

Comprehensive Database for Facial Expression Analysis

Takeo Kanade
The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA, USA 15213
tk@cs.cmu.edu
<http://www.cs.cmu.edu/~face>

Jeffrey F. Cohn
Department of Psychology
University of Pittsburgh
The Robotics Institute
Carnegie Mellon University
4015 O'Hara Street
Pittsburgh, PA, USA 15260
jeffcohn+@pitt.edu

Yingli Tian
The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA, USA 15213
yltian@cs.cmu.edu

Abstract

Within the past decade, significant effort has occurred in developing methods of facial expression analysis. Because most investigators have used relatively limited data sets, the generalizability of these various methods remains unknown. We describe the problem space for facial expression analysis, which includes level of description, transitions among expression, eliciting conditions, reliability and validity of training and test data, individual differences in subjects, head orientation and scene complexity, image characteristics, and relation to non-verbal behavior. We then present the CMU-Pittsburgh AU-Coded Face Expression Image Database, which currently includes 2105 digitized image sequences from 182 adult subjects of varying ethnicity, performing multiple tokens of most primary FACS action units. This database is the most comprehensive test-bed to date for comparative studies of facial expression analysis.

1. Introduction

Within the past decade, significant effort has occurred in developing methods of facial feature tracking and analysis. Analysis includes both measurement of facial motion and recognition of expression. Because most investigators have used relatively limited data sets, the generalizability of different approaches to facial expression analysis remains unknown. With few exceptions [10, 11], only relatively global facial expressions (e.g., joy or anger) have been considered, subjects have been few in number and homogeneous with respect to age and ethnic background, and recording conditions have been optimized. Approaches to facial expression analysis that have been developed in this way

may transfer poorly to applications in which expressions, subjects, contexts, or image properties are more variable. In addition, no common data exist with which multiple laboratories may conduct comparative tests of their methods. In the absence of comparative tests on common data, the relative strengths and weaknesses of different approaches is difficult to determine. In the areas of face and speech recognition, comparative tests have proven valuable [e.g., 17], and similar benefits would likely accrue in the study of facial expression analysis. A large, representative test-bed is needed with which to evaluate different approaches.

We first describe the problem space for facial expression analysis. This space includes multiple dimensions: level of description, temporal organization, eliciting conditions, reliability of manually coded expression, individual differences in subjects, head orientation and scene complexity, image acquisition, and relation to non-facial behavior. We note that most work to date has been confined to a relatively restricted region of this space. We then describe the characteristics of databases that map onto this problem space, and evaluate Phase 1 of the CMU-Pittsburgh AU-Coded Facial Expression Database against these criteria. This database provides a large, representative test-bed for comparative studies of different approaches to facial expression analysis.

2 Problem space for face expression analysis

2.1 Level of description

Most of the current work in facial expression analysis attempts to recognize a small set of prototypic expressions. These prototypes occur relatively

infrequently, however, and provide an incomplete description of facial expression [11]. To capture the subtlety of human facial expression, fine-grained description of facial expression is needed. The Facial Action Coding System [FACS: 4] is a human-observer-based system designed to detect subtle changes in facial features. Viewing videotaped facial behavior in slow motion, trained observers can manually FACS code all possible facial displays, which are referred to as action units (AU) and may occur individually or in combinations.

FACS consists of 44 action units. Thirty are anatomically related to contraction of a specific set of facial muscles (Table 1) [22]. The anatomic basis of the remaining 14 is unspecified (Table 2). These 14 are

referred to in FACS as miscellaneous actions. Many action units may be coded as symmetrical or asymmetrical. For action units that vary in intensity, a 5-point ordinal scale is used to measure the degree of muscle contraction.

Although Ekman and Friesen proposed that specific combinations of FACS action units represent prototypic expressions of emotion, emotion-specified expressions are not part of FACS; they are coded in separate systems, such as EMFACS [8]. FACS itself is purely descriptive and includes no inferential labels. By converting FACS codes to EMFACS or similar systems, face images may be coded for emotion-specified expressions (e.g., joy or anger) as well as for more molar categories of positive or negative emotion [13].

