Using Artificial Neural Networks and Feature Saliency Techniques for Improved Iris Segmentation

Randy P. Broussard, Lauren R. Kennell, David L. Soldan, and Robert W. Ives

***Abstract*—One of the basic challenges to robust iris recognition is iris segmentation. This paper proposes the use of a feature saliency algorithm and an artificial neural network to perform iris segmentation. Many current Iris segmentation approaches assume a circular shape for the iris boundary if the iris is directly facing the camera. Occlusion by the eyelid can cause the visible boundary to have an irregular shape. In our approach an artificial neural network is used to statistically classify each pixel of an iris image with no assumption of circularity. First, a feed-forward feature saliency technique is performed to determine which combination of features contains the greatest discriminatory information. Image brightness, local moments, local orientated energy measurements and relative pixel location are evaluated for saliency. Next, the set of salient features is used as the input to a multi-layer perceptron feed-forward artificial neural network trained for classification. Testing showed 96.46 percent accuracy in determining which pixels in an image of the eye were iris pixels. For occluded images, the iris masks created by the neural network were consistently more accurate than the truth mask created using the circular iris boundary assumption. Post-processing to retain the largest contiguous piece in the iris mask increased the accuracy to 98.2 percent.**

I. INTRODUCTION

The iris is the colored portion of the eye surrounding the pupil. Its textured pattern is stable over time, distinct from person to person, and distinct from left eye to right eye of the same person [1]. Human identification based on patterns within the iris is potentially one of the most accurate methods of biometric identification [1]. Iris identification still poses several challenges that must be overcome before a truly robust identification system can be fielded. One of the basic challenges is iris segmentation. This paper proposes the use of a feature saliency algorithm and an artificial neural network to perform iris segmentation.

Many current iris segmentation approaches assume a circular shape for the pupil and limbic boundaries that border the iris. Often the pupil and limbic boundaries are slightly elliptical and may not share the same centroid [1]. Occlusion by the eyelid and non-orthogonal viewing angles can also cause the visible boundary of the iris to have an irregular shape. Each of these issues can cause a segmentation system, based on a circular assumption, to produce inaccurate results.

Manuscript received January 31, 2007.

Randy P. Broussard is with the Weapons and Systems Engineering Department, United States Naval Academy, Annapolis, MD 21402 USA (e-mail: broussar@ usna.edu).

Lauren R. Kennell, David L. Soldan and Robert W. Ives are with the Electrical Engineering Department, United States Naval Academy, Annapolis, MD 21402 USA.

II. METHODOLOGY

To perform iris segmentation without the assumption of circularity, an artificial neural network is used to statistically classify each pixel of an iris image as one of three classes: iris, pupil, or other. First, a feed-forward feature saliency technique (described in section IV) is performed to determine which combination of features contains the most discriminatory information. Local features, including local image statistics (mean, standard deviation, skewness, and kurtosis) and local orientated energy measurements, are evaluated for saliency. Various neighborhood sizes are used to compute each local measurement. Feature saliency is performed using all local features, with each feature computed at all neighborhood sizes. The saliency process identifies a subset of features that jointly contains the most discriminatory information for the problem at hand. Next, this subset of features is used as the input to a multi-layer perceptron feed-forward artificial neural network trained for classification. The number of nodes within the hidden layer of the neural network is kept low to deter memorization and to decrease computational run time. The number of hidden nodes was varied for thoroughness, but six hidden nodes provided acceptable accuracy.

A subset of the University of Bath iris database [2] was used to train the neural network. The database contains 2000 grayscale near infrared images. The images consist of twenty pictures of each eye from 50 individuals. Testing was performed on a different subset of the database. Ground truth was generated using an in-house algorithm, which used local statistics and a circular iris boundary assumption [3]. The truth masks for images containing occluded and/or non-circular irises were corrected by hand to generate the training and test sets. The neural network was trained using the error back-propagation algorithm. Initial results indicated that the local brightness-based features did not hold the discriminatory information necessary to properly segment the images within the test set. Pixel location features, such as pixel distance and angle from the pupil center, were added. The addition of these features allowed the neural network to statistically determine the most likely locations of the iris boundaries and focus its

decision making on those pixels.

