Biometric Identification using Song-Based Blink Patterns

Tracy Westeyn, Peter Pesti, Kwang-Hyun Park & Thad Starner
GVU, College of Computing
Georgia Institute of Technology
Atlanta, GA 30332
{turtle, pesti, akaii, thad}@cc.gatech.edu

Abstract

In this paper we describe a system that uses patterns of eye blinks as a biometric. Each user chooses a personal blink pattern based on a song (for example "Jingle Bells"). To establish their identity, the user looks at the system's camera and blinks to the cadence of their chosen song. Computer vision is used to detect the blinks and to compare the blinked cadence to a database of stored blinked patterns to determine which song is being blinked, and therefore which user is performing the blinking. To make the system more secure, the system compares the characteristics of the blinking itself, the "blinkprint", to the user's stored blinkprint. This provides a verification check to help protect against the theft and subsequent use of a user's blink pattern. We discuss the possible use of an enrollment process that alerts new users when their new blink code is similar to other codes already in the database, and report on the results of a preliminary experiment using this new enrollment procedure.

1 Introduction

After September 11th, security measures in public places, such as airports, have increased. Keys, identity badges and numeric keypads are utilized to restrict access to sensitive areas, allowing only authorized personnel to gain entry. These physical systems rely on possessing specific knowledge (numeric codes), specific devices (keys and badges) or a combination of the two. In either case, physical systems rely on the manipulation or possession of physical devices and/or access to special knowledge and can easily be compromised by an unauthorized person obtaining that knowledge (learning the numeric codes) and/or acquiring the physical device (stealing a key). Physical devices give the illusion of security, as possession is the only thing a person needs to gain access to a restricted area. Biometrics rely on physiological and/or behavioral characteristics, rather than associating access with the possession of special items or knowledge. This, in turn, associates access more directly to a person by using personal characteristics that cannot be easily duplicated.

Biometrics are the use of intrinsic physiological and/or behavioral characteristics to verify the identity of an individual. Fingerprint identification, DNA identification and face recognition are three examples of well-known physiological biometrics. An example of a behavioral biometric is recognizing someone by the way they walk, known as gait recognition. Biometrics offer several security advantages. Unlike keys and passwords, physical biometric traits cannot be lost or forgotten by an individual. Most biometrics are also very difficult to steal. While one could easily steal a person's key, it would be more difficult to mimic someone's gait.

While biometrics provide several advantages, they are not without fault. Some biometrics are susceptible to spoofing. For example, 2D face recognition can be fooled using photographs (Ross, Jain & Qian, 2001). Biometrics can also suffer from lack of generality, as not all individuals possess certain biometrics indicators. For example, while everyone has fingerprints, not everyone's fingerprints are capable of conveying uniquely identifying information. Due to the frailties of using a single biometric indicator, research in the past few years has seen a heightened interest in the fusion of multiple biometric indicators to authenticate identity (Ross et al., 2001). Combinations of biometric indicators, such as face recognition and iris recognition (Wang, Tan & Jain, 2003), can help boost accuracy of identification. In addition, as the number of biometric indicators used increases, the chance of defeating the system decreases.

Another issue that is common to both single and multiple biometric indicator systems is the method of enrollment. Enrollment refers to the process of associating an identity with data exemplars produced from sensor readings and storing them in a database. In face recognition systems, enrollment associates the name of a person with images of her face. This stored data is the template to which subsequent readings will be compared to validate a person's identity. The enrollment process varies according to the type of biometric and recognition process being used. For example, 2D frontal face recognition requires the enrollment of a few static images with the person appearing in the same approximate location in each image. Enrollment for iris data, on the other hand, has strict image requirements with regard to pupil size and location in the image (Wang et al., 2003). Poor image quality can lead to enrollment failures. This can result in several attempts to enroll a single image, which can be both time consuming and frustrating to the user.

In previous work, we introduced blinking as a behavioral–based biometric (Westeyn & Starner, 2004). Our preliminary experiments suggest that blinking, as a biometric indicator, is ideal for fusion with face recognition. In this paper, we introduce an enrollment process for an identity verification system based on blink recognition that helps increase the viability of blinking as a biometric indicator. With increased verification accuracy, blinking could be combined with face recognition to allow non-obtrusive, hands-free biometric-based identification systems to be created. We also feel this enrollment procedure can be modified to create a usable, blink-based interface for assistive technologies. We will briefly revisit this idea later in the paper.

