

Automatic recognition of eye blinking in spontaneously occurring behavior

JEFFREY F. COHN

University of Pittsburgh, Pittsburgh, Pennsylvania

JING XIAO and TSUYOSHI MORIYAMA

Carnegie Mellon University, Pittsburgh, Pennsylvania

ZARA AMBADAR

University of Pittsburgh, Pittsburgh, Pennsylvania

and

TAKEO KANADE

Carnegie Mellon University, Pittsburgh, Pennsylvania

Previous research in automatic facial expression recognition has been limited to recognition of gross expression categories (e.g., joy or anger) in posed facial behavior under well-controlled conditions (e.g., frontal pose and minimal out-of-plane head motion). We have developed a system that detects a discrete and important facial action (e.g., eye blinking) in spontaneously occurring facial behavior that has been measured with a nonfrontal pose, moderate out-of-plane head motion, and occlusion. The system recovers three-dimensional motion parameters, stabilizes facial regions, extracts motion and appearance information, and recognizes discrete facial actions in spontaneous facial behavior. We tested the system in video data from a two-person interview. The 10 subjects were ethnically diverse, action units occurred during speech, and out-of-plane motion and occlusion from head motion and glasses were common. The video data were originally collected to answer substantive questions in psychology and represent a substantial challenge to automated action unit recognition. In analysis of blinks, the system achieved 98% accuracy.

Facial expression is one of the most powerful, natural, and immediate means for human beings to communicate their emotions and intentions. Facial displays indicate emotion (Ekman, 1993) and pain (Craig, Hyde, & Patrick, 1991), regulate social behavior (Cohn & Elmore, 1988; Eibl-Eibesfeldt, 1989; Fridlund, 1994), reveal brain function (Ekman, Davidson, & Friesen, 1990; Fox & Davidson, 1988) and pathology (Katsikitis & Pilowsky, 1988; Rinn, 1984), and signal developmental transitions in infants (Campos, Bertenthal, & Kermoian, 1992; Emde, Gaensbauer, & Harmon, 1976). To make use of the information afforded by facial expression, reliable, valid, and efficient methods of measurement are critical.

Human-observer-based (i.e., manual) methods of coding facial expression are labor intensive, semiquantitative, and difficult to standardize across laboratories or over time. Training is time consuming (approximately 100 h with the most descriptive methods), and coding criteria may drift with time (Bakeman & Gottman, 1986; Martin & Bateson, 1986). Implementing comprehensive systems has been reported to take up to 10 h of coding time per minute of behavior, depending on the comprehensiveness of the system and the density of behavior changes (Ekman, 1982). Such extensive effort discourages standardized measurement and may encourage the use of less specific coding systems with unknown convergent validity (Matias, Cohn, & Ross, 1989). These problems tend to promote the use of smaller sample sizes (of subjects and behavior samples), prolong study completion times, and thus, limit the generalizability of study findings.

Within the past decade, there has been significant effort toward automatic recognition of human facial expression, using computer vision. Several such systems (Essa & Pentland, 1997; Padgett, Cottrell, & Adolphs, 1996; Yacoub & Davis, 1997) have recognized, under controlled conditions, a small set of emotion-specified expressions, such as joy and anger. Others (Bartlett, Hager, Ekman, & Sejnowski, 1999; Lien, Kanade, Cohn, & Li, 2000; Tian,

Portions of the findings were presented at the International Conference on Methods and Techniques in Behavioral Research, Amsterdam, August 2002 and at the International Conference on Pattern Recognition, Quebec City, Canada, August 2002. This research was supported by Grant R01 MH51435 from the National Institute of Mental Health and CIA Contract 2000-A128400-000 to J.F.C. and T.K. We thank Paul Ekman, Mark Frank, Marian Bartlett, Kathleen Miritello, Javier Movellan, Pietro Perona, Terry Sejnowski, and Yaser Yacoub for invaluable assistance, comments, and suggestions. Paul Ekman and Mark Frank generously provided image data. Correspondence should be addressed to J. F. Cohn, Department of Psychology, 4327 Sennott Square, University of Pittsburgh, Pittsburgh, PA 15260 (e-mail: jeffcohn@pitt.edu).

Kanade, & Cohn, 2000, 2001, 2002) have achieved some success in the more difficult task of recognizing facial action units (AUs) of the Facial Action Coding System (FACS: Ekman & Friesen, 1978). AUs are the smallest visibly discriminable changes in facial expression. In a series of studies (Cohn, Zlochower, Lien, & Kanade, 1999; Lien et al., 2000; Tian et al., 2001, 2002), our group has developed a system that recognizes 20 of approximately 30 AUs that have a known anatomic basis and occur most frequently in emotion expression and paralinguistic communication (Kanade, Cohn, & Tian, 2000; Sayette, Cohn, Wertz, Perrott, & Parrott, 2001). AUs are recognized whether they occur alone or in combinations. The ability to recognize AUs, whether they occur alone or in combinations, is an important feature, because AUs may occur in thousands of combinations (Ekman, 1982) and the appearance and timing of AUs can vary, depending on the AUs with which they co-occur (analogous to coarticulation effects in speech).

A limitation of almost all research to date in automatic facial expression recognition is that it is limited to deliberate facial expression recorded under controlled conditions that omit significant head motion and other factors that complicate analysis. Automatic recognition of facial AUs in spontaneously occurring facial behavior presents multiple challenges. These include camera orientation, head motion, occlusion, differences between deliberate and spontaneous expression, and individual differences among subjects. We first will review each of these factors. We then will report one of the first attempts to automatically recognize AUs in spontaneous facial behavior during social interaction, with nonfrontal pose, moderate out-of-plane head motion, and moderate occlusion.

