Automated Classification of Contact Lens Type in Iris Images

James S. Doyle, Patrick J. Flynn, Kevin W. Bowyer
Department of Computer Science and Engineering
University of Notre Dame
{jdoyle6,flynn,kwb}@nd.edu

Abstract

Textured cosmetic lenses have long been known to present a problem for iris recognition. It was once believed that clear, soft contact lenses did not impact iris recognition accuracy. However, it has recently been shown that persons wearing clear, soft contact lenses experience an increased false non-match rate relative to persons not wearing contact lenses. Iris recognition systems need the ability to automatically determine if a person is (a) wearing no contact lens, (b) wearing a clear prescription lens, or (c), wearing a textured cosmetic lens. This work presents results of the first attempt that we are aware of to solve this three-class classification problem. Results show that it is possible to identify with high accuracy (96.5%) the images in which a textured cosmetic contact lens is present, but that correctly distinguishing between no lenses and soft lenses is a challenging problem.

1. Introduction

Flom and Safir [4] proposed the texture of the iris as a biometric modality in a 1987 patent. The first working iris biometric algorithm was developed by John Daugman [2] and is the basis for many of the commercial systems available on the market today. In [3], Daugman proposes that iris biometric systems "can tolerate a huge amount of corruption in iris images due to ... contact lenses ..." among other image artifacts.

Some commercial iris biometrics systems claim to have a method for detecting the presence of textured contact lenses, to prevent a spoofing attempt. Hard contact lenses, otherwise known as rigid gas-permeable lenses, also present some issues to iris biometrics system as they greatly alter the appearance of the iris texture. Some iris biometrics systems also detect the presence of hard contact lenses and attempt to mask out the affected region of the iris. Traditional thought is that clear, soft, non-cosmetic contact lenses do not alter the imaged iris texture enough to be of concern. As a result, we know of no iris biometrics system that attempts to detect or mitigate the effects of clear, soft, non-textured contact lenses.

Since the purpose of a prescription contact lens is to change the optical properties of the eye, it must, by definition, have some effect on the iris texture observed through it. Clear, soft, non-textured lenses are also able to move around slightly, resulting in a slightly different observed effect on the iris texture from acquisition to acquisition. Some soft lenses also have visible markings on them, which may be observed in different locations from acquisition to acquisition. Some lenses also have a noticeable boundary between the support region of the lens and the corrective region of the lens, which can alter the appearance of the iris texture.

Detection of the presence of a clear, soft, non-cosmetic contact lens is a first step to improving the usability of iris biometrics systems for contact lens wearers. One simple solution might be to change the decision threshold when a contact lens is detected such that the false non-match rate (FNMR) is identical to users who do not wear lenses. Detection is also a first step to performing any sort of image correction on images with contact lens artifacts.

This paper presents an approach to classifying an iris image into one of three categories: (1) textured contact lens, (2) soft, non-textured contact lens, and (3) no contact lens. Related work is outlined in Section 2. A detailed description of the problem is presented in Section 3. Section 4 describes the dataset used in this work and how the problem was addressed. Results of the experiment are presented in Section 5. Finally, concluding remarks are given in Section 6.

2. Related Work

It has long been believed that clear prescription contacts do not affect the accuracy of iris recognition. For example, an article appearing in IEEE Computer in 2000 stated, “Successful identification can be made through eyeglasses and contact lenses [9]. As recently as 2010, sources such as Wikipedia contained claims such as, “Iris recognition efficacy is rarely impeded by glasses or contact lenses [13]. Even currently, India’s UIDAI site contains the statement, “Iris recognition is rarely impeded by glasses or contact lenses...” [11].

Baker et al. [1] show that clear, soft, non-cosmetic lenses cause an increase in the FNMR. Using a dataset of images from the LG 2200 sensor, Baker et al found that the FNMR for contact wearers was twenty times higher than non-wearers. The same basic effect is observed with our dataset of images from the newer LG 4000 sensor (Figure 1). Sample images of a subject with and without a contact lens are shown in Figures 2(b) and 2(a) respectively.

Methods for detecting textured lenses have been well researched. He et al. [6] used Local Binary Patterns (LBP) as a feature extraction method and AdaBoost as a learning algorithm to build a textured lens classifier. Zhang et al. [14] investigated the use of Gaussian-smoothed and SIFT-weighted Local Binary Patterns to detect textured lenses in images acquired with multiple iris cameras.

