

Facial Recognition of Identical Twins

Matthew T. Pruitt, Jason M. Grant, Jeffrey R. Paone, Patrick J. Flynn
University of Notre Dame
Notre Dame, IN

{mpruitt, jgrant3, jpaone, flynn}@nd.edu

Richard W. Vorder Bruegge
Digital Evidence Laboratory, Federal Bureau of Investigation
Quantico, VA

Richard.VorderBruegge@ic.fbi.gov

Abstract

Biometric identification systems must be able to distinguish between individuals even in situations where the biometric signature may be similar, such as in the case of identical twins. This paper presents experiments done in facial recognition using data from a set of images of twins. This work establishes the current state of facial recognition in regards to twins and the accuracy of current state-of-the-art programs in distinguishing between identical twins using three commercial face matchers, Cognitec 8.3.2.0, VeriLook 4.0, and PittPatt 4.2.1 and a baseline matcher employing Local Region PCA. Overall, it was observed that Cognitec had the best performance. All matchers, however, saw degradation in performance compared to an experiment where the ability to distinguish unrelated persons was assessed. In particular, lighting and expression seemed to have affected performance the most.

1. Introduction

Biometric signatures are dependent on the assumption that each individual is unique. Identical twins can have biometric signatures that are very similar, especially when the signature is derived from a face image. While face recognition software exhibited poor performance in this context, there are other biometric modalities that can offer a performance increase at the cost of increased invasiveness. Jain and Prabhakar [6] showed that the false accept rate for an identical twin recognition task was found to be 2-6% higher than for a dataset with a normal twin-nontwin distribution.

The primary motivation of this work is to assess the accuracy of current generation facial recognition systems on a particularly challenging data set, containing twins. In particular, we employ PittPatt 4.2.1 [11], Cognitec 8.3.2.0 [3], and VeriLook 4.0 [7] as representative of the state of the art facial recognition systems. We also compare these commercial matchers against the LRPCA matcher [10] as a baseline

method.

This paper is organized as follows. Section 2 discusses related work in the performance of face recognition software on twins. Section 3 outlines the data set used and the parameters of the experiments performed in this paper. In Section 4 reviews the results of the experiments performed. Lastly, Sections 5 and 6 discuss the results and future work.

2. Background

The problem with discrimination of identical twins using facial recognition has been noted. Sun et al. [12] used the CASIA Multimodal Biometrics Database of Twins containing the 134 subjects, including both identical and non-identical twins, to perform recognition experiments. Using the Cognitec FaceVACS system, a true accept rate of approximately 90% at a false accept rate greater than 10% was obtained. In contrast to the work of Sun et al. [2], this study takes multiple experimental variables into consideration when characterizing performance. The resolution of the images used here is much higher, but the changes to values of facial pose variables (pitch, roll, yaw) are not considered. This problem has also been explored by Phillips et al. in [9]. Specifically, same day and cross year performance were explored using the top three submissions from the Multiple Biometric Evaluation (MBE). The names of the algorithms, however, were not given in the paper. The main conclusion of this paper was that the best performance came when all images were taken on the same day with a neutral expression and controlled lighting.

2.1. Matchers

2.1.1 LRPCA

Principal Component Analysis (PCA) for face recognition, as first described by Turk and Pentland in [13] represents face images as points in a high dimensional space determined via training. To perform a match, one can project a probe face into the space and take the label of the closest

projection of an existing face as the identity of the projected probe.

The Local Region Principal Component Analysis (LR-PCA) is a facial recognition algorithm developed by Colorado State University which initially was developed as a baseline algorithm for the Good, Bad, and the Ugly Face Recognition Challenge Problem. This PCA approach reduces the face to 128 by 128 pixels. Next, it segments the face region into 13 subregions including an eye and two eyebrow regions (for both the left and right eye); an upper, middle, lower left, and lower right regions of the nose; and the left, middle, and right regions of the mouth. These regions are normalized to attenuate variation in illumination. PCA is performed on these 13 subregions and the entire face block and score fusion is used to obtain an identity label.

2.1.2 PittPatt 4.2.1

Recognition experiments were performed using the Pittsburgh Pattern Recognition system known as PittPatt [11]. PittPatt is optimized to work on low resolution image data. It performs recognition in three stages at varying sizes (12, 20, and 25 pixels between the eyes). Texture information is not typically available at these resolutions. In version 2 of the Multiple Biometrics Grand Challenge (MBGC), PittPatt was found to have a true accept rate of approximately 90% at false accept rates as low as 0.1% using controlled lighting and images as small as 90 by 90 pixels [8]. On a large dataset of 1,600,000 images in the MBE [4], PittPatt yielded a rank one identification rate of only 62%, which was lower than most other performers.