Camera Orientation and Rigid Head Motion

Most research in facial expression recognition has been limited to image sequences in which the face is oriented toward the camera. Although frontal pose is a reasonable assumption when subjects are asked to model facial expressions or talk in front of a camera (as in “talking head” or TV news applications), in most social interaction research, cameras are oriented about 15°–30° to the side, in order not to interfere with face-to-face interaction. With variation in camera orientation, face appearance changes qualitatively. As one example, the relative symmetry of facial features is lost (Liu, Schmidt, Cohn, & Weaver, 2002), and facial features become occluded when camera orientation is no longer frontal. Systems that have been trained on relatively symmetric frontal views may generalize poorly to face images acquired from oblique camera angles.

A related problem is head motion. Moderate to large head motion is common in naturally occurring behavior and often accompanies change in expression. Kraut and Johnson (1979) found that smiling typically occurs while turning toward another person. Camras, Lambrecht, and Michel (1996) found that infant surprise expressions often occur as the infant pitches his or her head back. To control for head motion, previous research has used head-

mounted cameras (Pantic, 2002) or has selected image sequences in which head motion is absent (Bartlett et al., 1999; Donato, Bartlett, Hager, Ekman, & Sejnowski, 1999). When head motion is allowed, it typically is limited to motion that is parallel to the image plane of the camera (i.e., planar head motion; Tian et al., 2001; Yacoob & Davis, 1997). For head motion that is planar or nearly so, an affine or perspective transformation of images (Lien et al., 2000) can align images so that face position, size, and orientation are kept relatively constant across subjects and these factors do not interfere significantly with feature extraction. The problem occurs with out-of-plane motion of more than about 5°, which is typical in spontaneous facial behavior. Out-of-plane head motion is a particular challenge for computer vision systems, because the face looks different when the head pitches up or down or turns from side to side (i.e., yaw), which occludes facial features. A computer vision system must be robust to variations in camera orientation and out-of-plane head motion.

Deliberate Versus Spontaneous Facial Expression

Most facial expression data have been collected by asking subjects to perform a series of expressions. These directed facial action tasks typically are more intense and include more appearance changes than do those that occur spontaneously (Zlochower, 2001). They also are likely to differ in appearance and timing from spontaneously occurring behavior (Cohn & Schmidt, 2003). Deliberate and spontaneous facial behaviors are mediated by separate motor pathways—the pyramidal and extra-pyramidal motor tracks, respectively (Rinn, 1984). As a consequence, fine motor control of deliberate facial actions often is inferior to and less symmetric than that which occurs spontaneously. Many people, for instance, are able to raise their outer brows spontaneously while leaving their inner brows at rest; few can perform this action voluntarily. Spontaneous depression of the lip corners (AU 15) and raising and narrowing the inner corners of the brow (AU 1 + 4) are common signs of sadness. Without training, few people can perform these actions deliberately, which, incidentally, is an aid in lie detection (Frank & Ekman, 1997). Spontaneous smiles show a highly consistent relation between duration and amplitude that is characteristic of ballistic motion; for deliberate smiles, these parameters are uncorrelated, to which observers may be particularly sensitive (Cohn & Schmidt, 2003). Differences in the temporal organization of spontaneous and deliberate facial actions are particularly important to automatic facial expression analysis, in that many pattern recognition approaches, such as hidden Markov modeling, are highly dependent on the timing of appearance change.

Individual Differences Among Subjects

Face shape, texture, color, and facial and scalp hair vary with sex, ethnic background, and age (Bruce & Young, 1998). Infants, for instance, have smoother, less textured skin and often lack facial hair in the brows or

scalp. The eye opening and contrast between iris and sclera differ markedly between Asians and Northern Europeans, which may affect the robustness of eye tracking and of facial feature analysis more generally. Beards, eyeglasses, or jewelry may obscure facial features. Such individual differences in appearance can have important consequences for face analysis. Few attempts to study their influence exist. In our experience (Cohn, Tian, & Forbes, 2000), algorithms for optical flow and high-gradient component detection that have been optimized for young adults perform less well when used in infants. The reduced texture of infants' skin, their increased fatty tissue, juvenile facial conformation, and lack of transient furrows may all have contributed to the differences observed in face analysis between infants and adults.

In addition to individual differences in appearance, there are individual differences in expressiveness, which refers to the degree of facial plasticity, relative intensity of facial expression, bias for particular facial expressions, and overall rate of expression. Individual differences in these characteristics are well established and are important aspects of individual identity (Cohn, Schmidt, Gross, & Ekman, 2002; Moore, Cohn, & Campbell, 1997; Schmidt & Cohn, 2001). An extreme example of variability in expressiveness occurs in individuals who have incurred damage to the facial nerve or central nervous system (Van Swearingen, Cohn, & Bajaj-Luthra, 1999). To develop algorithms that are robust to individual differences in facial features and expressiveness, it is essential to include subjects of varying ethnic background, age, and sex and to allow for facial hair, jewelry, or eyeglasses that may obscure or hide facial features.

We report one of the first attempts to automatically recognize AUs—in particular, eye blinking—in spontaneous facial behavior during social interaction with non-frontal pose, moderate out-of-plane head motion, and moderate occlusion in diverse subjects. The image data were collected in order to investigate a socially and psychologically important topic, deception, rather than to test facial expression algorithms. We focus here on eye blinking. Measurement of blinking is important in several fields, including neurology, physiology, and psychology. Blink rate varies with physiological and emotional arousal, cognitive effort, and deception (Ekman, 2001; Holland & Tarlow, 1972; Karson, 1988).

The face analysis system recovers three-dimensional (3-D) motion parameters, stabilizes facial regions, extracts motion and appearance information, and recognizes AUs in spontaneous facial behavior. Manual processing is limited to marking several feature points in the initial image of the stabilized image sequence. All other processing is automatic. In an initial test, reported below, the system recognized blinks with 98% accuracy.