The remainder of the paper is arranged as follows: first, we briefly motivate using blinking as a biometric; second, we describe our enrollment procedure; third, we address the hardware requirements and algorithms used; fourth, we discuss, our pilot experiments and discuss results involving blink pattern recognition; fifth, we present potential applications utilizing this technology; and last we present related work, future work, and our conclusions.

2 Blinking as a Biometric

To motivate why blinking is a potentially interesting biometric, let us explore the implementation of a door-locking mechanism used in a research lab, which is controlled by biometric indicators. In addition to securing the door, the interface should not require the use of hands to unlock it, as people entering may have their hands full of equipment. Face recognition would be a good candidate for such an environment. Due to its reliance on just a camera, face recognition provides an unobtrusive method of biometric verification. Assuming the camera is mounted above the door, the user would not be required to interact with any additional devices. However, face recognition alone would not be sufficient to provide access to the lab door. Face recognition systems can be deceived by using photographs (Ross et al., 2001). An unauthorized person could show a photograph of an authorized person to the system's camera to gain access.

As mentioned previously, adding additional biometric indicators can combat fraudulent entry. Since the security system requires hands-free access, indicators such as hand geometry and fingerprints are not an option. Scanning the iris of a person's eye is one of the most accurate biometric indicators. However, interaction with the iris scanner and first time enrollment of a person's iris image into the database is tedious and time consuming (Wang et al., 2003). Instead of interacting with additional scanners, blink recognition can be performed using the camera already in place for the face recognition system.

Blinking is a biological function that can be both voluntary and involuntary. While some research has explored the use of involuntary blinks as a biometric indicator (Heishman, Duric & Wechsler, 2004), this work focuses on the use of voluntary blinking as a form of personal identification. Voluntary blinking as a behavioral biometric indicator measures how the eye region changes as individuals blink specific patterns. The patterns an individual blinks can be thought of as blinking to the cadence of songs, where the number of blinks and their relative spacing in the pattern is dictated by the rhythm of the song. This is illustrated in top half of Figure 1. In this image we see the first three measures of the holiday song "Frosty the Snowman" with indications of when to blink underneath. There are five blinks in this pattern with spaces of varying length between each blink. The change in notes is the indication of when to blink. The bottom half of Figure 1 is an example of another pattern with five blinks. The two patterns are distinguished by the difference of the temporal relationships between the blinks. In other words, given the same tempo, the blinks occur at different times in each pattern.

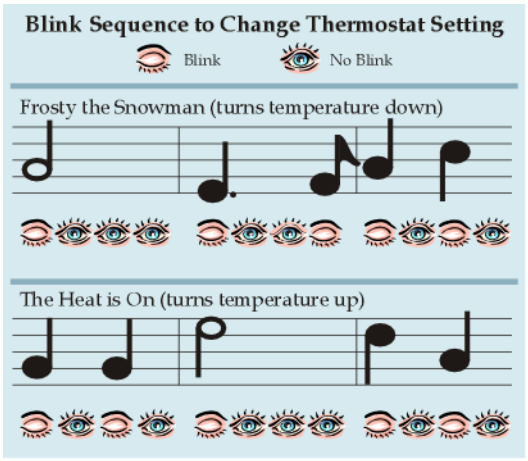

Figure 1: Two song-based blink sequences, each with five blinks. While each sequence has the same number of blinks, the temporal relationship allows two patterns to be distinguishable

Voluntary blinking as an identity indicator is a blending of biometric and knowledge based security. Each person selects a meaningful pattern and enrolls it into the system. Associating the pattern with a song may help the user remember the pattern. For identification systems, the blinking cadence chosen, in essence, is similar to a Personal Identification Number (PIN) – which associates identity with numerical patterns. However, unlike PINs, identifying the user with blink-based biometrics requires more information than just knowing the correct pattern. The blinked pattern must also be performed correctly. Correct performance of the pattern is dependent on the time between blinks, how long the eye is held closed at each blink, and other physical changes the eye undergoes (which are specific to the user) while blinking. We refer to these characteristics as a person's *blinkprint* – a "blinking fingerprint" – which is used to perform identification. Figure 2 shows a visualization of the difference between two users blinking the same pattern.