2.1.3 Verilook 4.0

The VeriLook Software Developer's Kit is a set of facial recognition tools developed by Neurotechnology. Neurotechnology claims that their software assures system performance and reliability by implementing live face detection, simultaneous multiple face recognition, and fast face matching in one-to-one and one-to-many modes [7]. The algorithm is able to perform simultaneous multiple face processing in a single frame and uses live face detection. At its worst, VeriLook obtained a False Reject Rate of 3.201% at a False Accept Rate of 0.1% in the Face Recognition Grand Challenge (FRGC) evaluation. .

2.1.4 FaceVACS 8.3.2.0

Cognitec is a commercially available face recognition system making use of a Face Visual Access Control System (FaceVACS). Cognitec has competed in several face recognition challenges. The general aim of Cognitec is to perform efficiently on large-scale databases with high performance. Cognitec software achieved a rank one identification rate of 83% on a dataset of 1,600,000 images in the MBE and on an even larger dataset of 1,800,000 images obtained a rank one identification rate of 88%.

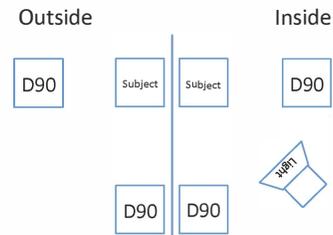


Figure 1. The 2009 setup. Frontal images were taken both indoors and outdoors using all four cameras.

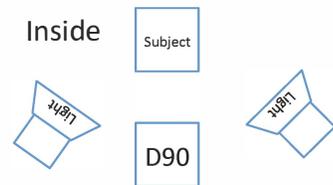


Figure 2. The 2010 setup. Images were taken indoors from a single camera.

3. Experiments

3.1. Data Set

A dataset containing images of twins was used. Images of subjects were acquired at the Twins Day festival [1] in 2009 using a three camera set up as seen in Fig. 1. The same images were acquired both inside and outside. Two frontal face photos were acquired from two separate cameras. The subject would look at camera three, then rotate to camera one. This would cause differences in the sensor and lighting on the face.

In 2010, a slightly different setup was used at the same festival. A five camera setup as seen in Fig. 2 was used with data taken indoors only. There were two frontal images taken in 2010, however, they were taken from the same sensor within seconds of each other. Subjects were instructed to stand up, turn around, and sit back down after the first photo was taken in order to take a second photo.

From this data set, only the frontal images, corresponding to a yaw of zero degrees were used. Due to the way the data was collected, galleries taken from different cameras could be used.

From 2009 a set of 5,800 images was used for this experiment. 1,635 images from 2010 were used as well. The images were converted to PPM format from the original Nikon RAW image format (NEF) taken using Nikon D90 cameras and had a resolution of 2848 columns and 4288 rows.

3.2. Using the matchers

3.2.1 PittPatt 4.2.1

PittPatt first creates sets of probe and gallery templates. This step takes the longest to complete. Images are loaded and PittPatt templates are generated, then added to a single gallery or probe file. After gallery and probe file creation is

complete, the probe and gallery templates are compared to each other in an "all versus all" fashion to produce a similarity matrix, which takes less than a minute.

3.2.2 VeriLook

To perform matching experiments, a Neurotechnology proprietary template was created for each enrolled image. After creating a template for each image, a single template was matched against all templates in the gallery. There was no distinction between templates used as probes or templates used in the gallery. The matching experiment returns a matrix of similarity scores which ranged from 0 to 180, with higher values corresponding to higher probable matches.

3.2.3 Cognitec

Cognitec first detects the eyes and then face locations in the image. The image with eye positions makes up the Primary Facial Data. Once a face is found, the image is normalized and standard preprocessing techniques are applied. Some of the preprocessing techniques include histogram equalization and intensity normalization. Features of the face are then extracted and stored as a vector in high-dimensional space. This vector represents the Secondary Facial Data. The similarity between two faces is then the distance between the two corresponding vectors. A correct verification can be made if the distance score exceeds a given threshold.

3.2.4 LRPCA

The LRPCA algorithm was executed in the partially automated mode; therefore, we preprocessed the set of images to determine the horizontal and vertical positions of each pair of eyes in the set. Signature sets were generated, which contained the list of images, their locations, and filetypes. Together, the signature sets and the table of eye locations, were passed to LRPCA which performed the rest of the matching experiment internally. Perfect matches received a score of 1 and the lowest match scores averaged approximately -0.5.

3.3. Variables

In these experiments, there are three variables.

- For lighting, controlled lighting images were acquired inside a tent using studio lamps. For uncontrolled lighting, images were taken outside.
- Eyeglasses were also used in the data acquisition phase. These were a set of zero-prescription glasses used to simulate normal eyewear.
- Subjects were prompted to express either a neutral expression or a smile.