One of the issues confronted when setting up the classification problem was determining the number of classes to classify. Initially, two classes were used. The classifier was asked to classify each pixel as iris or non-iris. This meant the dark pupil and the lighter portions of the image were included in the same class. Since the iris contains shades of gray somewhere between the dark pupil and light sclera (white portion of the eye), the non-iris class had a bimodal distribution. This distribution violated the Gaussian Probability Density Function (PDF) assumption made in the Bayesian classifier. The feature selection process thus performed sub-optimally and initial classifier results were less accurate than desired. To address this bimodal distribution issue, a 3-class problem was evaluated. The pupil was designated as the third class. This allowed the grayscale values within the two non-iris classes to form a more Gaussian distribution. Initial classification results showed a large accuracy improvement using the 3-class problem. In many cases, the pupil can be segmented more easily than the iris. Adding the assumption that the pupil center and radius are known, the pupil pixels can be removed from the training and test sets. Removing the pupil from consideration returned the problem to a 2-class problem. This modified 2-class problem produced the highest classification accuracy in initial testing and was used in the remainder of this research.

III. FEATURE SET

A total of 264 local features, including local image statistics (mean, standard deviation, skewness, and kurtosis) and local orientated energy measurements were evaluated for feature saliency. Each local feature was computed using multiple neighborhood sizes and shapes. Each feature listed in table 1 was computed using a neighborhood size of 3, 7, 11, 15, 19, 21, 25 and 29 pixels along the major axis. Five shaped neighborhoods were used to compute the local statistics. A square neighborhood centered about each pixel being measured was the first to be tested. Next a rectangular neighborhood whose major axis is parallel to the vector from the pupil center to the pixel being measured was added. This forms a radial local statistic. A rectangular neighborhood whose major axis is perpendicular to the radial local statistic was also added. For completeness, a vertical and a horizontal rectangular neighborhood were also evaluated. Figure 1 shows the five neighborhood shapes. Each local statistic computed at each neighborhood size/shape produced a single feature for a total of 160 local statistic features.

Local orientated energy measurements were computed using oriented band pass filters (similar to Gabor filters). Four orientations of the filters were computed (0, 45, 90 and 135 degrees) and the maximum response was selected as magnitude and direction of the local pixel energy as detected by the particular orientated filter. Each oriented filter is designed to have an affinity for objects of a particular size. Each filter's size affinity is determined by the size of the

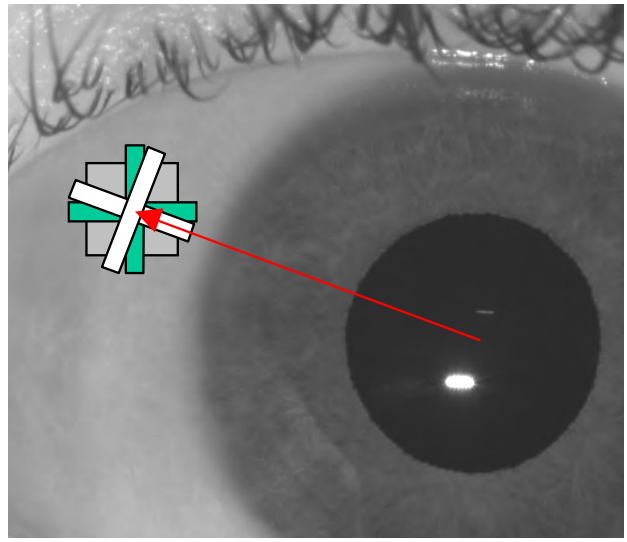

Fig. 1. Five neighborhood shapes used to compute local statistics

positive region within its filter kernel. We will refer to this positive region as the filter's receptive region size. Various receptive region sizes were computed and each size produced two features (magnitude and direction).