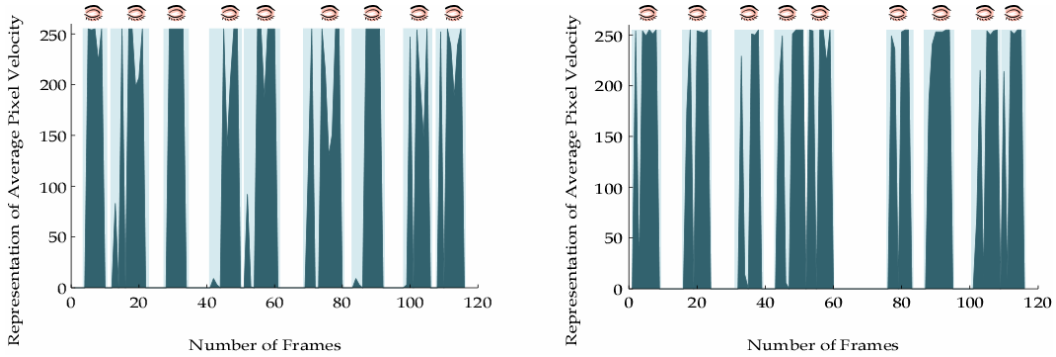

Figure 2: Visual representation of two distinct blinkprints (dark blue). Each blinkprint represents the same nine-blink pattern lasting 4.2 seconds performed by two different people. (the duration of each blink is shaded light blue)

Our previous work showed that if nine people blink exactly the same pattern, we could identify individuals with 82.02% accuracy. This accuracy increases to 95.57% if we use an identity ranking scheme, where the correct identity must be ranked as one of the top three choices. Blinking, therefore, can be a potential verification step for other identifying measures, similar to how a PIN helps verify the identity of a person using an Automatic Teller Machine card – while any person can be in possession of someone else's ATM card, only the owner of the card knows the PIN. Likewise, while anyone can blink someone else's pattern, only the person associated with the pattern will have the correct blinkprint .

Returning to our earlier example of the secure laboratory door, we will address the advantages that blinking provides. First, blinking can be used to localize the position of a person's face (Betke, Mullally, & Magee, 2000), allowing the system to be less obtrusive to the user. Blinking allows the use of face recognition without requiring the user to manually align himself with the camera. Second, as a behavioral-based biometric indicator, blinking

eliminates the possibility of compromising the system using static images. Possessing a photograph of an authorized individual is no longer sufficient to fool the system. With this system in place, a perpetrator attempting to spoof the system with a photograph must find a way to animate the eyes in the picture. Third, unlike a numeric code entered on a keypad, a blinking sequence cannot be observed by glancing over the person's shoulder. In order to observe the sequence, an observer must be positioned in the blinking person's line-of-sight, which will be noticed by the blinker. Fourth, assuming that a person's blink sequence was somehow compromised and that a perpetrator possessed a mask of that person capable of deceiving the face recognition, it is unlikely that the perpetrator would be able to reproduce the blinked sequence in such a way that it matched the blinkprint of the authorized user. The blinkprint of a sequence provides a biometric check performed against the pattern of the sequence – both the blink sequence and the blinkprint must agree on the identity of the user. Thus, blinking as a biometric provides verification for the identity classification produced by the other indicators. In our mask example, the two indicators (face recognition and blinked-pattern recognition) would be fooled by the mask and the blinked pattern. The blinkprint, however, would not agree with the classification of the other indicators and the perpetrator would be denied access.

An unauthorized person could potentially compromise the system by installing a stealth camera with a viewpoint similar to the camera over the door that would record an authorized user's blink pattern. However, in this scenario, the unauthorized person would still have to find a way to play back the video from the stealth camera in such a manner that interlacing was not visible from the door's camera. In addition, other video parameters, such as frame rate, would also have to be known and adjusted to match that of the security system's camera.