Examples of the different images used in this experiment can be seen in Fig. 3.

For the graphs in this paper, the legends refer to the systems by their system ID. Table 1 shows the translation between system id and system name.

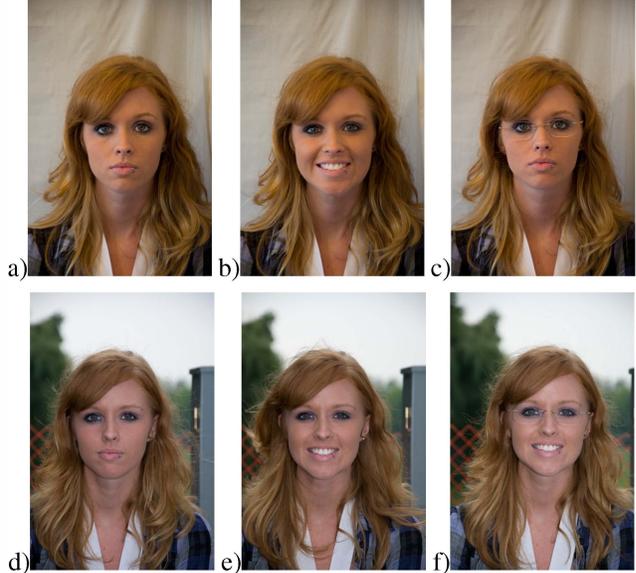


Figure 3. Images representing different covariate conditions acquired at Twins Days for a subject. (a) Inside with Neutral Expression (b) Inside with Smile (c) Inside with Neutral Expression and Non-Prescription Glasses (d) Outside with Neutral Expression (e) Outside with Smile (f) Outside with Smile and Non-Prescription Glasses

System ID	System Name
PFE3	LRPCA
VL40	VeriLook 4.0
C	Cognitec 8.3.2.0
P42	PittPatt 4.2.1

Table 1. Matchers and System IDs

4. Results

In these experiments, three covariates are examined, namely expression, lighting, and glasses. For the baseline experiments, the expression is neutral, the lighting is controlled, and the subject is not wearing any kind of glasses. These correspond to ideal conditions for matching. Bootstrapping was performed by randomly sampling the match and non-match scores separately to get a set 80% of the original size with the same mix of match and non-match scores.

After enrollment, an all versus all experiment was performed on all four matching systems. Not all matchers were able to enroll every image leading in a different number of entries in a similarity matrix for each matcher. All matchers were symmetric, so only the upper right half of the similarity matrix was stored.

4.1. Baseline Performance

Performance results are given using standard ROC plots with box and whisker plots staged at even intervals along the curve. Here, the whiskers represent a 95% confidence interval for the True Accept Rate at that False Accept Rate and

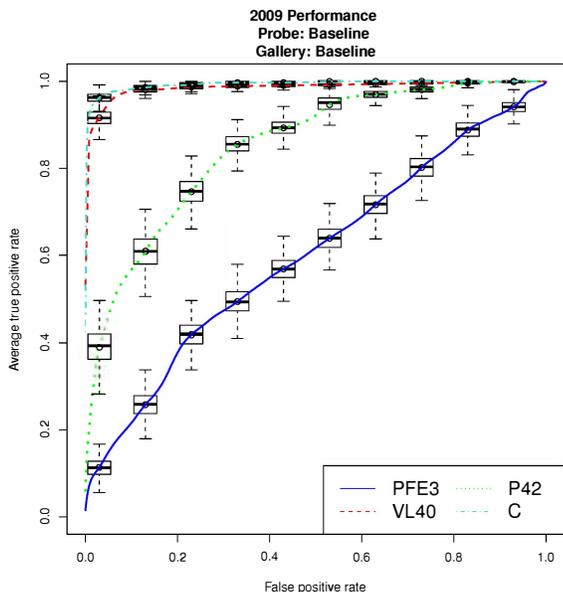


Figure 4. Performance of all matchers under the baseline conditions.

System	Probe	Gallery	Rate
Cognitec	Baseline	Baseline	0.03
LRPCA	Baseline	Baseline	0.43
VeriLook	Baseline	Baseline	0.05
PittPatt	Baseline	Baseline	0.24

Table 2. 2009 Baseline EER

the top and bottom of the box represent the upper and lower quartiles respectively. From the ROC curve in Fig. 4, we can immediately see that both Cognitec and VeriLook dominate PittPatt and LRPCA by a significant margin with both VeriLook and Cognitec being very close in performance. The whiskers of the box plots overlap between the Cognitec and VeriLook curves, however, indicating that there is no statistically significant difference between the performance of these matchers. The confidence intervals for PittPatt and LRPCA are significantly different; their confidence intervals are wider, though, indicating more variability and subsequently less reliability in the performance.