Measurements
Square Local mean
Square Local standard deviation
Square Local skewness
Square Local kurtosis
Radial Local mean
Radial Local standard deviation
Radial Local skewness
Radial Local kurtosis
Perpendicular Local mean
Perpendicular Local standard deviation
Perpendicular Local skewness
Perpendicular Local kurtosis
Vertical Local mean
Vertical Local standard deviation
Vertical Local skewness
Vertical Local kurtosis
Horizontal Local mean
Horizontal Local standard deviation
Horizontal Local skewness
Horizontal Local kurtosis
Local orientated energy magnitude
Local orientated direction
Radial Mean - Perpendicular Mean
Radial StdDev minus Perpendicular StdDev
Radial Skewness minus Perpendicular Skewness
Radial Kurtosis minus Perpendicular Kurtosis
Radial Mean minus Square Mean
Radial StdDev minus Square StdDev
Radial Skewness minus Square Skewness
Radial Kurtosis minus Square Kurtosis

Radial Mean divided by Iris Mean
Radial StdDev divided by Iris StdDev
Radial Skewness divided by Iris Skewness
Radial Kurtosis divided by Iris Kurtosis

Table 1. Measurements and neighborhood sizes used to create a feature set.

Several ratios and differences of the various features were also computed to form additional features, which have a degree of brightness invariance. Table 1 shows all measurements used to create the feature set.

Initial neural network test runs indicated that the 264 local features (based on brightness only) did not contain the discriminatory information required to accurately segment the iris. Classification accuracy was approximately 86 percent when compared to the truth masks. Visually, many of the segmented images were poor quality. Pixel location features were added to the feature set and the classification accuracy improved by more than 10 percent. The pixel location features consisted of the following measurements for each evaluated pixel.

Euclidean distance from pupil center
Euclidean distance from pupil boundary
Angle from pupil center
Vertical distance from pupil center
Horizontal distance from pupil center

Table 2. Pixel location features added to the 264 local features

The majority of the previous incorrect classifications were found to be located on the iris-limbic boundary. The 10 percent improvement gained from the pixel location features translated into a dramatic improvement visually.

Processing every pixel within each training set image proved to be impractical. A standard image from the University of Bath database has a resolution of 1280 by 960 pixels. This amounts to 1.2 million pixels per image or 24.6 million observations within the entire training set. To increase the computational tractability of the problem, each image was cropped to a 600 by 600 pixel region surrounding the pupil center. Selecting only every 5th pixel for processing additionally reduced the computation time. In total, the size of the training set was reduced to 1.44 million observations taken from 20 images which contained a mix of right and left eyes.

IV. FEATURE SALIENCY

Feature saliency was performed to determine a subset of features that jointly contain the greatest discriminatory information that will distinguish the iris from the rest of the eye. To find this feature subset, a feed-forward, Bayesian classifier based feature saliency technique was used. The feed-forward technique is a sub-optimal approach, but is computationally practical and delivers excellent results in practice. Using this approach, hundreds of thousands of observations, each containing hundreds of features, can be evaluated within an hour on a desktop PC. The Bayesian classifier equation used within the algorithm utilizes a simple Euclidean distance discriminate and a Gaussian PDF. The jth discriminate is computed over n features using equation (1) where μ_i and σ_i are the mean and standard deviation for the ith feature.. This Gaussian assumption worked well since all of the local statistical computations used a Gaussian PDF assumption.

$$D_j = \sqrt{\sum_n \frac{(x_i - \mu_i)^2}{\sigma_i^2}} \qquad (1)$$

To begin, the feed-forward technique evaluates the discriminatory power of each feature individually and keeps the single feature that demonstrates the greatest discriminatory power. Next, each feature is individually used in conjunction with the first feature and the pair that contains the greatest discriminatory power is selected. This process, of adding individual features to the group of previously selected features, is continued until no discriminatory power is gained by adding an additional feature to the previously selected features. Because this is a local search methodology, the results are considered sub-optimal. The resulting feature subset may be optimal, but is not guaranteed to be. Because of the large number of pixels in the training image set and the large number of features evaluated, a single global search could take weeks or months to complete. For this research problem, using a local search allowed the compute-evaluate-improve cycle to occur on a daily basis.