Kohli et al. [7] perform an analysis of the effects of various types of contact lenses on the performance if a commercial iris biometrics system. They present ROC curves demonstrating an improvement when lens detection is used to filter probe images.

3. Problem Definition

Textured contact lenses are designed to alter the appearance of the wearer’s eye, giving it a different color and/or texture. Unfortunately, they also greatly reduce the amount of genuine iris texture visible to iris recognition systems. Increasing the chance of a false non-match. Accordingly, these images should be rejected at the time of segmentation, before a template is generated for them.

The effect of soft lenses is much less obvious. The genuine iris texture is not concealed in the manner it is with textured contact lenses. However, the negative impact on iris biometrics systems by soft lens wearers has been documented [7, 1]. We found a similar degraded performance for clear prescription contacts with the dataset used in this work, as shown in Figure 1. Before an iris biometrics system can perform any corrective measures for subjects who wear soft, non-cosmetic contact lenses, it must first be able to determine which subjects are wearing such lenses.

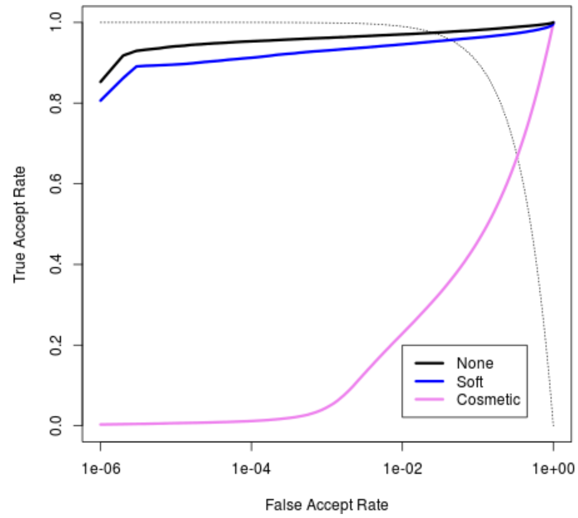

Figure 1. ROC curves for the three classes in the NDCCL12 dataset as determined by a commercially available matcher.

Sample images of an iris with no contact lens, a soft contact lens, and a textured lens can be found in Figure 2. Figure 2(b) demonstrates both how a soft lens can produce an artifact in the iris region of the image, and also how difficult it can be to detect the outer boundary of the soft lens.

4. Experimental Method

4.1. Data Set

Our results in this paper are based on the Notre Dame Cosmetic Contact Lenses 2012 [12] dataset, which to our knowledge is the only publicly-available iris image dataset that supports research in detection of contact lenses. The dataset consists of a training set of $3,000$ images and a verification set of $1,200$ images. All images were acquired with an LG 4000 [8] iris camera. Both the training set and the verification set are divided equally into three classes: (1) no contact lenses, (2) soft, non-cosmetic contact lenses, and (3) textured contact lenses. Categories (1) and (2) are balanced between male and female, and represent a variety of ethnicities. Category (3) images are predominantly Caucasian male.

All textured contact lenses in the dataset came from three major suppliers of textured lenses: Johnson&Johnson, Ciba Vision and Cooper Vision. Multiple colors were selected for each manufacturer and some