Table 1. FACS Action Units.

| AU | Facial muscle | Description of muscle movement |
|-----------|----------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| 1 | Frontalis, pars medialis | Inner corner of eyebrow raised |
| 2 | Frontalis, pars lateralis | Outer corner of eyebrow raised |
| 4 | Corrugator supercillii, Depressor supercillii | Eyebrows drawn medially and down |
| 5 | Levator palpebrae superioris | Eyes widened |
| 6 | Orbicularis oculi, pars orbitalis | Cheeks raised; eyes narrowed |
| 7 | Orbicularis oculi, pars palpebralis | Lower eyelid raised and drawn medially |
| 9 | Levator labii superioris alaeque nasi | Upper lip raised and inverted; superior part of the nasolabial furrow deepened; nostril dilated by the medial slip of the muscle |
| 10 | Levator labii superioris | Upper lip raised; nasolabial furrow deepened producing square-like furrows around nostrils |
| 11 | Levator anguli oris (a.k.a. Caninus) | Lower to medial part of the nasolabial furrow deepened |
| 12 | Zygomaticus major | Lip corners pulled up and laterally |
| 13 | Zygomaticus minor | Angle of the mouth elevated; only muscle in the deep layer of muscles that opens the lips |
| 14 | Buccinator | Lip corners tightened. Cheeks compressed against teeth |
| 15 | Depressor anguli oris (a.k.a. Triangularis) | Corner of the mouth pulled downward and inward |
| 16 | Depressor labii inferioris | Lower lip pulled down and laterally |
| 17 | Mentalis | Skin of chin elevated |
| 18 | Incisivii labii superioris and Incisivii labii inferioris | Lips pursed |
| 20 | Risorius w/ platysma | Lip corners pulled laterally |
| 22 | Orbicularis oris | Lips everted (funneled) |
| 23 | Orbicularis oris | Lips tightened |
| 24 | Orbicularis oris | Lips pressed together |
| 25 | Depressor labii inferioris, or relaxation of mentalis, or orbicularis oris | Lips parted |
| 26 | Masseter; relaxed temporal and internal pterygoid | Jaw dropped |
| 27 | Pterygoids and digastric | Mouth stretched open |

| | | |
|----|---------------------------------------------------------------------------------|--------------------|
| 28 | Orbicularis oris | Lips sucked |
| 41 | Relaxation of levator palpebrae superioris | Upper eyelid droop |
| 42 | Orbicularis oculi | Eyelid slit |
| 43 | Relaxation of levator palpebrae superioris; orbicularis oculi, pars palpebralis | Eyes closed |
| 44 | Orbicularis oculi, pars palpebralis | Eyes squinted |
| 45 | Relaxation of levator palpebrae superioris; orbicularis oculi, pars palpebralis | Blink |
| 46 | Relaxation of levator palpebrae superioris; orbicularis oculi, pars palpebralis | Wink |

Table 2. Miscellaneous Actions.

| AU | Description of Movement |
|----|-------------------------|
| 8 | Lips toward |
| 19 | Tongue show |
| 21 | Neck tighten |
| 29 | Jaw thrust |
| 30 | Jaw sideways |
| 31 | Jaw clench |
| 32 | Bite lip |
| 33 | Blow |
| 34 | Puff |
| 35 | Cheek suck |
| 36 | Tongue bulge |
| 37 | Lip wipe |
| 38 | Nostril dilate |
| 39 | Nostril compress |

2.2 Transitions among expressions

A simplifying assumption in previous research is that expressions are singular and begin and end from a neutral position. In reality, facial expression is more complex, especially at the level of action units. Action units may occur in combinations or show serial dependence. Transitions from action units or combination of actions to another may involve no intervening neutral state. Parsing the stream of behavior is an essential requirement of a robust facial analysis system, and training data are needed that include dynamic combinations of action units, which may be either additive or non-additive.