METHOD

Database

We used video data from a study of deception by Frank and Ekman (1997). The subjects were 20 young adult men. Data from 10 were

available for analysis. Seven of the 10 were Euro-American, 2 were African-American, and 1 was Asian. Two wore glasses. The subjects either lied or told the truth about whether they had stolen a large sum of money. Prior to stealing or not stealing the money, they were informed that they could earn as much as \$50 if successful in perpetuating the deception and could anticipate relatively severe punishment if they failed. By providing strong rewards and punishments, the manipulation afforded ecological validity for deception and for truth-telling conditions.

The subjects were video recorded using a single S-Video camera. Head orientation to the camera was oblique, and out-of-plane head motion was common. The tapes were digitized into 640×480 pixel arrays with 16-bit color resolution. A certified FACS coder at Rutgers University, under the supervision of Dr. Frank, manually FACS-coded start and stop times for all AUs in 1 min of facial behavior in the first 10 subjects. Certified FACS coders from the University of Pittsburgh confirmed all the coding.

In this report, we focus on automatic analysis of blinks (AU 45 in FACS). We included for analysis all instances of blink (AU 45) for which two independent teams of certified FACS coders agreed; 95% of those examined (167 blinks) met this criterion and were included in the analyses. The average number of blinks per subject was 16.70, with a standard deviation of 10.86. We also include the few instances of eyelid *flutter* that occurred, defined as two or more rapidly repeating blinks (AU 45) with only partial eye opening (AU 42 or AU 41) between them. In flutter, the modal interval between blinks was one frame, with a maximum of two frames. These instances of flutter occurred in 3 subjects, with 1 subject accounting for 12 of the 14 that occurred. For each subject, we included an equal number of nonblink intervals of equal duration, for comparison.

Overview of Face Analysis System

Figure 1 depicts an overview of the face analysis system (Automated Face Analysis, Version 3) used for automatic recognition of blinks (FACS AU 45) and flutter. A digitized image sequence is input to the system. The face region is delineated in the initial frame either manually or by using a face detector (Rowley, Baluja, & Kanade, 1998). The image in which the head is most upright is chosen as the reference image. Head motion (6 degrees of freedom [dfs]) is recovered automatically. With the recovered motion parameters, the face region is stabilized—that is, warped to a view of common orientation. Facial features are extracted in the image sequence, and AUs are recognized. The present study is limited to analysis of the eye region. The modules for feature extraction and analysis of other face regions (e.g., forehead and brows) are not considered.

Automatic Recovery of 3-D Head Motion and Stabilization of Eye Region

To estimate 3-D head motion, one option is to use an anatomically based face model in which the exact proportions of facial features are represented (DeCarlo & Metaxas, 1996; Essa & Pentland, 1997). To work well, such anatomic models require fitting a large number of parameters that are dependent on the exact shape of the individual face, which typically is unknown. These parameters must be precisely estimated throughout an image sequence. When precise initialization and tracking are not available, better recovery may be achieved by using a simpler 3-D geometric model. In our approach, we use a cylindrical head model to estimate the 6 dfs of head motion, whose parameters are horizontal and vertical position, distance to the camera (i.e., scale), pitch, yaw, and roll.

A cylindrical model is fit to the initial face region, and the face image is cropped and "painted" onto the cylinder as the initial appearance template. For any given frame, the template is the head image in the previous frame that is projected onto the cylindrical model. To estimate head motion, we first warp the head image to the cylindrical model with the pose being that of the template; that is, we assume that the pose has remained unchanged from the template

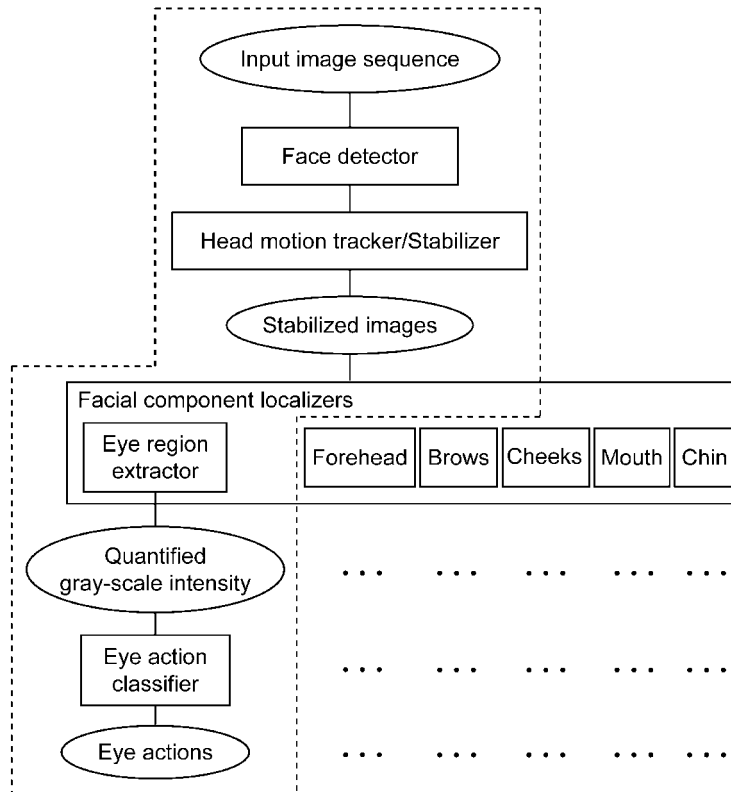


Figure 1. Overview of Automated Face Analysis, Version 3.

