation using Matlab R2006a on the digitized ultrasound images. Surface density was estimated as $(4/\pi)\times$boundary density [13], where boundary density is the ratio of the length of follicle boundary to the section area. Length of the follicle boundary was calculated by counting the number of pixels that make up the boundary of the follicle region. Numerical density (number of follicle regions per unit area) was obtained by counting the number of follicles in the ultrasound image. Follicle size was calculated by computing the average diameter of all the follicle regions in the given image. This was achieved by finding the equivalent diameter of a circle with the same area as a follicle region and was computed as $\sqrt{4 * Area/\pi}$.

The following five stereological features were used to construct a feature vector describing the follicles segmented from the input image: surface density (SD), volume density (VD), number of follicle regions per image (Profiles), mean follicle diameter (meanD), and maximum follicle di-

Table 2. Example feature vectors extracted from normal ovaries.

SD	VD	Profiles	meanD	maxD
0.018479	0.13671	6	33.765	72.234
0.013556	0.096059	3	40.547	49.392
0.025535	0.15763	4	35.038	66.593

ameter (maxD). These features were chosen because they characterize the most important aspects of follicles within the ovary and abnormalities in follicle morphology are a primary indication of PCO. Table 1 contains example feature vectors extracted from images of polycystic ovaries obtained after the segmentation and feature extraction phases. Table 2 contains example feature vectors extracted from normal ovaries.

4. Classification

The feature vectors obtained from the previous step were classified into one of the following two classes: (i) PCO present or, (ii) PCO absent. The linear discriminant classifier, k-nearest neighbor classifier (KNN), and Support Vector Machine (SVM) classifier were evaluated for their ability to correctly determine classes of feature vectors. The classification rates of each classifier were determined using the k-fold cross validation methodology. Fundamentals of pattern classification and details on the above classifiers are referenced in [4]. A review of validation techniques for medical image analysis, including k-fold cross validation, can be found in [3].

4.1. Classifiers

The classifiers that we used in this work are the linear discriminant classifier (the `classify` function in Matlab R2006a), k-nearest neighbor classifier (`knnclassify` in Matlab R2006a)), and the Support Vector Machine Classifier (function `svmclassify` in Matlab R2006a). The two former implementations are part of Matlab's statistics toolbox, and the latter is from Matlab's Bioinformatics toolbox.

4.2. Validation

Once a classifier model has been learned from the training patterns, its ability to classify new patterns can be assessed using cross validation techniques. This is accomplished by using only part of the available patterns for training. The remaining "test" data are used to test the performance of the learned model. Common types of cross validation methods are the holdout method, and the k-fold cross validation method [3].

In the holdout (or half-and-half) method, the data set is randomly split into a training set and a testing set. A model is learned from the training set and the validity of the model checked by determining the classification accuracy of the model using the testing set. Model accuracy is dependent on the particular split of the data. The disadvantage of the holdout method can be avoided by using the k-fold cross validation technique. In this method, the data set is divided into k folds, out of which $k-1$ folds are used as the training set, and the remaining fold is used as the testing set. The holdout method is performed k times, each using a different fold as the testing set thus eliminating the dependence on the division of the data points among the training and the testing sets. The classification accuracy is averaged over the k trials, and the variance decreases as k increases.

5. Experimental Setup

A total of 70 ovarian ultrasound images were obtained from the Women's Health Imaging Research Laboratory (WHIRL) in Saskatoon, Canada. This set contained images of both polycystic ovaries ($n = 33$) and normal ovaries ($n = 37$). A feature vector was extracted from each image using phases 1 and 2 of our algorithm (see Sections 2 and 3).

The 70 feature vectors were randomly divided into $k = 10$ folds for evaluation using the k-fold cross validation technique. Comparison of the classifier performance and classification results using this cross validation method are presented in Section 6.2.