An example of an additive combination is smiling (AU 12) with mouth opening, which would be coded as AU 12+25, AU 12+26, or AU 12+27 depending on the degree of lip parting and whether and how far the mandible was lowered. In the case of AU 12+27, for instance, the facial analysis system would need to detect transitions among all three levels of mouth opening while continuing to recognize AU 12, which may be simultaneously changing in intensity.

Non-additive combinations represent further complexity. Following usage in speech science, we refer to these interactions as co-articulation effects. An example is the combination AU 12+15, which often occurs in embarrassment. While AU 12 raises the cheeks, its action on the lip corners is modified by the downward action of AU 15. The resulting appearance change is highly dependent on timing. The downward action of the lip corners may occur either simultaneously or sequentially. To be comprehensive, a database should include individual action units and both additive and non-additive combinations, especially those that involve co-articulation effects. A classifier trained only on single action units may perform poorly for combinations in which co-articulation effects occur.

2.3 Deliberate versus spontaneous expression

Most face expression data have been collected by asking subjects to perform a series of expressions. These directed facial action tasks may differ in appearance and timing from spontaneously occurring behavior [5]. Deliberate and spontaneous facial behavior are mediated by separate motor pathways, the pyramidal and extra-pyramidal motor tracks, respectively [16]. As a consequence, fine-motor control of deliberate facial actions is often inferior and less symmetric to 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 to lie detection [5]. Differences in the temporal organization of spontaneous and deliberate facial actions are particularly important in that many pattern recognition approaches, such as Hidden Markov Modeling, are highly dependent on the timing of appearance change. Unless a database includes both

deliberate and spontaneous facial actions, it will likely prove inadequate for developing face expression methods that are robust to these differences.

2.4 Reliability of expression data

When training a system to recognize facial expression, the investigator assumes that training and test data are accurately labeled. This assumption may or may not be accurate. Asking subjects to perform a given action is no guarantee that they will. To ensure internal validity, expression data must be manually coded, and the reliability of the coding verified. Inter-observer reliability can be improved by providing rigorous training to observers and monitoring their performance. FACS coders must pass a standardized test, which ensures (initially) uniform coding among international laboratories. Monitoring is best achieved by having observers independently code a portion of the same data. As a general rule, 15% to 20% of data should be comparison coded. To guard against drift in coding criteria [12], re-standardization is important. In assessing reliability, coefficient kappa [7] is preferable to raw percentage of agreement, which may be inflated by the marginal frequencies of codes. Kappa quantifies inter-observer agreement after correcting for level of agreement expected by chance.

2.5 Individual differences among subjects

Face shape, texture, color, and facial and scalp hair vary with sex, ethnic background, and age [6, 23]. 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 facial feature analysis more generally. Beards, eyeglasses, or jewelry may obscure facial features. Such individual differences in appearance may have important consequence for face analysis. Few attempts to study their influence exist. An exception was a study by Zlochower [23]. They found that algorithms for optical flow and high-gradient component detection that had been optimized for young adults performed 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, morphology, frequency of intense expression, and overall rate of expression. Individual differences in these characteristics are well established and are an important aspect of individual identity [14]. (Incidentally, these individual differences could be used to augment the accuracy of face recognition algorithms). An extreme example of variability in expressiveness occurs in individuals who have incurred damage either to the facial nerve or central nervous system [16, 19, 21]. To develop algorithms that are robust to individual differences in facial features and behavior, it is essential to include a large sample of varying ethnic background, age, and sex, that includes people who have facial hair and wear jewelry or eyeglasses, and includes both normal and clinically impaired individuals.