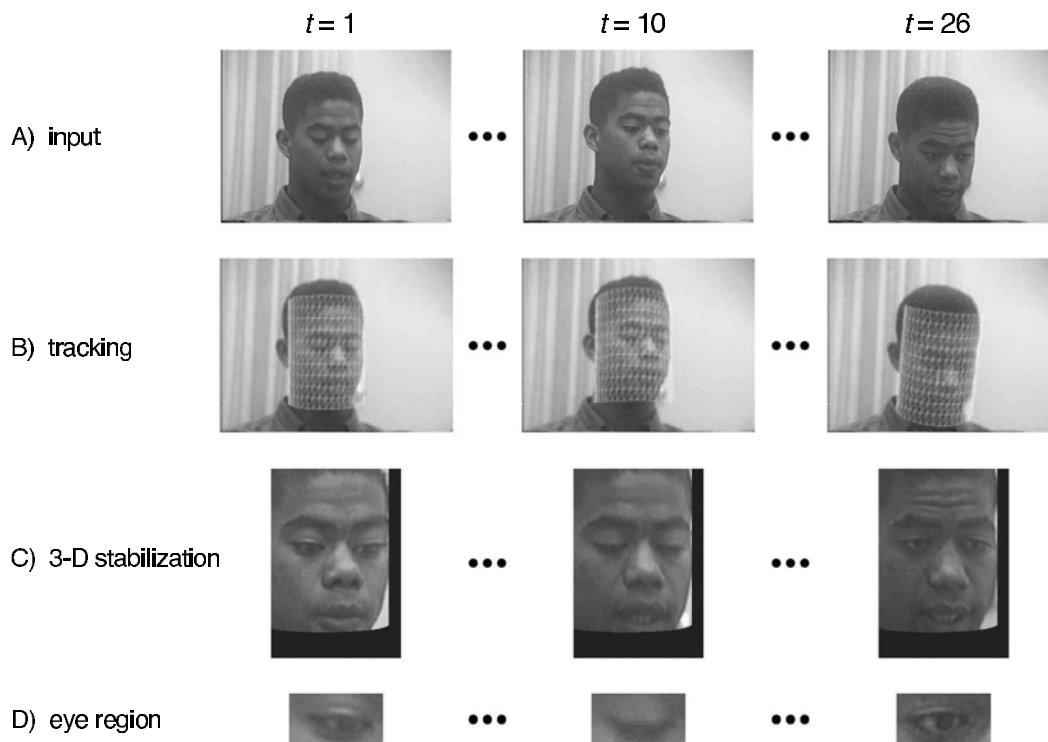


Figure 2. Automatic recovery of three-dimensional (3-D) head motion and image stabilization. (A) Frames 1, 10, and 26 from the original image sequence. (B) Automatic face tracking in the corresponding frames. (C) Stabilized face images. (D) Localized eye regions.




Figure 3. Upper and lower portions of the eye region in Frames 1, 5, 6, 7, and 9 from a stabilized image sequence in which the eye changes from open to closed.

to the current frame. We then compute the difference between the warped image and the template, to provide the initial estimate of pose. To obtain the final estimate of pose, we iterate this process to further refine the estimate by using a model-based optical flow algorithm (Lucas & Kanade, 1981; Xiao, Kanade, & Cohn, 2002). This algorithm is implemented with an iterative hierarchical image pyramid (Poelman, 1995) so that rapid and large head motion may be tracked. Although the algorithm assumes that lighting remains unchanged from one frame to the next, we have found that the algorithm is relatively robust to the typical variations in lighting found in indoor settings.

Major potential sources of error in estimating head motion include nonrigid motion (e.g., facial expression) and occlusion (e.g., a hand moving in front of the face). When they occur, some pixels in the template may change in the processed image or even disappear. We do not want these pixels to contribute to motion estimation. When estimating optical flow from one frame to the next, these pixels will have relatively large error. To minimize their influence, pixels with large error are given less weight in estimating motion between frames, so that they contribute less than do other pixels. The specific technique that we use is iteratively reweighted least squares (Black, 1992). In this way, the potential influence of nonrigid motion and occlusion are minimized in estimating rigid head motion.

Head templates change while tracking occurs. Once head pose is estimated in a new frame, the region facing the camera is extracted as the new template. Because head poses are recovered, using templates that are constantly updated, and the pose estimated for the current frame is used in estimating the pose in the next frame, errors would accumulate unless otherwise prevented. To solve this problem, the first frame and the initial head pose are stored as a reference. When the estimated pose for the new frame is close to the initial one, the system rectifies the current pose estimate by registering this frame with a reference one. The reregistration prevents errors from accumulating and enables the system to recover head pose when the face reappears after occlusion, such as when the head moves momentarily out of the camera's view. By reregistering the face image, the system can run indefinitely.

The system was tested successfully in image sequences that include maximum pitch and yaw as large as 40° and 75° , respectively, and a time duration of up to 20 min (Xiao et al., 2002). The precision of recovered motion was evaluated with respect to the ground truth obtained by a precise position and orientation measurement device with markers attached to the head and found to be highly consistent (e.g., for 75° yaw, absolute error averaged 3.86° ; for details, see Xiao et al., 2002). Although a head shape is not actually a cylinder, a cylinder model is found to be adequate and, indeed, contributes to system stability and robustness.

An example of system output can be seen in Figure 2. From the input image sequence (Figure 2A), the head is tracked, and its pose is recovered, as described above (Figure 2B). Once the head pose is recovered, we can stabilize the face region by transforming the image to a common orientation (Figure 2C) and then can localize a region of interest, which, in the present study, is the eye region (Figure 2D), as will be described below.

Eye Action Classification

The eye region consists of the iris, sclera, upper and lower eyelids, and eyelashes. If we divide the eye region into upper and lower portions (Figure 3), the gray-scale intensity distribution of the upper and lower portions would change as the eyelid closes and opens during blinking. Gray-scale intensity can range from 0 (black) to 255 (white). When the eye is open, the relatively light gray-scale intensity of sclera is averaged with the darker gray-scale values of the upper eyelashes, pupil, and iris. As the eyelid closes, gray-scale intensity in the upper eye region increases as the eyelashes move into the lower region and the pupil and iris become occluded by the relatively brighter (i.e., higher gray-scale intensity) skin of the now exposed upper eyelid. In pilot testing and in the present study, we found that even in dark-skinned subjects, gray-scale intensity in the upper eye region is greater when the eyes are closed rather than open. We can use this knowledge to automatically track closing and opening of the eye and to recognize blinks from nonblinks once the eye region is stabilized with respect to rigid motion.