To get a better estimate of how well separated these curves are, it is beneficial to look at the equal error rates as well. Looking at Table 2, it can be seen that Cognitec has a smaller equal error rate by a small margin. Nonetheless, both Cognitec and VeriLook are still very close to perfect. On the other hand, PittPatt and LRPCA both fall short here. A score of 0.43 for LRPCA indicates that LRPCA is performing just slightly better than random chance. The confidence bands for LRPCA actually intersect to the random line in some places.

One reason PittPatt may be performing so poorly when compared to the other matchers is that it is optimized for smaller faces. Using such large images can play a part

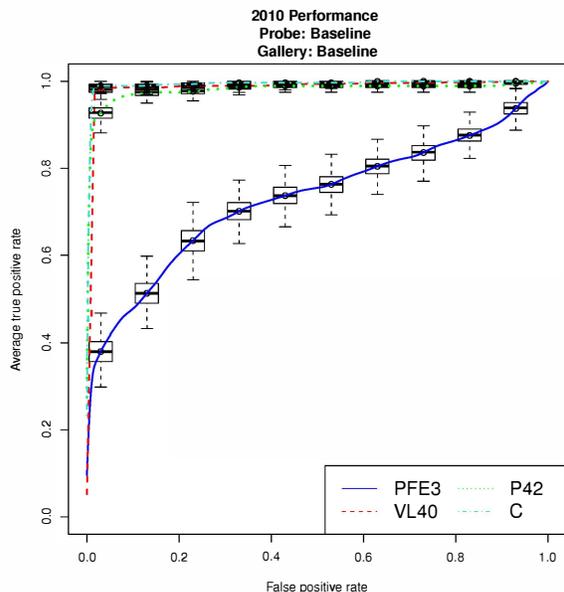


Figure 5. This ROC show the performance of all matchers under the baseline conditions. As can be seen in the graph, all matchers saw a performance boost over the 2009 data. This may be due to the face that the probe and gallery images from 2010 were taken within seconds of each other with the same camera and in the same position.

System	Probe	Gallery	Rate
Cognitec	Baseline	Baseline	0.01
LRPCA	Baseline	Baseline	0.31
VeriLook	Baseline	Baseline	0.02
PittPatt	Baseline	Baseline	0.05

Table 3. 2010 Baseline EER

in the poor performance on identical twins. Cognitec and VeriLook may be able to use the extra texture in the large number of pixels across the face.

We assumed that 2010 baseline data would follow the same trend as was seen in the 2009 baseline data. This is not the case. As can be seen in Fig. 5, all matchers saw a significant performance boost over the 2009 results. The confidence bands indicate no statistically significant difference between all three commercial matchers, perfect performance being within the confidence intervals. Even LRPCA sees a significant improvement in performance. VeriLook, Cognitec, and PittPatt perform nearly perfect in this case. A reason for the improvement may come from the way the cameras were set up at Twins Days. The two frontal images that are being compared in 2010 are taken from the same camera at the same position within seconds of each other and are likely to be more similar than images in the 2009 data where subjects had their second frontal image taken with a different sensor at a different relative position. This phenomenon can be seen in all ROC curves generated from 2009 data .

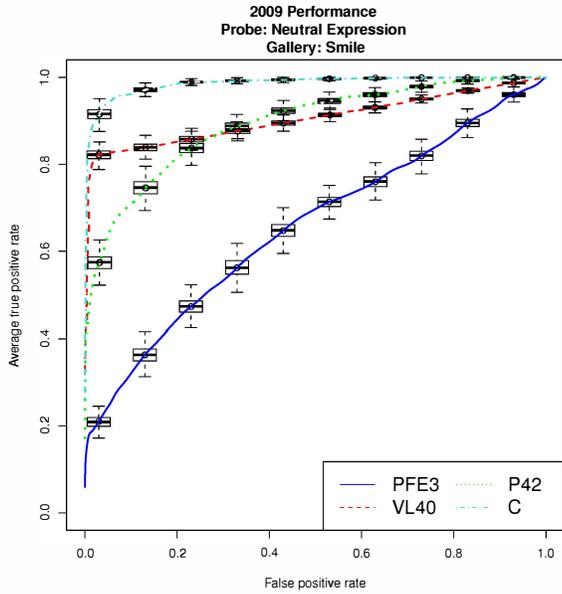


Figure 6. These ROC curves show the performance of the four matchers for 2009 data where the probe has a neutral expression while the gallery has a smile. There is degradation in performance across all four matchers from the 2009 baseline.

4.2. Expression

The expression covariate variable’s performance impact was measured under controlled indoor lighting where no glasses are worn. A change in expression can cause different wrinkles and textures to appear in the face that might have not been apparent in the baseline images.