2.6 Head orientation and scene complexity

Face orientation relative to the camera, presence and actions of other people, and background conditions may influence face analysis. In the face recognition literature, face orientation has received deliberate attention. The FERET data base [17], for instance, includes both frontal and oblique views, and several specialized data bases have been collected to try to develop methods of face recognition that are invariant to moderate change in face orientation [20]. In the face expression literature, use of multiple perspectives is rare and relatively less attention has been focused on the problem of pose invariance. Most researchers assume that face orientation is limited to in-plane variation [1] or that out-of-plane variation is small [10, 11]. In reality, large out-of-plane variation in head position is common and often accompanies change in expression. Kraut [9] found that smiling typically occurs while turning toward another person. Camras [2] showed that infant surprise expressions often occur as the infant pitches her head back. To develop pose invariant methods of face expression analysis, image data are needed in which facial expression changes in combination with significant non-planar change in pose.

Scene complexity, such as background and presence of other people, potentially influences accuracy of face detection, feature tracking, and expression recognition. Most databases use image data in which the background is neutral or has a consistent pattern and only a single person is present in the scene. In natural environments, multiple people interacting with each other are likely to be present, and their effects need to be understood. Unless this variation is represented in training data, it will be difficult to develop and test algorithms that are robust to such variation.

2.7 Image acquisition and resolution

Image acquisition includes properties and number of video cameras and digitizer, size of the face image relative to total image dimensions, and ambient lighting. All of these factors may influence facial expression analysis. Images acquired in low light or at coarse resolution can provide less information about facial features. Similarly, when face image size is small relative to total image size, less information is available. NTSC cameras record images at 30 frames per second, The implications of down-sampling from this rate are unknown. Many algorithms for optical flow assume that pixel displacement between adjacent frames is small. Unless they are tested at a range of sampling rates, robustness to sampling rate and resolution cannot be assessed.

Within an image sequence, change in head position relative to the light source and variation in ambient lighting have potentially significant effects on face expression analysis. A light source above the subject's head will cause shadows to fall below the brows, which can obscure the eyes, especially for subject's with more pronounced bone structure or hair. Methods that work well in studio lighting may perform poorly in more naturalistic lighting (e.g., through an exterior window) when angle of lighting changes across an image sequence.

Most investigators use single-camera set-ups, which is problematic when a frontal orientation is not required. With image data from a single camera, out-of-plane variation may be difficult to standardize. For more than small out-of-plane rotation, multiple cameras may be required. Multiple camera setups can support 3-D modeling and in some cases ground truth with which to assess the accuracy of image alignment.

Image resolution is another concern. Professional grade PAL cameras, for instance, provide very high resolution images. By contrast, security cameras provide ones that are seriously degraded. Although post processing may improve image resolution, the degree of potential improvement is likely limited. Also the effects of post processing for expression recognition are not known. Algorithms that work well at optimal resolutions of full-face frontal images and studio lighting can be expected to perform poorly when recording conditions are degraded or images are compressed. Without knowing the boundary conditions of face expression algorithms, comparative performance is difficult to assess. Algorithms that appear superior within one set of boundary conditions may perform more poorly across

the range of potential applications. Appropriate data with which these factors can be tested are needed.

2.8 Relation to non-facial behavior

Facial expression is one of several channels of nonverbal communication that may occur together. Contraction of the zygomaticus major (AU 12), for instance, often is associated with positive or happy vocalizations, and smiling tends to increase vocal fundamental frequency [3]. Few research groups, however, have attempted to integrate gesture recognition broadly defined across multiple channels of communication. An important question is whether there are advantages to early rather than late integration. Databases containing multi-modal expressive behavior afford opportunity for integrated approaches to analysis of facial expression, prosody, gesture, and kinetic expression.

2.9 Summary and problem statement

The problem space for facial expression includes multiple dimensions. To develop robust methods of facial expression analysis, these dimensions must be adequately sampled. In addition, to allow for comparative tests of alternative approaches to facial expression analysis, appropriate data must be made available to the face analysis community. To meet these needs, we have developed the CMU-Pittsburgh AU-Coded Facial Expression Database to serve as a test-bed for algorithm development and testing.