Automatic Feature Extraction in the Eye Region

The input face image sequence (Figure 2A) has been automatically processed to obtain the stabilized image sequence (Figure 2C, as described above). We then define the eye region by manually marking four feature points in the first frame of the stabilized image sequence. These feature points in the first frame are the inner and outer eye corners, the center of the upper eyelid, and just below the center of the lower eyelid. The pixel coordinates of these feature points in the first frame then are used to define the upper and lower eye regions in each frame of the stabilized image sequence. Figure 3 shows an example from Frames 1, 5, 6, 7, and 9 of a stabilized image sequence in which the eye changes from open to closed. Note




Figure 4. Blink detection algorithm. The number and duration of nonblinks, blinks, and flutters are detected by counting the number of peaks (N_p) and crossings (N_c) in the gray-scale intensity curves. If $N_p = 0$, there is a nonblink; if $N_p = 1$ and $N_c \neq 0$, a blink; if $N_p \geq 2$ and $N_c \neq 0$, a flutter.

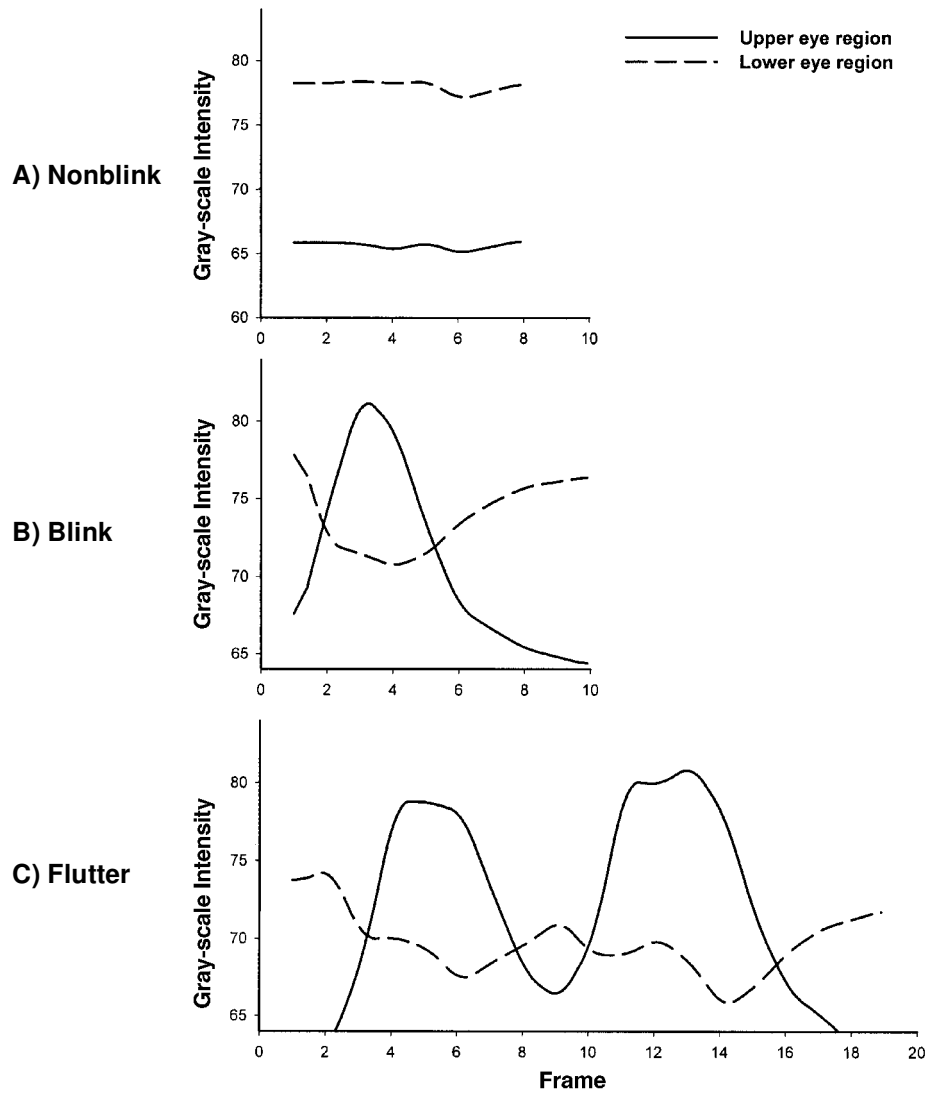


Figure 5. Examples of gray-scale intensity curves for nonblink, blink, and flutter.

that as the eye closes, the gray-scale intensity of the upper eye region becomes lighter as the upper eyelid progressively covers the iris and pupil and the eyelashes on the upper eyelid move into the lower eye region.

For now, we treat only the right eye (image left). The classification categories of eye actions are blink, flutter, and nonblink. For this classification, the average gray-scale intensity is calculated for the upper and the lower eye regions. When the eye is open, mean gray-scale intensity in the upper half is smaller than that in the lower half and reverses when the eye is closed. When mean intensities for the upper and the lower regions are plotted over time, they cross when an eye closes and opens. The intensity curve for the upper eye peaks when the eye is completely closed (AU 45). By counting the number of crossings (N_c) and the number of peaks (N_p) in the gray-scale intensity curves, we can detect the timing, number, and duration of eye blinking:

- $N_p = 0$ → Nonblink
- $N_p = 1$ and $N_c \neq 0$ → Blink
- and
- $N_p \geq 2$ and $N_c \neq 0$ → Flutter,

where N_p = number of peaks and N_c = number of crossings in the gray-scale intensity curves for the upper and the lower eye regions (Figure 4).