6. Results and Discussion

Since the segmentation algorithm used to identify follicles was a modified version of the algorithm in [10], a segmentation validation was performed to verify the accuracy of the modified algorithm. This process is detailed in Section 6.1. Section 6.2 presents the classification results for the three classifiers. Section 6.3 discusses results of a small experiment to classify the spatial distribution of follicles as random or peripheral (as described in Section 1) automatically using the linear discriminant and KNN classifiers.

6.1. Segmentation validation

The accuracy of follicle segmentation was measured by comparing manual segmentations generated by a human expert (ground truth) with our automatic segmentation results using the following similarity metrics: Hausdorff distance (HDist), mean distance (MDist), and DICE coefficient (DICE).

Table 3. Mean validation metrics and their standard deviations for the automatic segmentation algorithm over all images in the data set.

	Mean	σ
HDist	7.31mm	3.89mm
MDist	1.27mm	0.88mm
DICE	0.60	0.16
RR	83.1%	21.5%
MR	31.1%	23.3%

Hausdorff distance measures the largest minimum distance between a point on the automatically segmented region and all the points on the expert marked region and vice versa. It characterizes the maximum deviation of the segmentatation boundary from the ground truth. Let A denote the set of points in the automatically generated boundary and let G denote the set of points in the ground truth boundary. The Hausdorff distance is then given by

$$d(p, B) = \min_{b \in B} ||b - p||, \quad (4)$$

$$\text{HDist} = \max \left[\max_{a \in A}[d(a, G)], \max_{g \in G}[d(g, A)] \right], \quad (5)$$

where $d(p, B)$ is the minimum Euclidian distance between a point p and the boundary B. Thus, a smaller Hausdorff distance indicates a more accurate segmentation. The average Hausdorff distance over all images was 7.31mm (standard deviation $\sigma = 3.89$mm).

Mean distance (Mdist) is the average minimum distance between a point on A and the boundary G. It is formally defined as

$$\text{MDist} = \frac{1}{2} \left[\frac{1}{n_A} \sum_{a \in A} d(a, G) + \frac{1}{n_G} \sum_{g \in G} d(g, A)) \right], \quad (6)$$

where n_A is the total number of pixels in the segmented region and n_G is the total number of pixels in the expert traced region. The average MDist over all images was 1.27mm ($\sigma = 0.88$mm). This result, combined with the average Hausdorff distance above, indicates there were few long segments of follicle boundaries that exhibited significant deviation from the ground truth.

The DICE coefficient is defined as twice the ratio of the area of intersection of the automatically segmented region and the expertly segmented region to the total area of the automatically and expertly segmented regions,

$$DICE = \frac{2 \cdot |(\alpha \cap \gamma)|}{|\alpha| + |\gamma|}, \quad (7)$$

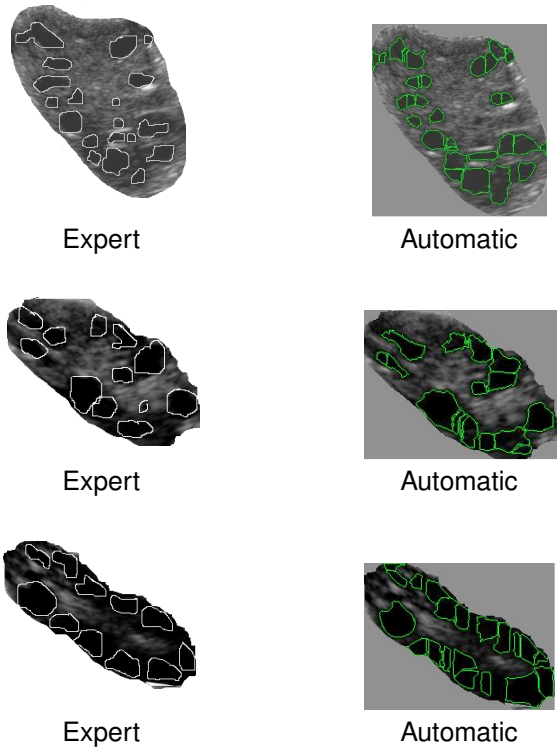

Figure 2. Examples of automatic segmentation results for polycystic ovaries.

where α is the set of pixels in the automatically segmented region and γ is the set of pixels in the expertly segmented region. This metric captures the amount ovarlap between the two regions; it is 0.0 if the regions are disjoint and 1.0 if they are identical. For our images, average DICE coefficient for a segmented image was 0.60 ($\sigma = 0.16$) which shows a significant percentage of overlap between the follicle regions of these two sets of images.