The discriminatory power of each subset of features was measured by using a Bayesian classifier. Each feature subset was evaluated on the entire training set and the subset producing the highest classification accuracy was deemed to contain the greatest discriminatory power of the subsets evaluated. Since a Bayesian classifier is a statistics based classifier, the selected subset is statistically the most salient feature subset of the subsets evaluated. A Bayesian classifier was selected because of its computational speed. Using an artificial neural network as a feature selection classifier was tested, but extended the processing time by several orders of magnitude. In our probative initial research, the slight classification accuracy gains provided by the neural network did not justify the immense time loss it introduced. A neural network based approach may be practical if executed on a high performance computing system, or when used for fine-tuning a production iris recognition system.

Features: 4, Percent correct = 90.19
Features: 4 136, Percent correct = 92.30
Features: 4 136 153, Percent correct = 93.47
Features: 4 136 153 201, Percent correct = 93.63
Features: 4 136 153 201 3, Percent correct = 94.03
Features: 4 136 153 201 3 7, Percent correct = 94.37
Features: 4 136 153 201 3 7 152, Percent correct = 94.55
Features: 4 136 153 201 3 7 152 5, Percent correct = 94.65
Features: 4 136 153 201 3 7 152 5 121, Percent correct = 94.78
Features: 4 136 153 201 3 7 152 5 121 22, Percent correct = 94.84

Fig. 2. An example of the feature saliency algorithm selecting features

Figure 2 shows an example of the feature saliency algorithm evaluating subsets of the 269 features for classification accuracy. The feature numbers correspond to the following features:

Feature 4: Euclidean distance from Pupil Center
Feature 136: Square Mean (region size = 25)
Feature 153: Perpendicular StdDev (region size = 29)
Feature 201: Radial StdDev - Square StdDev (region size = 11)
Feature 3: Euclidean distance from Pupil Boundary
Feature 7: Y Distance from Pupil Center
Feature 152: Perpendicular Mean (region size = 29)
Feature 5: Absolute value of angle from Pupil Center (0-180)
Feature 121: Vertical StdDev (region size = 21)
Feature 22: Vertical Skewness (region size = 3)

Fig. 3. Ten most salient features selected by the feature selection algorithm.

Using a Bayesian classifier and the features listed in Figure 3, the pixels in the 20-image test set could be classified with 94.84 percent accuracy when compared to the truth masks.

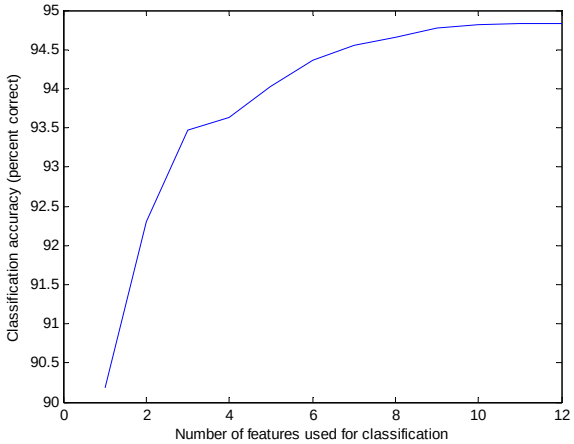

Fig. 4. Classification accuracy vs. number of features used for classification

From Figure 4 it can be seen that each additional feature provided diminishing incremental classification improvement. To speed processing, the top five or six features could be used with little degradation in classification accuracy. Typically the last few features in the salient feature subset provide classification of statistically outlying data points. They typically do not generalize well and add little accuracy when performing classification on the test set.