Figure 5 shows examples of the gray-scale intensity curves for nonblink (Figure 5A), blink (Figure 5B), and flutter (Figure 5C).

RESULTS

The accuracy of blink detection was assessed in two ways. First, we compared the sensitivity of the algorithm with quantitative changes in eye closure to that of human judges. Second, we compared the accuracy of automatic blink detection with that of human FACS coders.

To evaluate the sensitivity of the blink detection algorithm to variation in eye opening and closing, we compared automatic and manual (human) processing. Stabilized eye image sequences in which AU 45 (blink) occurred were analyzed using the algorithm described above. Gray-scale intensity of the upper eye region, as

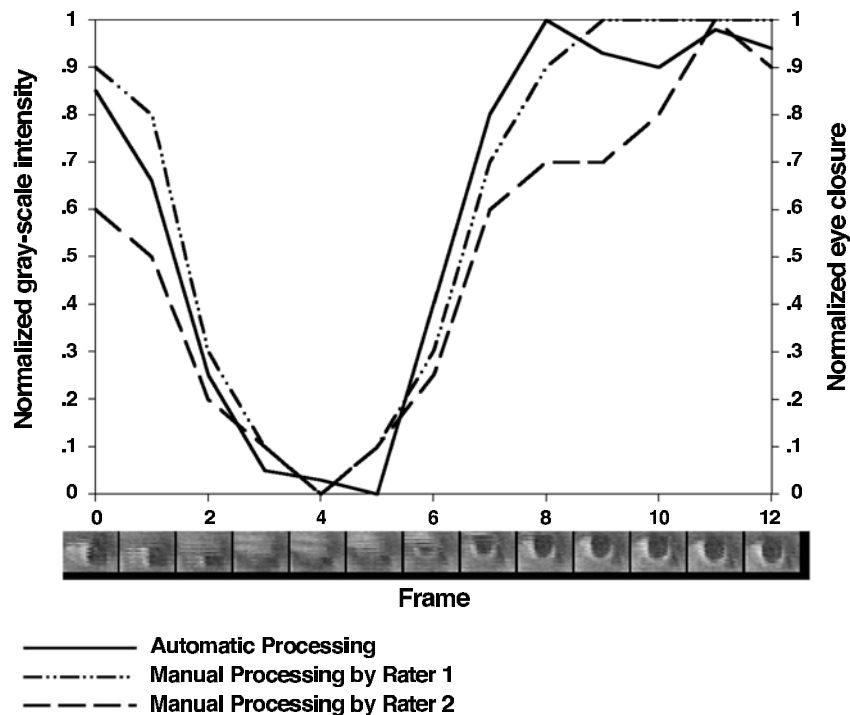


Figure 6. Comparison of manual ordering of eye closure by Rater 1 and Rater 2 and automatic ordering using automatic facial expression analysis.

determined automatically, was normalized over the range of 0 to 1. A normalized intensity of 1 corresponded to the open eye, and 0 corresponded to the closed eye. The digitized images then were randomly sorted. Two researchers, blind to the results of automatic processing, then manually sorted each sequence from eye open to closed to open. They next estimated the degree of eye opening on a scale from 0 (*eye closed*) to 1 (*eye open*). In each of the 10 sequences examined, the manual ratings of eye open by each of the two researchers were highly consistent with automatic measurement of gray-scale intensity. Figure 6 shows the graphical results for one of the 10 sequences examined. The manual ratings by each of the raters were highly consistent with those of the others and with the results of automatic processing of gray-scale intensity. In the example shown, the consistency between Rater 1 and automatic processing ($r = .95$) was comparable to the consistency found in manual ratings between Rater 1 and Rater 2 ($r = .96$). These findings were typical of the close correspondence between manual ratings of eye state and automatic processing.

Table 1 shows the recognition results for blink detection in all the image data, in comparison with the manual FACS coding. The algorithm achieved an overall accuracy of 98%. Six of 14 instances of flutter were incorrectly recognized as single blinks. Rapid transitions from AU 45 to AU 42 to AU 45, in which eye closure remains nearly complete, were occasionally recognized as a sin-

gle blink. Transitions from AU 45 to AU 41 (drooping upper eyelid) to AU 45 were more frequently detected. The measure we used (crossing of average intensities) was not consistently sensitive to the slight change between complete closure (AU 45) and closure that was nearly complete (AU 42). If blink and flutter are combined into a single category (which is common practice among FACS coders), classification accuracy of eye closure and opening was 100%.

DISCUSSION

This study is one of the first to attempt automatic AU recognition in naturally occurring facial behavior. All other work in automated facial expression recognition has been limited to analysis of deliberate facial expressions that have been collected under controlled conditions for purposes of algorithm development and testing. We an-

Table 1
Comparison of Manual FACS Coding and Automatic Recognition

Manual FACS Coding	Automatic Recognition		
	Blink (AU45)	Flutter	Nonblink
Blink (AU 45)	153	0	0
Flutter	6	8	0
Nonblink	0	0	167

Note—Overall agreement = 98% (kappa = .97). Combining blink and flutter agreement = 100%.

alyzed image data from Frank and Ekman (1997), who collected them under naturalistic conditions in the course of psychological research on deception, and not with the intention of automated facial expression analysis. We analyzed spontaneously occurring behavior, rather than posed expressions. The data presented significant challenges in terms of heterogeneity of subjects, brightness, occlusion, pose, out-of-plane head motion, and low intensity of AUs.

To meet these challenges, we developed Automated Face Analysis, Version 3, which automatically estimates 3-D motion parameters, stabilizes face images for analysis, and recognizes facial actions, using a face-component-based approach to feature extraction. We emphasized the aspect of automated analysis of feature extraction, localization, and tracking; manual processing was limited to feature marking in a single initial frame, and all other processing was fully automatic. This focus contrasts with recent efforts by Bartlett, Braathen, Littlewort, Sejnowski, and Movellan (2001) to analyze spontaneous facial behavior that requires manual labeling and registration of each image.