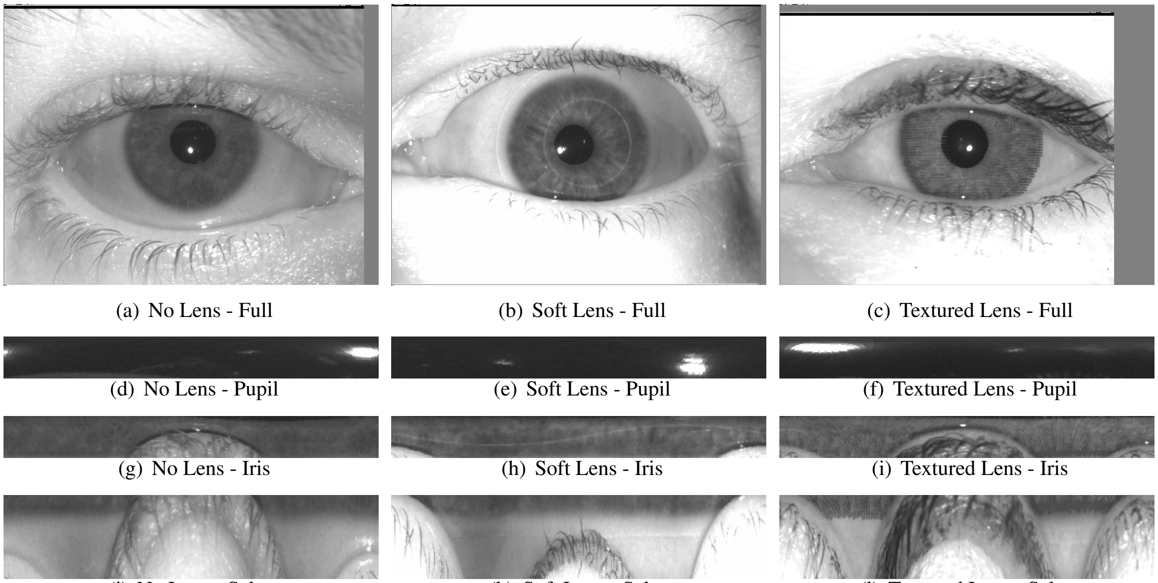

(a) No Lens - Full (b) Soft Lens - Full (c) Textured Lens - Full

(d) No Lens - Pupil (e) Soft Lens - Pupil (f) Textured Lens - Pupil

(g) No Lens - Iris (h) Soft Lens - Iris (i) Textured Lens - Iris

(j) No Lens - Sclera (k) Soft Lens - Sclera (l) Textured Lens - Sclera

Figure 2. Sample images for the three classes showing the original images and the unrolled sections from which the features were extracted. The no lens images were taken from sample 04233d2567. The soft lens images were taken from sample 04261d994. The textured lens images were taken from 07015d1314.

lenses were also toric.

The dataset provides segmentation information as center and radius for circles defining the pupillary boundary and the limbic boundary. The data is also predivided into 10 subject-disjoint and class-balanced folds for training evaluation.

4.2. Feature Extraction

The known segmentation provided by the NDCLD12 dataset is used to break each iris image into three regions: (1) pupil, (2) iris, and (3) sclera. The sclera region is extended slightly such that it includes 20 pixels of iris texture and not more than 60 pixels of scleral data for each angle. This attempts to capture contact lens boundaries that may have shifted into the iris region while also limiting the amount of eyelid and eyelash occlusion. Examples of the regions after extraction can be found in Figure 2.

Modified Local Binary Pattern analysis (described in Figure 3 and similar to [10]) is applied to each region of each image at multiple scales (shown in Figure 4) to produce feature values. The pupil, iris, and sclera all have significantly different appearances, and as such the binary pattern analysis is performed separately for each region. Unlike traditional LBP, this method does not decompose the image into blocks and independently analyze each block to construct a large feature vector. The kernel size for the binary pattern analysis is scaled from 1 to 15 in increments of 1 for a total of 15 different feature sets for each of three regions and 45 feature sets overall.

4.3. Model Training

Fourteen different classifiers, intentionally sampling a variety of different classifier technologies [5], were explored as possible approaches to train models on the feature sets. Each of the 45 feature sets described in Section 4.2 is treated as an independent dataset for the purposes of model training.

For data-mining algorithms that had tunable parameters, a parameter sweep was conducted with reasonable values. A breakdown of the parameters evaluated can be found in Table 1. The predefined folds in the ND-CCL12 dataset are used to evaluate the performance of each trained model by cross-fold evaluation. The parameters yielding the highest correct classification rate (CCR) for each algorithm were recorded, and a model was built with the same parameters but using all training data. This process resulted in an ensemble of trained models totaling $14 * 3 * 15 = 630$ to be evaluated on the verification set.

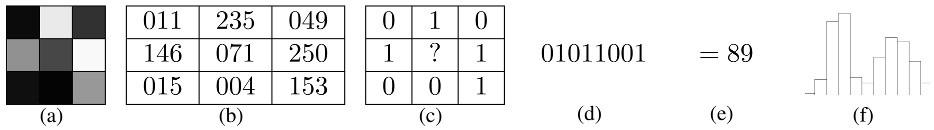

Figure 3. A $3x3$ neighborhood of pixels (a) with values (b) can be converted into a Binary Pattern by comparing perimeter values against the center value. A larger perimeter value gets replaced with a 1 and a smaller or equal perimeter value is replaced with a 0 in (c). Starting with the upper left neighbor and working clockwise, a binary number (d) is formed. The binary value is converted to decimal for a final scalar value (e) and added into the histogram of values for the entire image (f).