The expression experiments take two factors into account. The first is where the probe image has a neutral expression and the gallery image has a smiling face. This condition set represents the most difference between probe and gallery.

The ROC curve in Fig. 6 shows that the performance of Cognitec is statistically significantly better than all other matchers. VeriLook performs better than PittPatt when FAR is less than 0.2, well within a normal operating range.

Comparing this information with the baseline for 2009, the medians of all four curves has been shifted down indicating worse performance. VeriLook saw a statistically significant decrease in performance from the 2009 baseline data. This is again supported by the EER as seen in Table 4. The EER of VeriLook is significantly higher than that of the EER for VeriLook with the 2009 data while the other matchers are relatively unaffected.

When both the probe and gallery faces are smiling, it can be seen that both VeriLook and Cognitec perform statistically significantly better than PittPatt and LRPCA, even though there is no statistically significant difference between the performance of VeriLook and PittPatt. The statistically significant difference in performance for VeriLook during cross-expression matching indicates that expression

System	Probe	Gallery	Rate
Cognitec	Neutral Expression	Smile	0.05
LRPCA	Neutral Expression	Smile	0.39
VeriLook	Neutral Expression	Smile	0.16
PittPatt	Neutral Expression	Smile	0.19

Table 4. 2009 Neutral Expression vs. Smile EER

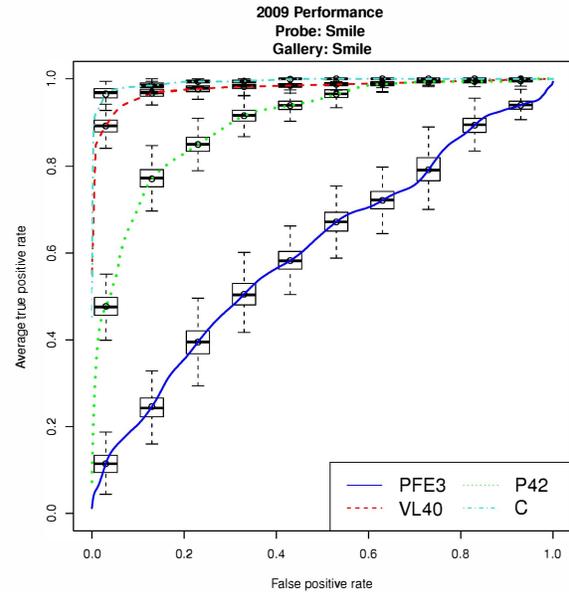


Figure 7. These ROC curves show the performance of the four matchers for 2009 data where the probe and gallery both are smiling. This shows better performance for comparisons where both probe and gallery are of the same expression rather than different expression.

System	Probe	Gallery	Rate
Cognitec	Smile	Smile	0.03
LRPCA	Smile	Smile	0.42
VeriLook	Smile	Smile	0.06
PittPatt	Smile	Smile	0.18

Table 5. 2009 Smile vs. Smile EER

change can potentially degrade performance. The same cannot be said when both probe and gallery have the same expression, regardless of whether or not that expression is neutral or a smile. This would suggest that for best matching results, the gallery should contain samples of multiple expressions, or at least trained with data of the expected expression for matching sessions.

The performance on the 2009 data where both probe and gallery were smiling did not show as much degradation as the performance of cross-expression matching. The EER, seen in Table 5 shows a similarity to the baseline data from 2009. None of the ROC curves, seen in Fig. 7, have any statistically significant difference from the ROC curves from baseline 2009 experiments.

While the 2010 data, where probe and gallery had different expressions, shows the same trend as seen in the 2010 baseline data, it can be seen that there is still a statistically significant difference in the performance of all three matchers, seen in Fig. 8. There was no statistically significant difference for the 2010 data where both probe and gallery were smiling for all commercial matchers from the 2010 baseline. This is particularly interesting considering the observation that this data should be inherently similar due to the conditions under which the data was acquired.

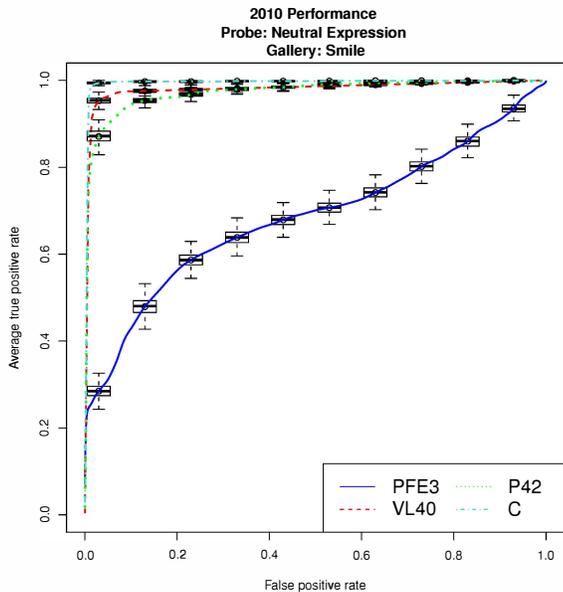


Figure 8. These ROC curves show the performance of the four matchers for 2010 data where the probe has a neutral expression while the gallery has a smile.

