



Imaging Facial Physiology for the Detection of Deceit

P. TSIAMYRTZIS

Department of Statistics, Athens University of Economics and Business, 76 Patission Street, 104 34 Athens, Greece
pt@aueb.gr

J. DOWDALL, D. SHASTRI AND I.T. PAVLIDIS

Department of Computer Science, University of Houston, 501 Philip G. Hoffman Hall, Houston, TX 77204-3010
jbdowdal@central.uh.edu
djshastr@bayou.uh.edu
ipavli@central.uh.edu

M.G. FRANK

*School of Informatics/Department of Communication, University of Buffalo, State University of New York,
Buffalo, NY 14260*
mfrank83@buffalo.edu

P. EKMAN

Department of Psychiatry, UC San Francisco, San Francisco, CA-94143-0984
paul@paulekman.net

Received March 25, 2005; Revised August 21, 2005; Accepted November 8, 2005

First online version published in June, 2006

Abstract. Previous work has demonstrated the correlation of increased blood perfusion in the orbital muscles and stress levels for human beings. It has also been suggested that this periorbital perfusion can be quantified through the processing of thermal video. The idea has been based on the fact that skin temperature is heavily modulated by superficial blood flow. Proof of this concept was established for two different types of stress inducing experiments: startle experiments and mock-crime polygraph interrogations. However, the polygraph interrogation scenarios were simplistic and highly constrained. In the present paper, we report results derived from a large and realistic mock-crime interrogation experiment. The interrogation is free flowing and no restrictions have been placed on the subjects. Additionally, we propose a new methodology to compute the mean periorbital temperature signal. The present approach addresses the deficiencies of the earlier methodology and is capable of coping with the challenges posed by the realistic setting. Specifically, it features a tandem CONDENSATION tracker to register the periorbital area in the context of a moving face. It operates on the raw temperature signal and tries to improve the information content by suppressing the noise level instead of amplifying the signal as a whole. Finally, a pattern recognition method classifies stressful (Deceptive) from non-stressful (Non-Deceptive) subjects based on a comparative measure between the entire interrogation signal (baseline) and a critical subsection of it (transient response). The successful classification rate is 87.2% for 39 subjects. This is on par with the success rate achieved by highly trained psycho-physiological experts and opens the way for automating lie detection in realistic settings.

Keywords: deception detection, polygraph, thermography, facial tracking, contact-free physiological monitoring

1. Introduction

Stress has long been associated with deceptive behavior during interrogation. It emanates from a brain center and manifests itself in the peripheral senses through a variety of physiological signatures, including perspiration, pulse, and breathing rate (Sokolov and Cacioppo, 1997; Drummond and Lance, 1987). Polygraph examinations are based exactly on this principle. However, polygraph technology is based on contact sensors and heuristic analysis of resultant signals by experts (Kleinmuntz and Szucko, 1984). A consequence of this is that the comfort of subjects is compromised, which is very important in psycho-physiology, as it is believed to contaminate the experiment (Yankee, 1965). The analysis of the results is slow because it is performed manually by several experts who later have to reconcile their findings. The processing of a typical 10 min interrogation session may take several hours. Finally, due to the heuristic nature of the analysis the outcome cannot be embraced with high confidence and always has a suggestive nature (Committee to Review the Scientific Evidence on the Polygraph, 2003).

Traditionally, polygraph examination is used by the U.S. Government to check periodically the veracity of employees that hold security clearances. It is also favored in some instances by private companies for screening during the hiring process. This is primarily for potential employees who will assume sensitive positions. Due to the problems of the existing polygraph technology, polygraph examinations are inadmissible in court. In recent years and with the intensification of anti-terrorism efforts, there has been a clear desire to improve or replace the polygraph technology with something better. There is a real need for administering real-time, highly automated veracity tests inside the country and around the globe, where human intelligence is being collected. At the same time, similar technology may be useful for quick screening of potential suspects in ports of entry (Knight, 2004).

Pavlidis et al. have discovered a physiological signature on the face directly associated with stress levels (Pavlidis et al., 2000). This physiological response could be considered to be part of the 'fight or flight' syndrome triggered by the autonomic nervous system, whereby blood redistributes peripherally towards musculoskeletal tissue (Martini et al., 2001). Specifically, Pavlidis et al. demonstrated that during stress produced

by startle stimuli in the lab, subjects exhibited elevated blood perfusion in the orbital muscle area, which resulted in localized elevated temperature. They further suggested that such a heat signature can be captured by a highly sensitive thermal imaging system and analyzed using pattern recognition methods. Based on this principle, they have later developed an imaging system for quantifying stress during polygraph examinations. The system was put into comparative test with traditional polygraphy in an experiment designed and implemented by the Department of Defense Polygraph Institute (DOD PI). The accuracy of the two modalities was found to be equivalent (around 80%), thereby establishing the feasibility of the idea (Pavlidis et al., 2002).