3 Enrollment Procedure

The previous example illustrates how blinking, used in combination with face recognition, has the potential to provide an unobtrusive identification system. The only interaction required by the user is to pause before the door, look in the direction of the camera and blink. We have developed an enrollment system interface designed to facilitate ease of use and help improve the accuracy of the system. The system provides an interactive enrollment process that allows a user to gauge how unique her blinked cadence is when compared to the patterns enrolled by other users. In essence this is a measure of how likely the system is to misclassify her pattern. If her pattern is too similar to those already enrolled, she is prompted to enroll a different pattern (we will explain later how this does not compromise security). While enrolling her pattern, the system also provides the user with an indication of how much variation is occurring within her enrollment. In other words, it shows how consistently the user is blinking her pattern. If the variation is too high, then the user has picked a pattern she is incapable of reproducing and the system prompts her to select a new pattern.

The specific hardware requirements and algorithmic details of the process will be discussed in the next section. To understand how the process works from a user's standpoint, it is sufficient to know that there is a camera and a monitor to provide visual feedback to the user. We will now present the details of our interface which is currently designed for research purposes only. The interface is designed to facilitate our research of patterned blink recognition and currently violates many human-computer interaction design principles. The design of the interface for every-day use will be the focus of future research. The research interface is composed of three parts: one section shows video feedback, and two output windows show statistical data and the state of the user enrollment process. Figure 3 is a screen shot of the research interface. When the user first approaches the system, the video window, which displays the image seen by the camera, is grayed. When the user moves close enough to the camera to trigger detection, the video window displays video from the camera with a superimposed box around the user's head. A square in the upper left-hand corner of the video window indicates when the system has successfully detected the user's face and eyes. A yellow square indicates that two or three quick blinks are required to assist the detection. This will often be accompanied by a mismatch between the superimposed square and the actual location of the user's head. A green square indicates that the system has successfully detected the face and eyes, and is ready to enroll a blinked pattern.

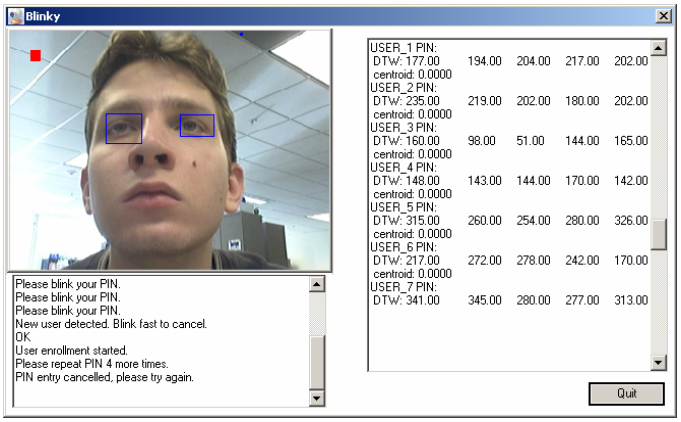

Figure 3: Interface for blink—based enrollment system. The right panel provides intraclass and interclass variability statistics. The left panel displays live-video feedback and the state of the enrollment process.

A user is required to repeat their chosen pattern a few times during the enrollment process. The number of repetitions required is dependent on how similar the pattern is to the existing patterns of other users and how consistently the user blinks the pattern. After each repetition of their pattern, the statistical window displays an indication of the intraclass variability (how consistent the pattern is within the user's own examples) and the interclass variability (the dissimilarity of the user's pattern compared to other enrolled patterns). It is important to note that the indication of interclass variability does not allow one user to feign another user. Similarity with other enrolled patterns can be a result of similarity in blinked cadence (ignoring the biometric check) or it can be a failure of the biometric information to clearly distinguish between individuals (ignoring the similarity of the pattern). In the case of aborting an enrollment due to low interclass variability, the user is given no indication which of the two cases caused the enrollment failure, only that she must select a different cadence.

4 Blink Detection

In order to test our enrollment system in real-world settings, we have extended our previous prototype video capture system (Westeyn & Starner, 2004) to include continuous, real-time blink detection and alignment-free capture. Our system also allows gradual drift of the user's head while blinking. The minimal hardware required by our system is a low-resolution camera, a standard computer (Pentium 4 1GHz with 512MB RAM) and a component providing user feedback (e.g. a monitor).With this equipment we are capable of performing all of the necessary processing for enrollment at 30 frames per second. The remainder of this section discusses how the enrollment system detects the blinked sequence performed by the user, the presence of a user and the location of the user's eyes.