4.3. Lighting

For the lighting covariate analysis, subjects had images taken both indoors under studio lighting and outdoors with uncontrolled lighting. In this instance, uncontrolled lighting can range anywhere from direct sunlight to the amount of light visible through the clouds on a rainy day.

As the setup between 2009 and 2010 changed, outdoor images were only acquired in 2009. This means, unfortunately, that only an analysis of the 2009 data was possible for the lighting covariate. Just like the expression covariate, though, we have images where probe and gallery were taken under different lighting schemes and images where both probe and gallery were taken under uncontrolled lighting schemes.

Looking at the ROC curves for 2009 data, from Fig. 9, it can be seen that Cognitech and VeriLook perform well as compared to LRPCA and PittPatt. In the case of lighting, Cognitech actually performs statistically significantly better than VeriLook with low variability in both of the curves. Comparing this data to the baseline, VeriLook did not perform statistically significantly worse than VeriLook from

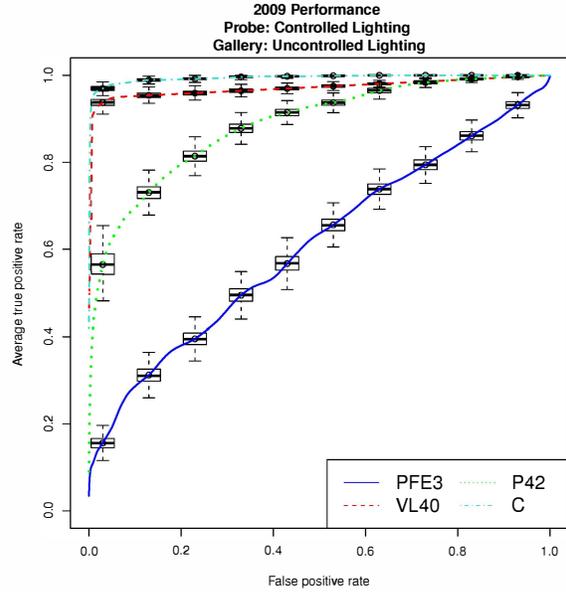


Figure 9. These ROC curves show the performance on 2009 data where the probe was taken under controlled lighting and the gallery was taken under uncontrolled lighting. There is only slight degradation from the baseline for the same year.

System	Probe	Gallery	Rate
Cognitech	Controlled	Uncontrolled	0.03
LRPCA	Controlled	Uncontrolled	0.43
VeriLook	Controlled	Uncontrolled	0.05
PittPatt	Controlled	Uncontrolled	0.20

Table 6. 2009 Controlled Lighting vs. Uncontrolled Lighting EER

the 2009 baseline, even though it was statistically significantly worse than Cognitech for this experiment. In fact, no curve performed statistically significantly worse than its corresponding curve from the 2009 baseline, although there was more variability. In Table 6, it can be seen that Cognitech, VeriLook, and LRPCA all have the same EER as their 2009 baseline counterpart.

Looking back at the performance from the cross-expression experiment with 2009 data, cross lighting does not seem to effect matcher performance as much as cross-expression matching. VeriLook is statistically significantly better for cross-lighting. The rest of the curves are not statistically significantly different.

Comparing the performance of cross-lighting conditions to same lighting conditions where both probe and gallery were taken under uncontrolled lighting, we can see that it is better to have at least one image, probe or gallery, be taken under controlled lighting. Performance for all commercial matchers saw a statistically significant drop in performance from both the cross-lighting experiment and the baseline performance for 2009 as can be seen in Fig. 10. The EER for all commercial matchers was greater across the board as well. The overlap of scores distributions was

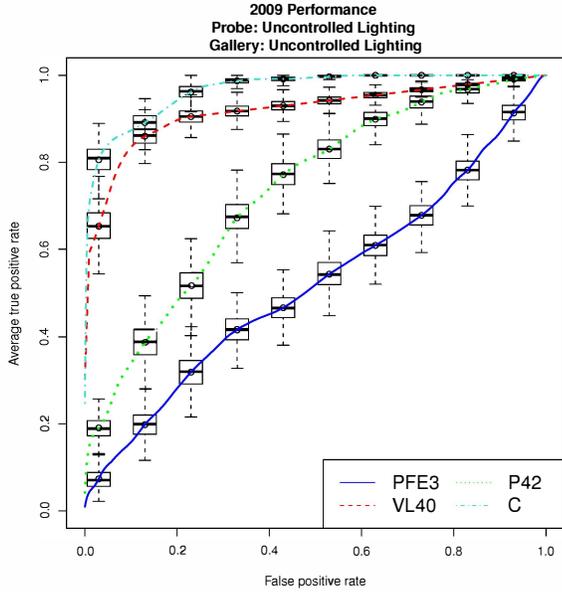


Figure 10. These ROC curves show the performance on 2009 data where both probe and gallery were taken under uncontrolled lighting. A degradation in performance can be seen from both the baseline for 2009 and the situation where at least one image was taken under controlled lighting.