Automated Face Analysis, V.3, successfully recognized blinks from nonblinks for all the examples in the database. It also was able to distinguish flutter, with lower accuracy. We found that automated analysis of facial images still presents significant challenges. Many previously published algorithms, including our own, that worked well for frontal faces and good lighting conditions fail with images under nonfrontal facial poses, full 6 *df* head motions, and ordinary lighting. Precise and reliable extraction and localization of features is the key to the success of automated FACS coding, for facial expression and emotion analysis. The 3-D model-based stabilization technique presented here for stabilizing the arbitrary and unknown head motion is one such example. In the next phase of our research, we will expand the size and diversity of our database of FACS-coded spontaneous facial behavior and increase the number and complexity of AUs that can be recognized automatically in this context. We also will include explicit analysis of the timing of facial actions in relation to context and communicative intention (Cohn & Schmidt, 2003; Schmidt & Cohn, in press). Although challenges remain, these findings support the feasibility of developing and implementing comprehensive, automated facial expression analysis in research on emotion and nonverbal communication.

REFERENCES

- BAKEMAN, R., & GOTTMAN, J. M. (1986). *Observing behavior: An introduction to sequential analysis*. Cambridge: Cambridge University Press.
- BARTLETT, M. S., BRAATHEN, B., LITTLEWORT, G., SEJNOWSKI, T. J., & MOVELLAN, J. R. (2001). *A comparative study of alternative FACS coding algorithms* (Tech Rep. MPLAB-TR-2001-06). La Jolla: University of California at San Diego, Institute for Neural Computation.
- BARTLETT, M. S., HAGER, J. C., EKMAN, P., & SEJNOWSKI, T. J. (1999). Measuring facial expressions by computer image analysis. *Psychophysiology*, **36**, 253-263.
- BLACK, M. (1992). *Robust incremental optical flow*. Unpublished doctoral dissertation, Yale University.
- BRUCE, V., & YOUNG, A. (1998). *In the eye of the beholder: The science of face perception*. Oxford: Oxford University Press.
- CAMPOS, J., BERTENTHAL, B., & KERMOIAN, R. (1992). Early experience and emotional development: The emergence of wariness of heights. *Psychological Science*, **3**, 61-64.
- CAMRAS, L. A., LAMBRECHT, L., & MICHEL, G. F. (1996). Infant "surprise" expressions as coordinative motor structures. *Infant Behavior & Development*, **20**, 183-195.
- COHN, J. F., & ELMORE, M. (1988). Effect of contingent changes in mothers' affective expression on the organization of behavior in 3-month-old infants. *Infant Behavior & Development*, **11**, 493-505.
- COHN, J. F., & SCHMIDT, K. (2003). The timing of facial motion in posed and spontaneous smiles. In *Proceedings of the International Conference on Active Media Technology*, **2**, 57-72.
- COHN, J. F., SCHMIDT, K., GROSS, R., & EKMAN, P. (2002). Individual differences in facial expression: Stability over time, relation to self-reported emotion, and ability to inform person identification. *Proceedings of the International Conference on Multimodal User Interfaces*, **4**, 491-496.
- COHN, J. F., TIAN, Y. L., & FORBES, E. (2000). Detection, tracking, and classification of facial action units in infant and caregiver facial expression [Abstract]. In *Proceedings of the International Conference on Infant Studies*.
- COHN, J. F., ZLOCHOWER, A., LIEN, J., & KANADE, T. (1999). Automated face analysis by feature point tracking has high concurrent validity with manual FACS coding. *Psychophysiology*, **36**, 35-43.
- CRAIG, K. D., HYDE, S. A., & PATRICK, C. J. (1991). Genuine, suppressed and faked facial behavior during exacerbation of chronic low back pain. *Pain*, **46**, 161-171.
- DECARLO, D., & METAXAS, D. [N.] (1996). The integration of optical flow and deformable models with applications to human face shape and motion estimation. *Proceedings of the IEEE Conference on Computer Vision & Pattern Recognition*, **15**, 189-195.
- DONATO, G. L., BARTLETT, M. S., HAGER, J. C., EKMAN, P., & SEJNOWSKI, T. J. (1999). Classifying facial actions. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, **21**, 974-989.
- EIBL-EIBESFELDT, I. (1989). *Human ethology*. Hawthorne, NY: Aldine de Gruyter.
- EKMAN, P. (1982). Methods for measuring facial action. In K. R. Scherer & P. Ekman (Eds.), *Handbook of methods in nonverbal behavior research* (pp. 45-90). Cambridge: Cambridge University Press.
- EKMAN, P. (1993). Facial expression and emotion. *American Psychologist*, **48**, 384-392.
- EKMAN, P. (2001). *Telling lies*. New York: Norton.
- EKMAN, P., DAVIDSON, R. J., & FRIESEN, W. V. (1990). The Duchenne smile: Emotional expression and brain physiology II. *Journal of Personality & Social Psychology*, **58**, 342-353.
- EKMAN, P., & FRIESEN, W. (1978). *Facial action coding system*. Palo Alto, CA: Consulting Psychologists Press.
- EMDE, R. N., GAENSBAUER, T. J., & HARMON, R. J. (1976). *Emotional expression in infancy: A biobehavioral study* (Psychological Issues, Vol. 10, No. 1). New York: International Universities Press.
- ESSA, I., & PENTLAND, A. (1997). Coding, analysis, interpretation and recognition of facial expressions. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, **7**, 757-763.
- FOX, N., & DAVIDSON, R. J. (1988). Patterns of brain electrical activity during facial signs of emotion in ten-month-old infants. *Developmental Psychology*, **24**, 230-236.
- FRANK, M., & EKMAN, P. (1997). The ability to detect deceit generalizes across different types of high-stake lies. *Journal of Personality & Social Psychology*, **72**, 1429-1439.
- FRIDLUND, A. J. (1994). *Human facial expression: An evolutionary view*. San Diego: Academic Press.
- HOLLAND, M. K., & TARLOW, G. (1972). Blinking and mental load. *Psychological Reports*, **31**, 119-127.
- KANADE, T., COHN, J. F., & TIAN, Y. (2000). Comprehensive database for facial expression analysis. *Proceedings of the IEEE Conference on Automatic Face and Gesture Recognition*, **4**, 46-53.
- KARSON, C. N. (1988). Physiology of normal and abnormal blinking. *Advances in Neurology*, **49**, 25-37.
- KATSKITIS, M., & PILOWSKY, I. (1988). A study of facial expression in