When blinking song-based sequences, patterns of rhythm are indicated not by how long the eye is kept closed, but by the duration of time the eye remains open between blinks. This allows users to blink rhythms using a blink that is natural for them. Therefore, during a blink, the eyelid is always in motion. We use the well-known technique of optical flow (Lucas & Kanade, 1981) to detect frames of video where this motion occurs. Using optical flow helps to provide robustness by producing features that remain somewhat consistent across varying environments.

Optical flow represents the movement of a pixel as a velocity vector expressing the magnitude (ρ) and direction of change (θ) from a previous frame. Although some head movement is inevitable, the motion of blinking will be much more rapid in comparison. This fact allows us to filter out pixels with low velocities, keeping only those associated with blinking. The mean velocity of the pixels for each frame is calculated, then used to determine whether that frame was part of a blink. Once data for an entire pattern is collected, the number of blinks, spaces and the duration of each are calculated.

Betke et al. showed that the motion of the eyelids is sufficient for detection of the eyes' location in video (Betke et al., 2000). By observing large changes in optical flow and subsequent pairings of the highest velocities, we have developed a method that allows any user to approach the system and begin blinking, regardless of their position in the camera image. In our system, a user provides two rapid blinks to allow the system to identify when a user has

approached the system and approximately where the user's eyes are located relative to the camera. While less than ideal, the calibration step also allows successful blink detection even when motion is occurring behind the user. Thus the system will not accidentally interpret this background motion as a blink. It should be noted that the user is not required to keep his head perfectly aligned after the calibration process. The system can tolerate slow drifts in the position of the user's head.

5 Blinked Pattern Classification and Blinkprint Verification

This section will describe how the enrollment system produces its intraclass and interclass variability measure when registering a new pattern. A sequence of blinks is identified by the temporally ordered collection of interval lengths between blinks. This encodes the number of blinks and the duration (in frames of video) of the length between the blinks. For example, the string "6,4,12,10,12,32" would represent the 7-blink sequence for "The Star Spangled Banner." Because there are six elements in the string we can infer that there are seven blinks in this pattern. Inspection of the relative sizes of the intervals demonstrates that the first three blinks in the sequence (blink, 6 frames, blink, 4 frames, blink) occurred rapidly compared to the rest of the sequence which was more evenly spaced.

After a sequence is detected, it is encoded according to the process described above. An observed sequence of blinks is mapped onto a person's identity by using a K-nearest neighbor (KNN) classification scheme – the identity of a pattern is determined by the classification of the K enrollment patterns it most closely matches. Because of temporal variations, the sequences cannot be compared directly. For example, just because a string has seven elements does not mean the string represents "The Star Spangled Banner." Not only must the number of elements match, but the elements of the string must also encode the same relative duration of two quick blinks followed by five longer blinks. A direct comparison would also have difficulty dealing with missing or extraneous blinks that might occur in the sequence as a result of system error, user error or both. If performed quickly enough, two blinks can often appear as one blink to the segmentation algorithm.

Dynamic time warping (DTW) (Darrell & Pentland, 1993) allows two temporal sequences to be compared, and can be used as a metric for nearest neighbor classification. DTW provides a measure of similarity between two sequences which factors in both the length of the sequence and the temporal relation of the elements in the sequence. In order for a match to score well, it must have the same number of blinks and the same relative duration of space between each blink in the sequence. DTW is also flexible enough to cope with missing or extraneous blinks included into the sequence. Assuming that the relative timing is intact, until the extraneous blink occurs, the score will only be slightly affected. The correct sequence plus or minus a blink should still match better than other sequences in the database.

Once the similarity of the enrolled pattern to other patterns has been determined, the system checks the biometric similarity of the pattern. The optical flow features used to detect a blink also encode the blinkprint information. For example, rapidly changing pixels can be a result of shape changes in the eyelid or other physical movements of the eyelid as a blink is performed. Our previous work (Westeyn & Starner, 2004) suggested that the best recognition rates occur when the optical flow features from a person's enrollment examples are modeled as a Gaussian and used to provide a maximum likelihood estimate for subsequent examples. This means that each person is represented by a probabilistic function over optical flow features, ρ and θ, representing the intrinsic characteristics of how they blink. Each time a person enrolls a pattern, the similarity measure with respect to another blinkprint is a measure of the likelihood that the probabilistic function representing the blinkprint will generate those same features.