3 The CMU-PITTSBURGH AU-Coded Face Expression Image Database

Our interdisciplinary research group of psychologists and computer scientists is developing a large, representative facial expression database for use in both training and testing of algorithms for facial expression analysis. In this section we first describe the CMU-Pittsburgh AU-Coded Face Expression Database. We then evaluate the database against the criteria presented above, and discuss current and future work.

3.1 Description of database

Facial behavior was recorded in 210 adults between the ages of 18 and 50 years. They were 69% female, 31% male, 81% Euro-American, 13% Afro-American, and 6% other groups (Table 3). They were observed in an observation room equipped with a chair on which to sit and two Panasonic WV3230 cameras, each connected to a Panasonic

AG-7500 video recorder with a Horita synchronized time-code generator. One of the cameras was located directly in front of the subject, and the other was positioned 30 degrees to the subject's right. An example of image data from the CMU-Pittsburgh AU-Coded Facial Expression Database can be seen in Figure 1. For approximately one third of subjects, ambient room lighting augmented by a high-intensity lamp was used. For the other two thirds, two high-intensity lamps with reflective umbrellas were used to provide uniform lighting.

Table 3. CMU-Pittsburgh AU-Coded Facial Expression Database.

| <u>Subjects</u> | |
|----------------------------|---------------------------------------------------|
| Number of subjects | 210 |
| Age | 18-50 years |
| Women | 69% |
| Men | 31% |
| Euro-American | 81% |
| Afro-American | 13% |
| Other | 6% |
| <u>Digitized sequences</u> | |
| Number of subjects | 182 |
| Resolution | 640x490 for grayscale 640x480 for 24-bit color |
| Frontal view | 2105 |
| 30-degree view | Videotape only |
| Sequence duration | 9-60 frames/sequence |
| <u>Action Units</u> | |
| From Table 1 | All but AU 13 |
| From Table 2 | AU 8, 38, and 39 |

Subjects were instructed by an experimenter to perform a series of 23 facial displays; these included single action units and combinations of action units. Although each display began and ended in a neutral face, for combinations the timing of action units could be examined. Sixty subjects performed head rotation to 30 degrees with facial expression, which was recorded with both cameras.

Image sequences for frontal views were digitized into either 640x490 or 640x480 pixel arrays with 8-bit gray-scale or 24-bit color values. To date, 1917 image sequences from 182 subjects have been FACS coded for either target action units or the entire sequence. Thirty-degree views are available on videotape. Approximately fifteen percent of the 1917 sequences were comparison coded by a second certified FACS coder. Inter-observer agreement was quantified with coefficient kappa, which is the proportion of agreement above what would be expected to occur by chance [7]. The mean kappas for inter-observer agreement were 0.82 for action units coded at apex and 0.75 for frame-by-frame coding. For those action units that were coded only at

apex and not from beginning to end, we are performing additional coding. We also are increasing the number of action units that have been coded for intensity.



Figure 1. Frontal and 30-degree to the side views available in CMU-Pittsburgh AU-Coded Facial Expression Database.

The CMU-Pittsburgh AU-Coded Face Expression Image Database includes all of the action units listed in Table 1 either alone or in combination with other AUs, with the exception of AU 13, which most people find difficult to perform voluntarily and occurs infrequently in spontaneous behavior. Miscellaneous actions listed in Table 2 are less well represented. The database has been used in several studies [e.g., 10, 11, 18]. For information about use of the database, please contact the second author.

3.2 Database evaluation

With 1917 sequences from 182 male and female subjects of varying ethnic background, the database provides a sufficiently large, representative test-bed for assessing facial behavior and individual-difference factors. The level of facial expression description supports analyses of single action units and both additive and non-additive combinations. By applying EMFACS, emotion-specified expressions (e.g., joy, anger) may be analyzed as well. Action unit combinations representing joy, surprise, sadness, disgust, anger, and fear are included in adequate numbers. These prototypic expressions may be further combined to form aggregates of positive and negative expression. The database includes subjects of varying skin color and facial conformation, so that the influence of these factors may be examined.