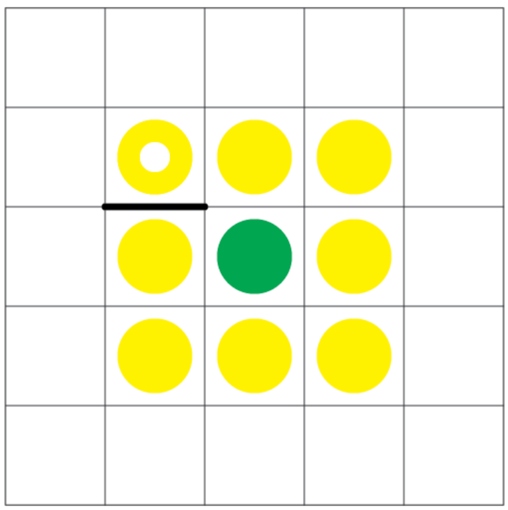

(a) $3x3$ kernel with $scale = 1$

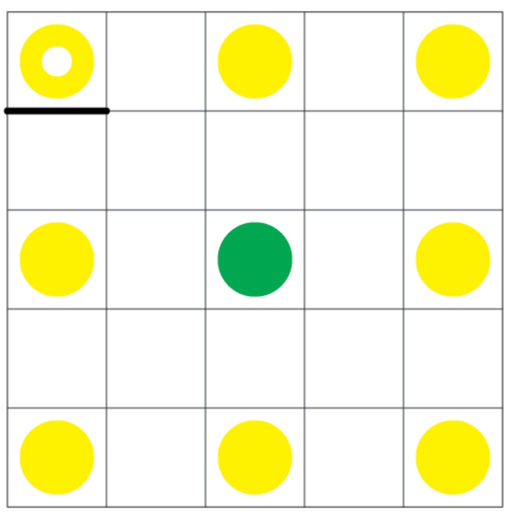

(b) $3x3$ kernel with $scale = 2$

Figure 4. Sample kernels for the modified Local Binary Pattern feature extraction.

4.4. Model Evaluation

To improve the overall CCR of the ensemble of models, an exhaustive comparison of all possible model combinations would ideally be performed. However, this would have resulted in a comparison space of 2^{630} combinations. To reduce the comparison space to a tractable problem, sub-groups of models were considered in a 3-step hierarchical tournament. Each ensemble of 14 models constructed for each of the 45 region-scale combinations was evaluated as an independent group, for a manageable comparison space of $45 * 2^{14} = 737,280$ combinations. This resulted in 45 ensembles of models to be exhaustively compared in a second round of the model evaluation, which while more manageable than before, was still intractable for the scope of this research. Additionally, many of the 45 ensembles achieved 100% CCR on the training data, making further evaluation on the training data irrelevant.

Table 1. Classifiers and Parameter Combinations

Classifier	**Params**	**# Permuts**
Naive Bayes	1	2
Logistic	0	1
M.layer Perceptron	7	2,160
KStar	2	50
IBk	4	52
Bagging	4	80
Logit Boost	8	6,480
JRip	4	2,662
ZeroR	0	1
OneR	1	10
NNge	2	64
J48	3	570
Random Tree	5	1,080
Random Forest	4	576

5. Experimental Results

The performance of the 45 ensembles is evaluated on the $1,200$ image verification set provided as part of the NDCCL12 dataset. For each image of the verification set, a prediction and a confidence is output by each of the 45 model ensembles. A final prediction for each image is decided by taking the maximum of the sum of confidences for each ensemble for each class. The final ensemble resulted in a CCR of over 70% on the three-class problem. The accuracy on detecting instances of textured contact lenses was quite high. Over 96.5% of the images containing a textured contact lens were correctly detected. The CCR on a 2-class[1] problem of deciding whether or not an iris image contains a textured lens was nearly 98%. The confusion matrices can be found in Figure 5.