Figure 2 depicts example automatic segmentations. The output of automatic segmentation is similar to the manually traced expert segmentation. Table 3 summarizes the segmentation validation results. A common mode of error is the division of single follicles into multiple regions.

The recognition rate (RR) and misidentification rate (MR) were also computed in order to characterize the ability of the segmentation algorithm to locate follicle regions regardless of the boundary accuracy. Recognition rate is the ratio of the number of actual follicles in the ground truth that correspond to some automatically segmented follicle to the actual number of expert traced follicles. The average recognition rate for all our images was 83.1% ($\sigma = 21.5\%$). The original method of [10] achieved a RR of 78% ($\sigma = 21\%$).

The misidentification rate is defined as the proportion of

Table 4. Performance results for the linear discriminant (LDC), k-nearest neghbor (KNN) and support vector machine (SVM) PCO classifiers.

Classifier	CorrectRate	Sensitivity	Specificity
LDC	92.86%	0.9091	0.9459
KNN	91.43%	0.9394	0.8919
SVM	91.43%	0.9091	0.9189

the total number of segmented regions that did not correspond to an expert-identified follicle (false positives). The average MR for our images was 31.1% ($\sigma = 23.3\%$). The original algorithm had an MR of 29% ($\sigma = 25\%$) [10].

6.2. Classification results

The accuracy of the classification of the feature vectors as polycystic or normal by the three classifiers as determined by the k-fold cross validation method is given in Table 4. Sensitivity is the proportion of polycystic ovaries for which there was a positive test. The linear discriminant classifier exhibited a sensitivity of 0.909 and the KNN classifier exhibited a sensitivity of 0.9394.

Specificity is defined as the proportion of disease-free ovaries for which there was a negative test. Consistently high rates of specificity were exhibited by all three classfiers; the linear discriminant classifier achieved the best sensitivity of 0.9459. Thus, the rates of both false positives and false negatives are low for all three classifiers.

The column labeled CorrectRate in Table 4 is the overall classification rate and indicates the percentage of ovaries for which a correct classification was made. The linear discriminant classifier produced the highest classification rate of 92.86%, while the KNN and SVM classifiers made correct decisions 91.43% of the time.

6.3. Follicle distribution

An algorithm was devised to automatically classify the spatial distribution of follicles in an ovary as either random or peripheral using the segmented follicle regions as input, as described in Section 1.

For each ovary, the centroid of each segmented region was found. The mean centroid (centroid of the centroids), denoted m_c, was also computed. The mean and standard deviation of the distances between each region centroid and m_c were used as features for this classification. Higher order moments such as skew and kurtosis of region centroid distances to m_c were not found to improve the classification results presented below. Of the 33 polycystic ovaries in our data set, only the 23 which exhibited either peripheral or random distributions were used; 12 had random distributions, and 11 had peripheral distributions. In Section 1 it was mentioned that there are other possible types of distributions; the remaining 10 ovaries in the data set were of these types and were not used. Features were extracted and classification was performed using the linear discriminant classifier and the KNN classifier.

Table 5. Results for classification of follicle distribution as peripheral or random.

Classifier	CorrectRate	Sensitivity	Specificity
LDC	81.25%	0.9167	0.5000
KNN	62.50%	0.7500	0.2500

Table 5 shows the results of the follicle distribution classification using a linear discriminant classifier, and k-nearest neighbor classfier as determined by 10-fold cross validation. A correct classification rate of 81.25% was obtained using the linear discriminant classifier.