Lighting conditions and context were relatively uniform. Out-of-plane head motion was small to mild. Images were acquired using S-Video cameras (approximately 425 lines per frame) and digitized at frame rate omitting odd fields. The image data may be transformed or down-sampled to examine the effects of reduced resolution, sampling rate, image compression, and luminance.

The database has several limitations in its present form. Intensity scoring of action units is incomplete, and for many sequences only target frames rather than the entire sequence have been coded. As noted above, continued coding is underway in order to further improve the data. Another limitation is the lack of spontaneous expressions. Because deliberate and spontaneous expressions may have different appearance and timing, it is important to have adequate numbers of each.

One solution, which we are pursuing, is to examine our videotapes for instances of spontaneously occurring action units that occurred during the experimental session. We now have a large sample of spontaneous smiles (AU 12) and related action units (e.g., AU 6), and these will be added to the database.



Figure 2. Automated face detection and facial-feature and eye tracking in image sequence obtained from synchronized cameras and split-screen generator.

3.3 Database extension.

Emotion analysis. We are making arrangements to include data from our studies of emotion processes. In these studies, we have elicited emotion responses in infants, children, and adults and recorded their facial and vocal behavior [e.g., 3, 23]. Figure 2 from [18] shows an example of automated face detection and feature tracking in an image sequence of face-to-face interaction between a mother and her infant. The image was acquired from two synchronized cameras and a split-screen special-effects generator, which is commonly used by psychologists who study social interaction. The image suggests challenges in several dimensions of

the face analysis problem space. The face images and facial features, especially in the infant, are small relative to the image size, the infant's face has low texture, some shadows occur, and the likelihood of sudden and large motion, occasional occlusion, and moderate out-of-plane motion is high. These are challenging problems for which appropriate training and testing data are critical.

Surgical application. Another data source is facial behavior from patients who have experienced damage to the facial nerve or the higher brain centers that control facial behavior [16]. An example can be seen in Figure 3 from [21]. Notice the mild asymmetry in repose due to muscle weakness and the more marked asymmetry that occurs in the second frame. The inclusion of clinical data such as these challenges assumptions of symmetry, which are common when working with directed facial action task images from subjects who have normal facial function.




Figure 3. Automated facial feature tracking in image sequence from subject with facial nerve impairment.

Omni-view facial analysis. We plan to expand our ability to analyze facial expression from multiple perspectives. Our research team has a method referred to as virtualized reality that can integrate input from dense arrays of over 50 cameras [15]. An example of a face image from multiple camera views can be seen in Figure 4. The subject is shown from the perspective of six cameras. Using virtualized reality, intermediate views may be generated for all possible perspectives. This approach affords accurate 3-D modeling of facial expression and the necessary ground-truth with which to test image alignment algorithms that are based on single-camera data.

4. Conclusion.

The problem space for facial expression analysis includes multiple dimensions. These include level of description, transitions among expressions, distinctions between deliberate and spontaneous expressions, reliability and validity of training and test data, individual differences among subjects in facial features and related characteristics, head orientation and scene complexity, image

characteristics, and relation to other non-verbal behavior. Development of robust methods of facial expression analysis requires access to databases that adequately sample from this problem space. The CMU-Pittsburgh AU-Coded Facial Expression Image Database provides a valuable test-bed with which multiple approaches to facial expression analysis may be tested. In current and new work, we will further increase the generalizability of this database.



Figure 4. Face images obtained from omniview camera. (6 images are shown here out of 50 that were available).

5. Acknowledgements

This research was supported by grant number R01 MH51435 from the National Institute of Mental Health.