V. NEURAL NETWORK AND POST PROCESSING

A multi-layer perceptron (MLP) feed-forward artificial neural network was used for classification of the test set. An artificial neural network can offer a performance advantage over a Bayesian classifier if the feature distributions do not match the Bayesian PDF assumption [10]. The Bayesian classifier used in the feature selection process assumed the data to have a Gaussian distribution. The error back-propagation training algorithm used to train the MLP traverses an error surface to minimize classification error [9]. In this process, no PDF assumption is made. If the training and test features do not have a true Gaussian distribution, the artificial neural network should perform better [10].

The MLP was trained on the 20 image training data set consisting of 1.44 million observations. The 10 features, selected in the feature saliency process (Fig 3), were used as input to the MLP. To remove any effects caused by the random neural network internal weight settings at the start of training, the neural network was trained 10 times and the weights from the best training run were retained for use in the testing phase.

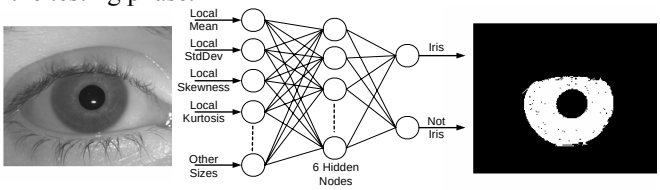

Fig. 5. Using an Artificial Neural Network to segment the iris.

Figure 5 demonstrates how the neural network was configured to segment the iris. For each pixel in the iris image, the selected features were computed and presented to the neural network for classification. To decrease computational run time, every second pixel was processed with little noticeable degradation in mask quality.

VI. RESULTS

The test set contained 40 images composed of left and right eyes of 13 individuals. No training images were contained in the test set. The artificial neural network achieved 97.32 percent classification accuracy on the test set. The neural network provided a 2.49 percent improvement over the Bayesian classifier results. This improvement indicates that some of the selected features do not have a true Gaussian distribution. The number of nodes within the hidden layer of the neural network was kept low to deter memorization and to decrease computational run time. Six hidden nodes were sufficient to achieve the 97.32 percent classification accuracy.

Post processing was used to increase the accuracy and quality of the iris mask. A number of the incorrectly

classified pixels were not contiguously connected to the largest grouping of pixels in the mask. Locating and retaining only the largest group of contiguous pixels increased the iris mask accuracy to 98.2 percent. Additional accuracy could have been pursued by using basic morphological operations, such as open or close, but no attempt was performed.

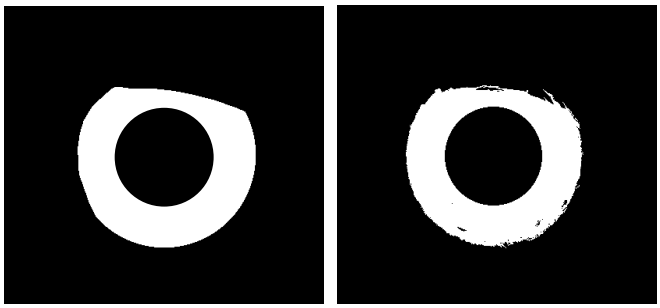

Fig. 6. Examples of a truth mask (left) and a neural network produced iris mask (right). The original image for this eye is shown in figure 7. The large pupil opening in the iris mask is due to the removal of the pupil area from the training and test sets.

In many images, the location of the iris boundaries was subjective due to the gradual color transition of the limbic boundary. This subjective nature of this boundary called into question the accuracy of the truth masks used to measure classification correctness. For many images, visual inspection often provided greater insight to correctness than numerical accuracy. Figure 6 shows a manually modified truth mask and the corresponding neural network produced iris mask. No attempt was made to locate the pupil boundary. The large pupil opening in the iris mask was created when the pupil area was removed from the training and test sets to create a two-class classifier problem.