The DOD PI experiment was small in scale (20 subjects) and highly contrived. Questions were posed to subjects multiple times and there were long pauses between each question. The subjects were required to give a binary answer ('yes/no') to impose uniformity on the time scale (Pavlidis and Levine, 2002). The subjects were also instructed to stay as still as possible. This is quite apart from the natural interrogation scenario where there are no motion restrictions and the flow is continuous and non-uniform. Actually, in practice the pace of an interrogation typically accelerates or decelerates at times to build up psychological pressure.

Recently, the Communication Lab at Rutgers University designed and executed a large mock-crime experiment where all the aforementioned restrictions during the interrogations were removed. The Computational Physiology Lab in the University of Houston joined this effort early in the data collection and collected facial thermal video along with audio for all those subjects run over a consecutive 9 month period (the entire collection period was 19 months). The analysis of the thermal imagery demonstrated the limitations of the earlier methodology proposed in the DOD PI experiment. Then, the problem of signal fading was addressed through the transformation from the temperature to the blood perfusion domain via bioheat modeling. This magnified the scale of the phenomenon but also amplified the noise levels.

In this paper, we propose an entirely different methodology, which addresses the issue of automatic deception detection under the most realistic conditions and for a statistically significant number of subjects. The only similarity between the present and past approaches is that the periorbital region continues to be




Figure 1. A comparison of extracted temperature signals over time from the periorbital region of a sample subject. The red (dark) signal was extracted using a stationary rectangle. The green (light) signal was extracted using our proposed tracking method, tandem assisted tracking. Both regions were selected from the subject's periorbital region from the first frame of the measurement. Notice the high variability of the untracked region's extracted temperature signal, which would render it unacceptable for physiological interpretation.

the region of interest. Additionally, we employ a tracking method to facilitate the extraction of the temperature measurement from the periorbital region, without which any measurements taken would be marred by noise to the point of being useless for physiological interpretation (see Fig. 1).

In Section 2 we provide an overview of the experiment and the proposed imaging solution. In Section 3 we describe the facial and periorbital tracking algorithm, which enables acquisition of temperatures from the tissue of interest despite head motion. In Section 4 we describe the physiological measurement and our noise suppression methodology. In Section 5 we present the pattern recognition algorithm that operates on the filtered temperature signal to perform the binary classification ('Deceptive' or 'Non-Deceptive'). In Section 6 we discuss the experimental design and results and ponder the strong and weak points of our approach. We conclude the paper in Section 7.

2. Overview of Our Approach

Our objective is to extract the mean temperature signal in the periorbital area of the subject through the

course of the interrogation. Fig. 2 depicts the system architecture that realizes this objective. Specifically, we have a specialized tissue tracker that registers the subject's periorbital region despite the presence of motion.

Even with the tracking, the extracted temperature signal is still marred by noise. While some systemic noise from the electronics of the thermal imaging system is present, the major source of noise stems from imperfections in the tracker. Therefore, a noise suppression algorithm (filter) is necessary to improve the quality of the signal. Finally, we have a pattern recognition algorithm that compares the overall outlook of the filtered signal (baseline) versus specific instances of high psycho-physiological value (transient response). This comparison constitutes the basis of a binary classification rule ('Deceptive' or 'Non-Deceptive').

3. Tandem Facial Tracking

Sequential Monte Carlo methods have been employed for tracking in computer vision applications for nearly




Figure 2. System architecture.

a decade (Doucet et al., 2001). The CONDENSATION algorithm (Isard and Blake, 1998a) formalized their introduction to the field, providing an effective and easily implemented framework for tracking. This work was followed by innovations that incorporated information fusion (Isard and Blake, 1998b), reduced computational cost (MacCormick and Isard, 2000), and facilitated multi-object tracking (Isard and Maccormick, 2001; Maccormick and Blake, 2000). The majority of research was concentrated on the problem of tracking in the visible spectrum, but there has also been investigation into utilizing the thermal portion of the infrared spectrum (Eveland et al., 2003).

Typically, in the aforementioned methods, tracking accuracy can fluctuate, and as long as it is regained




Figure 3. Out of plane rotation of the head.

momentarily, the effect on end applications (e.g., surveillance) is minimal. However, in physiological monitoring applications such as ours, even temporary loss of tracking creates serious problems. It results in spikes in the temperature signal of the tracked tissue, which are falsely indicative of strong physiological responses. This may mislead the pattern recognition algorithm. Tracking is challenging for the periorbital region, since it is a small tissue area and its projection is greatly affected by any rotation out of the projected image plane (see Fig. 3).