System	Probe	Gallery	Rate
Cognitec	Uncontrolled	Uncontrolled	0.12
LRPCA	Uncontrolled	Uncontrolled	0.49
VeriLook	Uncontrolled	Uncontrolled	0.14
PittPatt	Uncontrolled	Uncontrolled	0.33

Table 7. 2009 Uncontrolled vs. Uncontrolled EER

greater as well, especially for LRPCA and PittPatt.

4.4. Eyewear

4.4.1 Zero-Prescription Glasses

These results are presented under controlled indoor lighting conditions where the subject is wearing a neutral expression. While the effect of eyewear has been studied in other papers, it has not been determined whether it is the frames of the glasses that cause a change in performance or the change in interocular distance associated with prescription lenses. The glasses used here were designed to reduce the visibility of the frames to elucidate the effect of interocular distance. As can be seen in Fig. 11, we see Cognitec and VeriLook outperforming PittPatt and LRPCA. While Cognitec and VeriLook are not statistically significantly different, PittPatt’s performance is statistically significantly worse than the other two. The interesting fact about this graph comes when we compare the performance with the 2009 baseline performance.

While Cognitec and VeriLook are not statistically significantly different that the performance in the 2009 baseline,

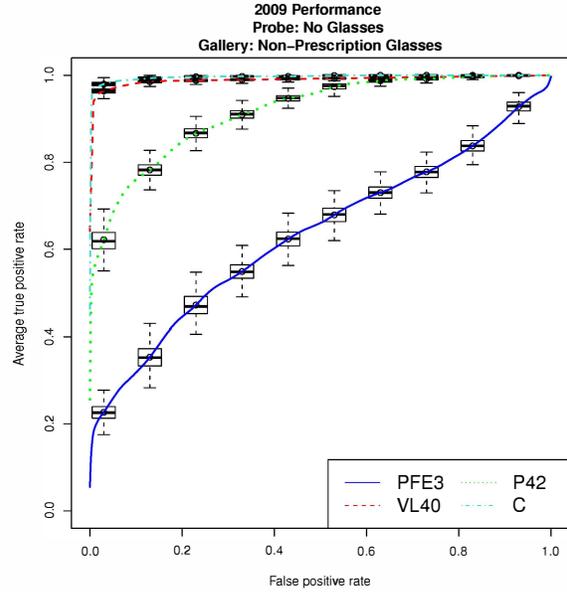


Figure 11. These ROC curves show performance on 2009 data for a probe that had no glasses and a gallery wearing non-prescription glasses. There is minimal performance degradation from the baseline.

System	Probe	Gallery	Rate
Cognitec	No Glasses	Glasses	0.02
LRPCA	No Glasses	Glasses	0.40
VeriLook	No Glasses	Glasses	0.04
PittPatt	No Glasses	Glasses	0.17

Table 8. 2009 No Glasses vs. Non Prescription-Glasses EER

PittPatt actually performs better here by a statistically significant margin. As was seen with the other covariates, such as expression or lighting, intuition would seem to suggest deviating from the baseline would cause performance to be less than or equal to the baseline performance. This result, however, shows the contrary. When one image, probe or gallery, has an image with glasses, the performance seems to be better.

While this may seem better, there could be other explanations for this result. Glasses can cause an increase in failures to enroll, but we have not been able to confirm that in this experiment. If this effect did occur, then there would be less data in this experiment than there was in the baseline, which would cause a statistical irregularity. As a result, it remains to be determined if glasses can improve the performance of the PittPatt algorithm.

When we look at the performance where both probe and gallery are wearing non-prescription glasses, seen in Fig. 12, the scores are not statistically significantly different from the 2009 baseline nor significantly different from the cross-glasses experiment.