6 Enrollment Experiments and Discussion

There are two key aspects to the system we have described. The first is the ability to accurately detect and recognize blinked patterns. The second is the enrollment rejection process where the system assists the user in selecting patterns that are easily distinguished by the system. A enrollment pattern is rejected when it has a high degree of similarity with one or more previously enrolled patterns. The pattern could have a timing similar to a previous pattern (e.g. two people blinking the same song cadence) or the blinkprint could be too similar to that of another user. However, in a security system, the user would simply be told that his pattern was too close to one currently in the database. Thus, the user would remain unaware of the blinkprint concept.

A new blink pattern could also be rejected if the user could not reliably reproduce it. As part of the enrollment process, the user would be required to repeat the pattern a number of times to try to capture the variability in his performances. If the variability is high enough to suggest confusion with other users, the pattern should be rejected and the user asked to choose another. Alternatively, a supervisor could make suggestions on how to blink a given pattern more consistently.

The enrollment process, and this method in general, assumes that the variability of a given user's blink pattern and his blinkprint will be small compared to the differences with other users' patterns and blinkprint. Here, we begin to examine this assumption.

In our pilot experiments we manipulated both the number of examples required for enrollment and the similarity metric used for classification (and enrollment rejection). Manipulating the number of enrollment samples should illuminate the effects of intrapersonal variability on the recognition rates. Different similarity metrics were explored for two reasons. First, we wanted to determine empirically which parameters resulted in the highest accuracy. Second, we wanted to determine the correct thresholds for the enrollment process. In other words, how much similarity to other patterns is tolerated before a pattern is rejected. Unfortunately our experiments were limited (even for a pilot study) by the small number of available participants. We recruited 4 participants for our studies but were forced to remove one participants' data due to an error in experimental procedure. For all experiments, each user had to enroll the same blink-based cadence into the system – an 11-blink pattern to the cadence of Jingle Bells. The users then attempted three "logins" to the system using their blink pattern. These logins were attempted 30 minutes after enrollment. This exercise allowed us to gauge the recognition accuracy of the system and determine if a person could repeat their pattern after some time had elapsed.

Table 1: Intraclass standard deviation

	5 examples	10 examples
Participant 1	3.50	4.20
Participant 2	6.71	10.2
Participant 3	4.11	24.0

Table 1 shows the standard deviation of the distances between the login attempt patterns to the enrolled patterns. The distances between an enrollment set of five patterns and a set of ten patterns are compared. For participants 2 and 3 the variability increased between the 5 and 10 example enrollments. Participant 3's results showed a significant increase in spread. Yet, participant 1 had little change. Perhaps participant 1 had more musical experience and could produce the eyeblink timings more precisely than participants 2 and 3. Or perhaps participant 3 would require more practice before he could reduce the variability of his patterns. One could imagine changing the enrollment system such that the user would be required to practice until the distance metric showed that he could blink his pattern consistently. However, would a user tolerate this process? Another option is to collect data, both at enrollment and every time the system is used, until a good sampling of the eyeblink pattern's distribution is obtained. As long as the pattern and the blinkprint could still be distinguished from other users adequately, this approach could be used. The next experiment begins to examine this possibility.

Table 2: Average accuracy for an enrollment set of 5 examples

	K=1	K=5	Biometric Verification
Participant 1	100%	100%	66%
Participant 2	100%	66%	100%
Participant 3	100%	100%	66%

We measured the average accuracy of the system given the three login attempts of each of the three users.
Due to the large variance of the 10 example enrollment for participant 3, we decided to perform classification using only the 5 example enrollment data. For each experiment, we determined the recognition rate using the K-nearest neighbors among the enrollment examples (for K=1 and K=5). We also determined if the biometric verification agreed or disagreed with the k-nearest classification. Table 2 reports the average accuracy for three trials of this experiment, and Table 3 is the corresponding confusion matrix.

In this study, accuracy was highest when basing classification off of the closest enrollment example (K=1). However, this type of classification is often sensitive to noise or outliers during enrollment. We hypothesize that classification using K=5 will perform better for a larger set of participants. The system makes an error when K=5 is used for KNN, mistaking participant 2 for participant 1. However, the confusion matrix for the biometric verification indicates that it recognized all of participant 2's trials correctly. Thus, if the KNN step is combined with the verification step, the system can avoid this error.