6. References

- [1] M.S. Bartlett, J.C. Hager, P. Ekman, T.J. Sejnowski, Measuring facial expressions by computer image analysis. *Psychophysiology*, 36:253-264, 1999.
- [2] L.A. Camras, L. Lambrecht, L., G.F. Michel. Infant "surprise" expressions as coordinative motor structures. *Journal of Nonverbal Behavior*, 20:183-195, 1966.
- [3] J.F. Cohn, G.S. Katz, Bimodal expression of emotion by face and voice. *ACM and ATR Workshop on Face/Gesture Recognition and Their Applications*, 41-44, 1998.
- [4] P. Ekman, W.V. Friesen, *Facial Action Coding System*, Consulting Psychologist Press, Palo Alto, CA, 1978.
- [5] P. Ekman, E. Rosenberg (Eds.), *What the face reveals*. NY: Oxford University, 1997.
- [6] L.G. Farkas, I.R. Munro, *Anthropometric Facial Proportions in Medicine*, Springfield, Illinois: Charles C. Thomas, 1987.
- [7] J.L. Fleiss, *Statistical Methods for Rates and Proportions*, NY: Wiley, 1981.
- [8] W.V. Friesen, P. Ekman, EMFACS-7: *Emotional Facial Action Coding System*, Unpublished manuscript, University of California at San Francisco, 1983.
- [9] R.E. Kraut, R. Johnson, Social and emotional messages of smiling: An ethological approach. *Journal of Personality and Social Psychology*, 37:1539-1523, 1979.
- [10] J.J.J. Lien, T. Kanade, J.F. Cohn, C.C. Li, Detection, tracking, and classification of subtle changes in facial expression. *Journal of Robotics and Autonomous Systems*, in press.
- [11] J.J.J. Lien, T. Kanade, J.F. Cohn, C.C. Li., Automated facial expression recognition. *Proceedings of the Third IEEE International Conference on Automatic Face and Gesture Recognition (FG'98)*, 390-395, 1998.
- [12] P. Martin, P. Bateson, *Measuring behavior: An introductory guide*. Cambridge: Cambridge University, 1986.
- [13] R. Matias, J.F. Cohn, S. Ross, A comparison of two systems to code infants' affective expression. *Developmental Psychology*, 25: 483-489, 1989.
- [14] A.S.R. Manstead, Expressiveness as an individual difference. In R.S. Feldman & R. Rime (Eds.), *Fundamentals of nonverbal behavior*, 285-328. NY: Cambridge University, 1991.
- [15] P.J. Narayanan, P. Rander, T. Kanade, Constructing virtual worlds using dense flow. *Proceedings of the International Conference on Computer Vision '98*, 3-10, 1998.
- [16] W.E. Rinn, The neuropsychology of facial expression: A review of the neurological and psychological mechanisms for producing facial expressions. *Psychological Bulletin*, 95:52-77, 1984.
- [17] S.R. Rizvi, P.J. Phillips, H. Moon, The Feret verification testing protocol for face recognition algorithms. *Proceedings of the Third International Conference on Automatic Face and Gesture Recognition*, 48-55, 1998.
- [18] Y.L. Tian, T. Kanade, J.F. Cohn, Robust lip tracking by combining, shape color and motion. *Asian Conference for Computer Vision, 2000*.
- [19] J.M., VanSwearingen, J.F. Cohn, A. Bajaj-Luthra, Specific impairment of smiling increases severity of depressive symptoms in patients with facial neuromuscular disorders. *Journal of Aesthetic Plastic Surgery*, in press.
- [20] T. Vetter, Learning novel views to a single face image. *Proceedings of the IEEE International Conference on Automatic Face and Gesture Recognition '95*: 22-29, 1995.
- [21] G.S. Wachtman, J.F. Cohn, J. VanSwearingen, E.K. Manders, Pixel-wise tracking of facial movement by computer image processing. *Robert H. Ivy Society of Plastic & Reconstructive Surgeons*, Pittsburgh, 1999.
- [22] A.J. Zlochower, *Deciphering emotion from the face: An evaluation of FACS, EMG, and computer-vision based approaches to facial expression analysis*. Unpublished manuscript, University of Pittsburgh, 1997.
- [23] A.J. Zlochower, J.F. Cohn, J.J. Lien, T. Kanade, A computer vision based method of facial expression analysis in parent-infant interaction, *International Conference on Infant Studies*, Atlanta, Georgia, April 1998.