The results reported above which show little or no degradation in performance with eyeglasses contradict earlier

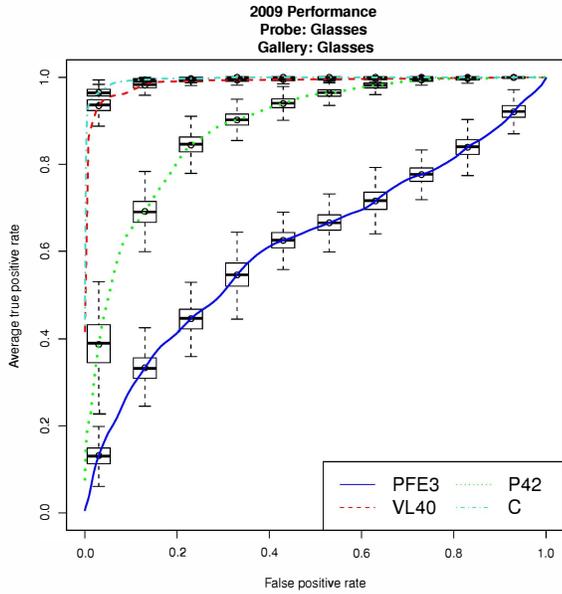


Figure 12. These ROC curves show performance on 2009 data where both probe and gallery were wearing non-prescription glasses. There is slightly more degradation in performance from where at least one image did not wear glasses.

System	Probe	Gallery	Rate
Cognitec	Glasses	Glasses	0.03
LRPCA	Glasses	Glasses	0.39
VeriLook	Glasses	Glasses	0.05
PittPatt	Glasses	Glasses	0.20

Table 9. 2009 Non Prescription vs. Non Prescription EER

results. This may indicate that the degradation observed in previous investigations may reflect the presence of the frames, more than the eyeglasses themselves.

5. Discussion

These experiments showed that there will be a need for better techniques to differentiate between twins. While current technologies can distinguish between twins most of the time under near ideal conditions, as the imaging variables between probe and gallery vary, the accuracy of these systems can decrease, as seen here and as seen in [5] with the FaceVACS system. In addition, the false accept rate under which one obtains these recognition rates is very high. The most significant variables that can affect recognition systems seems to be expression and lighting. By using marks on the face as recognition features, however, these variables would be less noticeable. The glasses would not be as much of a problem either since the eye area is masked.

6. Current and Future Work

While there is both 2009 and 2010 data currently available, the number of twins captured in both acquisitions is

still insufficient to perform a statistically significant comparison to see whether or not aging can be a factor in the matching phase. For future papers, having more cross year data can be used to see the effect of aging on twins recognition which may turn out to be significant due to the fact that as twins age, their features are more influenced by natural processes leading to different features for each twin.

References

- [1] Twins days festival. <http://www.twinsdays.org>.
- [2] T. Bracco. Our 'not so impossible' mission. *Network World*, 2000.
- [3] C. S. Corporation. Cognitec brochure. <http://www.cognitec-systems.de>, Sep 2010.
- [4] P. Grother, G. Quinn, and P. Phillips. Report on the evaluation on 2d still-image face recognition algorithms. Technical report, NIST IR 7709, 2010.
- [5] A. Jain and U. Park. Facial marks: Soft biometric for face recognition. *IEEE International Conference on Image Processing*, pages 37–40, November 2009.
- [6] A. Jain and S. Prabhakar. Can identical twins be discriminated based on fingerprints? Technical Report MSU-CSE-00-23, Michigan State University, 2002.
- [7] NEUROtechnology. VeriLook SDK Brochure. <http://www.neurotechnology.com>.
- [8] P. Phillips. Still face challenge problem: Multiple biometric grand challenge preliminary results of version 2. http://face.nist.gov/mbgc/2009/FACE_V2_FINAL.pdf, December 2009.
- [9] P. Phillips, P. Flynn, K. Bowyer, R. V. Bruegge, P. Grother, G. Quinn, and M. Pruitt. Distinguishing identical twins by face recognition. In *2011 IEEE Conference on Automatic Face Gesture Recognition and Workshops (FG 2011)*, March 2011.
- [10] P. J. Phillips, J. R. Beveridge, B. A. Draper, G. Givens, A. J. O'Toole, D. S. Bolme, J. Dunlop, Y. M. Lui, H. Sahibzada, and S. Weimer. An introduction to the good, the bad, & the ugly face recognition challenge problem. In *The 9th IEEE Conference on Automatic Face and Gesture Recognition*, Santa Barbara, CA, March 2011.
- [11] H. Schneiderman, N. Nechyba, and M. Sipe. PittPatt 4.2. <http://pittpatt.com>, 2010.
- [12] Z. Sun, A. A. Paulino, J. Feng, Z. Chai, T. Tan, and A. K. Jain. A study of multibiometric traits of identical twins. In *Proceedings of the SPIE, Biometric Technology for Human Identification VII*, volume 7667, pages 76670T–12, 2010.
- [13] M. Turk and A. Pentland. Face recognition using eigenfaces. *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 586–591, 1991.