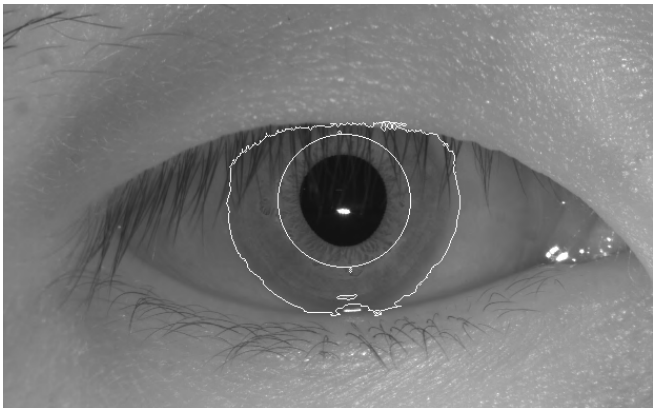

Fig. 7. The perimeter of a neural network produced iris mask overlaid onto the original iris image. Note this eye has a second fold beneath the outer, upper eyelid, which was correctly segmented by the neural network.

Figure 7 shows the iris mask perimeter overlaid onto the original image of the eye. This is an image of an Asian eye, which is often more difficult to segment than eyes of other races. Eyelashes also occlude this iris. This was an expected result, since no attempt was made to remove eyelashes from the truth masks.

Figure 8 shows two examples of iris segmentation. The upper eye demonstrates performance on a light colored iris and the lower demonstrates dark iris performance. The numerical accuracy for each image is 97.9 and 97.6 percent respectively, as compared to the truth mask. The truth mask

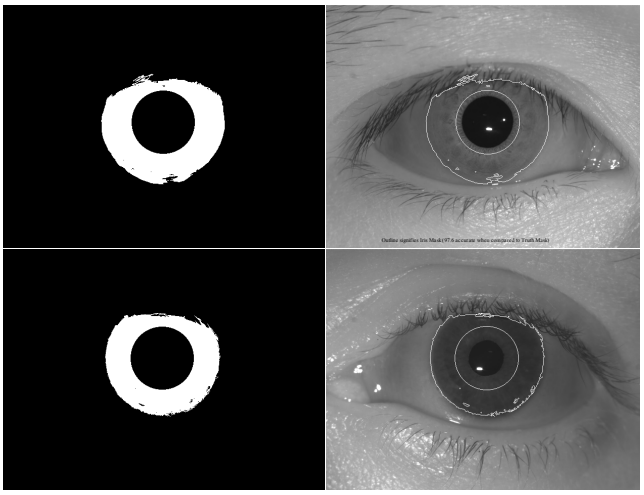

Fig. 8. Example of typical accuracy achieved on the 40-image test set. The neural network created iris mask is on the left. On the right is the outline of the iris mask overlaid onto the iris image. Both a dark colored and a light colored iris are shown for comparison.

and neural network produced iris mask are presented for comparison. The images in figure 8 represent typical results from the test set. Space precludes including more of the test set images in this paper. Most iris masks for non-occluded images had accuracies slightly above the 98.2 percent average and most occluded images had accuracies slightly below the average. The test set was limited to 40 images due to the time required to create non-circular truth masks by hand. Sample images from the remainder of the database gave visually similar results.

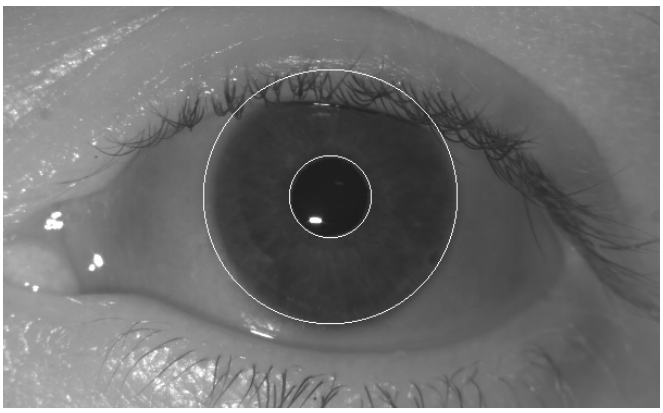

Fig. 9. Example of circular boundary assumption segmentation technique use in Kennell, et al.