	None	**Soft**	**Textured**	**Total**
None	262	192	13	467
Soft	114	201	1	316
Textured	24	7	386	417
Total	400	400	400	1,200

(a) 3-Class Problem

	None/Soft	**Textured**	**Total**
None/Soft	396	13	409
Textured	4	387	391
Total	400	400	800

(b) 2-Class Problem (Textured/Non-Textured)

	None	**Soft**	**Total**
None	219	179	398
Soft	160	209	369
Total	379	388	767

(c) 2-Class Problem (Soft/None)

Figure 5. Confusion matrices for the 3-class problem and the 2-class problem. For each matrix, the true class is presented vertically and the resulting classification is presented horizontally. For (c) there were 33 images classified as Textured lenses which were left out of the confusion matrix.

6. Discussion

One major contribution of this work is that it is the first to tackle the important three-class problem of categorizing iris images as no lens, clear lens, or textured lens. Iris recognition systems need the ability to automatically perform this classification, so that appropriate action can be taken.

Another contribution of this work is that we were able to obtain a 96.5% correct detection of iris images that contain a textured contact lens. The textured lenses represented a sampling of different color textured lenses for major manufacturers. This work is the first to use a publicly-available dataset with no-lens, soft-lens, and textured-lens groups.

While we were able to detect 262 of 400 no-lens images in the separate test set, the overall accuracy of distinguishing between no-lens images and clear-lens images is still too low.

Further experimentation will attempt to improve the soft lens detection rate. More advanced features than the simple binary patterns evaluated in this work may yield higher correct classification rates. Future work will simplify features and classifiers to only those actually needed.

[1] Translating the dataset into a two-class problem while maintaining balanced classes required half of the no lens and half of the soft lens data to be removed from consideration. To preserve the distribution of no lens / soft lens, male / female, and left / right eye, eight counters were maintained while traversing the original ordering of the validation set file. Only images encountered on the even numbers were retained. For instance, the male-right-soft counter was incremented every time a right-eye image of a male subject wearing a soft contact lens the counter was increased. If the counter was odd after the increment, the image was included in the dataset, otherwise it was excluded. A listing of the images that were included in the two class problem can be obtained via email.

References

[1] S. Baker, A. Hentz, K. Bowyer, and P. Flynn. Degradation of iris recognition performance due to non-cosmetic prescription contact lenses. *Computer Vision and Image Understanding*, 114(9):1030–1044, 2010.

[2] J. Daugman. How iris recognition works. *IEEE Transactions on Circuits and Systems for Video Technology*, 14(1):21–30, Jan. 2004.

[3] J. Daugman. Recognising persons by their iris patterns. *Advances in Biometric Person Authentication*, pages 783–814, 2005.

[4] L. Flom and A. Safir. Iris recognition system, Feb. 3 1987. US Patent 4,641,349.

[5] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten. The weka data mining software: an update. *SIGKDD Explor. Newsl.*, 11(1):10–18, nov 2009.

[6] X. He, Y. Lu, and P. Shi. A fake iris detection method based on fft and quality assessment. In *Chinese Conference on Pattern Recognition*, pages 1–4. IEEE, 2008.

[7] N. Kohli, D. Yadav, M. Vatsa, and R. Singh. Revisiting iris recognition with color cosmetic contact lenses. In *IAPR 6th International Conference on Biometrics*, 2013.

[8] LG. Lg 4000 camera, October 2011. `http://www.lgiris.com`.

[9] M. Negin, T. Chmielewski Jr, M. Salganicoff, U. von Seelen, P. Venetainer, and G. Zhang. An iris biometric system for public and personal use. *Computer*, 33(2):70–75, 2000.

[10] T. Ojala, M. Pietikäinen, and D. Harwood. A comparative study of texture measures with classification based on featured distributions. *Pattern Recognition*, 29(1):51–59, 1996.

[11] UIDIA. Biometric devices, December 2012. `http://www.uidai.gov.in/index.php?option=com_content&view=article&id=168&Itemid=165`.

[12] University of Notre Dame Computer Vision Research Lab. ND Cosmetic Contact Lenses 2012 Dataset, December 2012. `http://www.nd.edu/~cvrl/CVRL/Data_Sets.html`.

[13] Wikipedia. Iris recognition, October 2010. `http://en.wikipedia.org/wiki/Iris_recognition`.

[14] H. Zhang, Z. Sun, and T. Tan. Contact lens detection based on weighted lbp. In *20th International Conference on Pattern Recognition (ICPR)*, pages 4279–4282. IEEE, 2010.