Table 3: Confusion matrices for user identification with KNN and Biometric Verification methods. (vertical axis is true identity, horizontal axis is classified identity)

K=1	P1	P2	P3
P1	3	0	0
P2	0	3	0
P3	0	0	3

K=5	P1	P2	P3
P1	3	0	0
P2	1	2	0
P3	0	0	3

Bio.	P1	P2	P3
P1	2	0	1
P2	0	3	0
P3	0	1	2

While the intraclass variability for a given eyeblink pattern may seem high for individual users, the results from the preliminary experiments above show promise for recognition. Even with each participant blinking the same song cadence, the system could distinguish between patterns surprisingly well. Given that our previous work (Westeyn & Starner, 2004) showed over 99% accuracy when distinguishing between 10 different patterns from one user, blink patterns seem very separable in general. In addition, the verification stage seems appropriate for the task. However, a large test with many users is needed to further characterize the system. For this test, providing users with more training during enrollment may reduce their variability and provide a better sense of how the system would behave if deployed in a real scenarios.

7 Potential Applications

In our previous work (Westeyn & Starner, 2004) we demonstrated the potential of song-based blinking as an assistive technology interface. Research suggests that blinking patterns can be an effective form of communication for a certain population of the disabled community (Cook & Hussey, 1984). For example, in the United States two percent of brainstem stroke victims (2,500 people per year) survive paralyzed in the locked-in state (Barnett, Mohr, Stein & Yatsu, 1992). Patients in this condition, known as Locked-in Syndrome, suffer from complete paralysis of all voluntary muscles except for those that control eye movement. People suffering from such crippling diseases still possess the cognitive abilities for communication and other functions, but they lack a method for interacting with the people and environment surrounding them.

Our previous work suggested that such individuals could potentially interact with their environment by using song-based blink patterns to control devices. To help aid in recall, the songs selected should have some meaningful association to the device it interfaces. For example, a user might blink to the cadence of "Frosty the Snowman" to lower the temperature setting of the thermostat (see Figure 1). For some users, blinking could be their only method of communication – errors in recognition could cause considerable frustration for the user. These users might be willing to sacrifice pattern flexibility for accuracy.

Our enrollment system is designed to provide feedback to help choose distinguishable patterns. Currently when an individual enrolls a new pattern into the system, the system determines the similarity of the new pattern to patterns existing in the current database. If the pattern matches too closely with existing patterns (or blinkprint), the system can recommend that the user enroll a different pattern. This process helps the system to remain more robust by avoiding potential errors in classification. Removing the biometric verification process of the enrollment system allows it to be used as an intelligent enrollment for blink-based device interfaces. The modified version of our enrollment system could play an important part in the development of usable, blink-based interfaces for assistive technology.

8 Related Work

The system presented in this paper shares similar research interests with blink-based interfaces, eye-related biometrics, and multimodal biometric systems. In 2001, Grauman et al. (Grauman, Betke, Gips & Bradski, 2001) constructed an assistive technology system that accurately tracks the eyes and measures the duration of eye blinks. Their system uses image correlation with eye templates and is capable of continuously tracking eyes and detecting blinks in real time (27 to 29 fps). The correlation between the image and templates of open and closed eyes distinguishes between natural "short" eye blinks and voluntary "long" eye blinks, allowing the use of deliberate blinks to trigger a mouse click. Their system requires no special training and can continually detect blinks in a desk environment setting with an accuracy of 96.5%. The system can correctly classify detected blinks as natural or voluntary blinks 93% of the time.

Recent work (Heishman et al., 2004) examined eye related biometrics for the identification and classification of specific affective and cognitive states. They explored features of the eye region (irises, pupils, eyelids, eye folds, eyebrows, and blink characteristics) and determined the information content each provided for capturing the cognitive state of the user. The information provided by these features is positively correlated to both the attention level of the user and the level of fatigue experienced by the user. This observation is also supported by studies conducted by the FAA (Stern, Boyer, & Schroeder, 1994) on pilots and air traffic controllers. Their results suggest that the duration and frequency of involuntary blinks are related to the level of fatigue experienced by the individual.