For comparison, Figure 9 presents segmentation results from an algorithm described in [3]. Segmentation was performed on the same image as shown in Figure 8. The algorithm assumes circular boundaries for the pupil and

limbic boundaries. Note the elliptical shape of the limbic boundary. Modifications to this method include straight lines to approximate the upper and lower eyelids [11] or curved lines [1]. Neither of these modifications addresses elliptical boundaries caused by off-axis viewing angles.

VII. CONCLUSION

Feature selection can be used to increase performance and decrease processing time. The neural network improved segmentation accuracy over the circular assumption based segmentation algorithm. The neural network produced higher accuracy numbers than the Bayesian statistical classifier using the same features. On occluded images, the iris masks created by the neural network were consistently more accurate than the iris mask created using the circular iris boundary assumption. On non-occluded images, the neural network classified irregular shaped boundaries with accuracies in excess of 98.2 percent. The 1.8 percent error achieved on the test set was predominately in the occluded images. The majority of the errors occurred at locations where a large grouping of eyelashes existed on the eyelid and iris boundary. In some images, the neural network exhibited problems distinguishing between eyelash and iris. This result was expected, since no effort was made to remove eyelashes from the training set. Visually, the neural network segmentation results approached the perceived accuracy of the manually created truth masks.

REFERENCES

[1] J. Daugman, "High Confidence Visual Recognition of Persons by a Test of Statistical Independence", Pattern Analysis and Machine Intelligence, IEEE Transaction on, Vol. 15, No. 11, pp. 1148-1161, 1993.

[2] Monro, D. M., Rakshit, S., and Zhang, D, University of Bath, U.K. Iris Image Database, http://www.bath.ac.uk/elec-eng/pages/sipg/irisweb.

[3] Lauren R. Kennell, Robert W. Ives, Ruth M. Gaunt, "Binary morphology and local statistics applied to iris segmentation for recognition," *Proceedings of the 13th Annual International Conference on Image Processing*, Oct. 2006, in press.

[4] John Daugman, "Biometric personal identification system based on iris analysis," United States Patent, Patent Number 5,291,560, 1994.

[5] R.P. Wildes, J.C. Asmuth, G.L. Green, S.C. Hsu, R.J. Kolczynski, J.R. Matey, and S.E. McBride, "A Machine Vision System for Iris Recognition", Mach. Vision Application, Vol. 9, pp.1-8, 1996.

[6] W.W. Boles and B. Boashash, "A Human Identification Technique Using Images of the Iris and Wavelet Transform", *IEEE Transactions on Signal Processing*, Vol. 46, No. 4, pp. 1185-1188, 1998.

[7] Y. Zhu, T. Tan, and Y. Wang, "Biometric Personal Identification Based on Iris Patterns", *15th International Conference on Pattern Recognition*, Vol. 2, pp. 801–804, 2000.

[8] R.C. Gonzales and R.E. Woods, *Digital Image Processing*, Second Edition, Prentice Hall, Upper Saddle River, New Jersey, 2001.

[9] Cybenko, G. "Approximation by Superpositions of Sigmodial Functions", Mathematics of Control, Signals, and Systems, 1982

[10] Ruck, D. W., et al., "The multiplayer perceptron as an approximation to a Bayes optimal discriminant function", IEEE Transactions on Neural Networks, Vol 4, pp. 296-298, 1990

[11] Masek L., Kovesi P., MATLAB Source Code for a Biometric Identification System Based on Iris Patterns. The School of Computer Science and Software Engineering, The University of Western Australia. 2003.