- Parkinson's disease using a novel microcomputer-based method. *Journal of Neurology, Neurosurgery, & Psychiatry*, **51**, 362-366.
- KRAUT, R. E., & JOHNSON, R. (1979). Social and emotional messages of smiling: An ethological approach. *Journal of Personality & Social Psychology*, **37**, 1539-1553.
- LIEN, J. J. J., KANADE, T., COHN, J. F., & LI, C. C. (2000). Detection, tracking, and classification of subtle changes in facial expression. *Journal of Robotics & Autonomous Systems*, **31**, 131-146.
- LIU, Y., SCHMIDT, K., COHN, J. F., & WEAVER, R. L. (2002). Human facial asymmetry for expression-invariant facial identification. *Proceedings of the IEEE International Conference on Automatic Face & Gesture Recognition*, **5**, 208-214.
- LUCAS, B. D., & KANADE, T. (1981). An iterative image registration technique with an application in stereovision. *Proceedings of the International Joint Conference on Artificial Intelligence*, **7**, 674-679.
- MARTIN, P., & BATESON, P. (1986). *Measuring behavior: An introductory guide*. Cambridge: Cambridge University Press.
- MATIAS, R., COHN, J. F., & ROSS, S. (1989). A comparison of two systems to code infants' affective expression. *Developmental Psychology*, **25**, 483-489.
- MOORE, G. A., COHN, J. F., & CAMPBELL, S. B. (1997). Mothers' affective behavior with infant siblings: Stability and change. *Developmental Psychology*, **33**, 856-860.
- PADGETT, C., COTTRELL, G. W., & ADOLPHS, B. (1996). Categorical perception in facial emotion classification. *Proceedings of the Cognitive Science Conference*, **18**, 249-253.
- PANTIC, M. (2002). Facial gesture recognition from static dual-view face images. *Proceedings of the International Conference on Measuring Behavior*, **4**, 195-197.
- POELMAN, C. J. (1995). *The paraperspective and projective factorization methods for recovering shape and motion* (Tech. Rep. CMU-CS-95-173). Pittsburgh: Carnegie Mellon University, Robotics Institute.
- RINN, W. E. (1984). The neuropsychology of facial expression: A review of the neurological and psychological mechanisms for producing facial expressions. *Psychological Bulletin*, **95**, 52-77.
- ROWLEY, H. A., BALUJA, S., & KANADE, T. (1998). Neural network-based face detection. *IEEE Transactions on Pattern Recognition & Machine Intelligence*, **20**, 23-38.
- SAYETTE, M. A., COHN, J. F., WERTZ, J. M., PERROTT, M. A., & PARROTT, D. J. (2001). A psychometric evaluation of the facial action coding system for assessing spontaneous expression. *Journal of Nonverbal Behavior*, **25**, 167-185.
- SCHMIDT, K. L., & COHN, J. F. (2001). Human facial expressions as adaptations: Evolutionary questions in facial expression. *Yearbook of Physical Anthropology*, **44**, 3-24.
- SCHMIDT, K. L., & COHN, J. F. (in press). Signal characteristics of spontaneous facial expression: Automatic movement in solitary and social smiles. *Biological Psychology*.
- TIAN, Y. L., KANADE, T., & COHN, J. F. (2000). Eye-state detection by local regional information. In T. Tan, Y. Shi, & W. Gao (Eds.), *Advances in multimodal interfaces: ICMI 2000* (pp. 143-150). New York: Springer-Verlag.
- TIAN, Y. L., KANADE, T., & COHN, J. F. (2001). Recognizing action units for facial expression analysis. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, **23**, 97-116.
- TIAN, Y. L., KANADE, T., & COHN, J. F. (2002). Evaluation of Gabor-wavelet-based facial action unit recognition in image sequences of increasing complexity. *Proceedings of the IEEE International Conference on Automatic Face & Gesture Recognition*, **5**, 229-234.
- VAN SWEARINGEN, J. M., COHN, J. F., & BAJAJ-LUTHRA, A. (1999). Specific impairment of smiling increases severity of depressive symptoms in patients with facial neuromuscular disorders. *Journal of Aesthetic Plastic Surgery*, **23**, 416-423.
- XIAO, J., KANADE, T., & COHN, J. F. (2002). Robust full motion recovery of head by dynamic templates and re-registration techniques. *Proceedings of the IEEE International Conference on Automatic Face & Gesture Recognition*, **5**, 163-169.
- YACOOB, Y., & DAVIS, L. (1997). Recognizing human facial expression from long image sequence using optical flow. *IEEE Transactions on Pattern Recognition & Machine Intelligence*, **18**, 636-642.
- ZLOCHOWER, A. (2001). *Automated face analysis distinguishes deliberate from spontaneous smiles on the basis of quantitative differences in intensity and asymmetry*. Unpublished doctoral dissertation, University of Pittsburgh.

(Manuscript received September 17, 2002;
revision accepted for publication March 29, 2003.)