7. Discussion and Conclusion

A method for automatically discriminating between normal and polycystic ovarian follicle morphology was presented in three phases: follicle segmentation using a region growing algorithm, quantification of the attributes of the segmented follicles using stereology, and classification of the resulting feature vectors as either normal or polycystic. The linear discriminant classifier outperformed KNN, and SVM classifiers, but all the three classifiers performed well. Classification rates were 92.86%, 91.43%, 91.43% respectively.

A classifier for the distribution of the follicles inside polycystic ovaries was developed. It used two features, the mean and standard deviations of the distances of the centroids of individual follicles to the mean centroid. A linear discriminant classifier had a classification rate of 81.25% as determined by 10-fold cross validation.

A criticism of this work might be that the follicle segmentation algorithm which was used, despite having the best known follicle recognition rate, lacks robustness. Our group is currently investigating more robust follicle segmentation algorithms.

Our results offer the promise of deploying a robust automated screening system for PCO which will improve the rapidity and accuracy of diagnosis of Polycystic Ovarian Syn-

drome and facilitate reduced danger from the severe complications that can arise from an undiagnosed condition.

8 Acknowledgements

This research was supported by grants from the Natural Sciences and Engineering Research Council of Canada, the Canadian Institutes of Health Research, and the University of Saskatchewan. The authors would like to thank the anonymous reviewers for their helpful and constructive feedback.

References

[1] A. U. Bako, S. Morad, and W. A. Atiomo. Polycystic ovary syndrome: An overview. *Reviews in Gynaecological practice*, 5(2):115–122, 2005.

[2] A. H. Balen, J. S. E. Laven, S.-L. Tan, and D. Dewailly. Ultrasound assessment of the polycystic ovary: international consensus definitions. *Human Reproduction Update*, 9(6):505–514, 2003.

[3] K. W. Bowyer. Validation of medical image analysis techniques. In M. Sonka and J. M. Fitzpatrick, editors, *Handbook of Medical Imaging*, volume 2, chapter 10. SPIE-International Society for Optical Engineering, 2000.

[4] R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern Classification*. Wiley-Interscience, second edition, 2000.

[5] D. W. Fawcett. *Bloom and Fawcett: A textbook of Histology*. Hodder Arnold, 12th edition, 1994.

[6] S. Franks. Medical progress: Polycystic ovary syndrome. *New England Journal of Medicine*, 333(13):853–861, 1995.

[7] C.-F. Jiang and M.-L. Chen. Segmentation of ultrasonic ovarian images by texture features. In *Proceedings of the 20th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, volume 2, pages 850–853, 1998.

[8] A. Krivanek and M. Sonka. Ovarian ultrasound image analysis: Follicle segmentation. *IEEE transactions on medical imaging*, 17(6):935–944, 1998.

[9] R. Muzzolini, Y.-H. Yang, and R. Pierson. Multiresolution texture segmentation with application to diagnostic ultrasound images. *IEEE transactions on medical imaging*, 12(1):108–123, 1993.

[10] B. Potočnik and D. Zazula. Automated analysis of a sequence of ovarian ultrasound images. Part I: segmentation of single 2D images. *Image and Vision Computing*, 20(3):217–225, 2002.

[11] B. Potočnik and D. Zazula. Automated analysis of a sequence of ovarian ultrasound images. Part II: prediction-based object recognition from a sequence of images. *Image and Vision Computing*, 20(3):227–235, 2002.

[12] G. E. Sarty, W. Liang, M. Sonka, and R. A. Pierson. Semi-automated segmentation of ovarian follicular ultrasound images using a knowledge-based algorithm. *Ultrasound in Medicine and Biology*, 24(1):27–42, 1998.

[13] E. Weibel. *Stereological Methods. Vol. 1: Practical Methods for Biological Morphometry*. Academic Press, London, U.K., 1979.