Ross et al. (Ross et al., 2001) explore the fusion of multimodal biometric indicators. They describe a framework for combining data from multiple biometric indicators, such as person's face, hand geometry, and fingerprints. This combination of indicators can occur on three possible levels: the feature extraction level, the matching level and the decision-making level. Fusion at the feature extraction level combines the feature vectors produced by each biometric sensor into a single feature vector. Fusion at the matching level combines similarity scores from each indicator into an overall score. Fusion at the decision level combines accept/reject decisions from each individual indicator into an overall accept/reject decision. For their experiments, a weighted sum at the matching level yielded the best results. Wang et al. (Wang et al., 2003) explore the fusion of face recognition and iris recognition. They report increased accuracy when fusing information at the feature extraction level instead of using the matching level.

9 Future Work

The work presented in this paper is a pilot experiment designed to understand the parameters and limitations of our current prototype. In the immediate future we plan to conduct a series of experiments to test both the accuracy of the system and the utility of our enrollment procedure. Specifically we would like to confirm that the proposed enrollment procedure helps improve recognition accuracy. We would also like to conduct experiments exploring the saturation point of our enrollment system. In other words, how does the accuracy of the system degrade as the number of users and enrolled patterns increases. We will also explore how difficult is it for a new user to enroll as the size of the database increases. Once the stability of the system has been accessed, we would then like to conduct experiments addressing the usability of the system. We would like to determine its usability both as a component of a larger multimodal biometric system and as an interface for assistive technologies.

In addition to testing the system, we also plan to develop a blinking interface toolkit. Our primary goal with this toolkit is to aid in the development of assistive technology interfaces. This toolkit will provide modules necessary for processing and acquiring blinked patterns. Ideally this will be sensor independent to allow blink-based applications that are not dependent on a camera.

10 Conclusions

We have described a hands-free method of identifying a user through a PIN-like blink pattern based on the cadence of songs. Characteristics unique to each user allows the system to learn a "blinkprint," which helps avoid spoofing of the system by someone who has stolen a user's blink pattern. We also presented an enrollment process that alerts a new user if the blink pattern they use is likely to be confused with other blink patterns in the system. Preliminary experiments suggest that the enrollment process may produce a more accurate system if the feedback given encourages users to be more consistent in their blinking during enrollment.

11 Acknowledgments

Funding was provided in part from the GVU Seed Grant program and by NSF CAREER grant # 0093291.

References

Lucas, B., & Kanade, T. (1981). An iterative image registration technique with an application to stereo vision. In *IJCAI81* (p. 674-679).

Barnett, H., Mohr, J., Stein, B., & Yatsu, F. (1992). *Stroke: Pathophysiology, diagnosis, and management*. Churchill, Livingstone, 2nd Ed.

Betke, M., Mullally, W., & Magee, J. (2000, June). Active detection of eye scleras in real time. In *IEEE CVPR workshop on human modeling, analysis and synthesis.*

Cook, A., & Hussey, S. (1984). *Assistive technologies principles and practice, second edition*. Mosby.

Darrell, T., & Pentland, A. (1993). Space-time gestures. *Proc. Comp. Vis. and Pattern Rec.*, 335-340.

Grauman, K., Betke, M., Gips, J., & Bradski, G. R. (2001, Dec). Communication via eye blinks detection and duration analysis in real time. In *IEEE CVPR.*

Heishman, R., Duric, Z., & Wechsler, H. (2004). Using eye region biometrics to reveal affective and cognitive states. In *Proc. of CVPR workshop on face processing in video (FPIV'04), Washington DC, 2004.*

Lucas, B., & Kanade, T. (1981). An iterative image registration technique with an application to stereo vision. In *Ijcai81* (p.674-679).

Ross, A., Jain, A., & Qian, J. (2001, May). *Information fusion in biometrics* (Tech. Rep. No. MSU-CSE-01-18). East Lansing, Michigan: Department of Computer Science, Michigan State University.

Stern, J., Boyer, D., & Schroeder, D. (1994). *Blink rate as a measure of fatigue* (Tech. Rep. No. DOT/FAA/AM-94/17). FAA Civil Aeromedical Institute.

Wang, Y., Tan, T., & Jain, A. K. (2003). Combining face and iris biometrics for identity verification. In J. Kittler & M. S. Nixon (Eds.), *AVBPA* (Vol. 2688, p. 805-813). Springer.

Westeyn, T., & Starner, T. (2004). Recognizing song-based blink patterns: Applications for restricted and universal access. In *Sixth IEEE international conference on automatic face and gesture* (p. 717-722).