Even when tracking is maintained, its accuracy typically degrades over time. To address this issue, we have introduced the concept of *tandem tracking*. Simultaneously with the periorbital region we track a central facial region that is rich in contrasting features and more invariant to out of plane rotations than the periorbital region. The quality of tracking for this central region is better on average and can be used to correct the estimate of the periorbital tracker. This leads to the extraction of a superior quality temperature signal (see Fig. 4).

The scheme works as follows:

1. The operator selects two rectangular regions on the initial frame of the thermal clip—the central facial and periorbital.
2. The two regions are tracked by two independent CONDENSATION trackers. The feedback measurement for each tracker is based on a template




Figure 4. Example of tandem assisted tracking resulting in a superior temperature extraction. On the graph the green (light) line is the temperature extracted for each frame using the tandem assisted tracking, and the red (dark) line is the temperature extracted for each frame using the normal tracking. Note the highlighted portion of the graph where there is a spike at frame 8934 in the temperature extracted from the normal tracker. This is a result of a tracking failure from the normal tracker because of an abrupt movement by the subject. The tandem assisted tracking was able to correctly track the subject through this motion and therefore extract the correct temperature at this frame. The grey scale thermograph on the right shows both the tandem assisted (green/light), and the normal (red/dark) periorbital trackers.

extracted from the initial rectangular regions. In essence, this is a sub-sampling of the rectangular region of interest. We have determined experimentally that a surprisingly small sub-sample of the original area produces the same tracking result as the full area.

3. The relative spatial position of the central and periorbital regions is established in the initial frame.
4. In subsequent frames the CONDENSATION trackers come up with independent position estimates. The position estimate of the central facial tracker is computed first and is used to influence the periorbital region tracker. Since the relative position of the periorbital region to the central facial region has been established in the initial frame, the position probability suggested by the central tracker for the periorbital region is fused into the CONDENSATION framework for the independent estimate of the periorbital tracker itself. The combined estimate provides for a robust solution and helps the periorbital tracker to overcome local extrema (see Fig. 5).

3.1. Information Fusion

The concept of using the results from one tracker to boost the performance of another tracker has been previously introduced in the literature. One formula-

tion of this within the context of the CONDENSATION algorithm is the concept of information fusion through *Importance Sampling* introduced in Isard and Blake (1998b) where the performance of a high level tracker, that is a tracker which parameterizes the target object's state in a precise manner, is aided by utilizing information from a low level tracker, that is a tracker which parameterizes the target object in a general manner. We have modified this approach by utilizing the information from a robust high-level tracker (the central region tracker) to boost the stability of another high-level tracker (the periorbital region tracker). The benefit of the tandem relationship is enhanced by tuning the dynamics model for the central region tracker to be more tolerant to large changes in motion (by increasing the stochastic component of its dynamics model), and tuning the dynamics model for the periorbital region tracker to allow for more localization (by decreasing the stochastic component of its dynamics model).

Thus far we have omitted the details of how the central region tracker influences the periorbital region tracker. Incorporating this influence into our algorithm is a two step process. First, we need to derive the information about the periorbital region tracker's posterior density from the central region tracker. Then, we need to have the periorbital region tracker utilize this information in a meaningful way. For our framework let us assume that $B_n = (b_1, b_2, \dots, b_k)^T$ is the general parameterized representation of the periorbital region

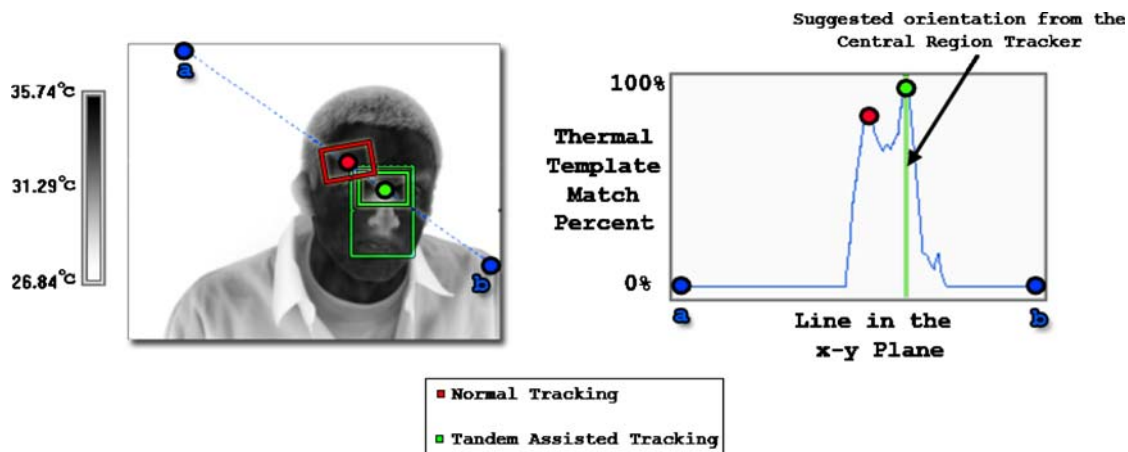


Figure 5. Tandem Facial Tracking: The subject in the thermograph (shown here in gray-scale) on the left has had his periorbital region tracked by both the normal (red/dark) and the tandem assisted (green/light) trackers. The normal periorbital region tracker has become confused because of a local maximum in the template match response space. However, the periorbital region tracker that is assisted by the central region tracker in the tandem configuration has maintained the proper track. This is because of the suggestion of the central region tracker, which is shown in the graph on the right as the vertical green (light) line. The response of the template matching along the blue line with endpoints a and b in the thermograph (see the graph on the right) has been generated to highlight this point.

tracker at frame n , for parameters b_1, b_2, \dots, b_k , where k can be as large as necessary to describe the tracker for a given representation. Similarly, let us assume that $C_n = (c_1, c_2, \dots, c_k)^T$ represents the central region tracker at frame n . In our representation we model the tracked regions by their center of mass positions and orientations. The state space \mathbb{S} for these is given in Eq. (1).

$$\mathbb{S} = \{(x, y, \theta) \mid x, y \in \mathbb{R}, \theta \in [0, 2\pi)\} \quad (1)$$

Where x and y are the coordinates of the center of mass of the tracked region, and θ is the angle of the tracked region with respect to the x -axis of the image coordinate system. The central region tracker uses its current state, C_n , coupled with information derived during initialization, C_0 and B_0 , to derive the information about the possible state of the periorbital region tracker B_n (see Eq. (2)).

$$\hat{B}_n = C_n + (B_0 - C_0) \quad (2)$$

Where \hat{B}_n represents the projected current state of the periorbital tracker based on the current state of the central region tracker. \hat{B}_n is generated from the maximum likelihood sample of the central region tracker. We use the maximum likelihood sample instead of the k -highest weighted samples in order to exert all the influence of the tandem tracker in one location. Thus, we explore the parameterization space around \hat{B}_n to the

maximum extent possible by virtue of the stochastic component of the dynamics model.

Once \hat{B}_n has been calculated, it must be incorporated into the CONDENSATION framework for the periorbital tracker. For this purpose we modify the selection step of the CONDENSATION framework to automatically pre-select a portion of its samples based on this new knowledge. The information fusion methodology is depicted in Fig. 6, where Z is an image and X is the object parameterization; $p(X_{t-1}|Z_{t-1})$ is the posterior probability from the previous time step derived from the sample set $\{(s_{t-1}^{(i)}, \pi_{t-1}^{(i)}), i = 1, \dots, N\}$, where $s^{(i)}$ is one of N samples with corresponding weight $\pi^{(i)}$; $p(X_t|Z_{t-1})$ is the prior for time-step t ; $p(Z_t|X_t)$ is the observation density at time step t ; $p(X_t|Z_t)$ is the posterior probability for the current time step derived from the sample set $\{(s_t^{(n)}, \pi_t^{(n)}), n = 1, \dots, N\}$.

We assign a percent weight α to the additional information from the central face region tracker. During the selection step of the CONDENSATION framework we choose $(1-\alpha)\%$ of the samples normally and $\alpha\%$ based on the additional information. All of the samples selected using the additional information are set equal to its suggested state, then during the prediction step only the stochastic component of the dynamics model is applied to these samples. We assign $\alpha = 20\%$, which we experimentally justified as large enough (greater than 5% in Fig. 7) to derive the benefit from the tandem tracker's influence and small enough to allow the




Figure 6. Modifications to the CONDENSATION framework. On the left is the traditional implementation of the CONDENSATION framework, on the right is the modified version of the CONDENSATION framework. Notice that during the selection step for the modified version, the additional information is used as the basis for producing new samples, which would have not been spawned if we used the traditional method.




Figure 7. The periorbital tracker's mean error over 200 frames of video for different percentage weights given to the tandem tracker's influence. The error was calculated as the Euclidean distance between the maximum likelihood sample of the periorbital tracker and the true location of the target. The experiment was performed using an artificial tracking environment described in detail Section 3.3.

periorbital tracker to function as a normal CONDENSATION tracker.

In Fig. 7 it is apparent that the periorbital tracker needs to have only a small percentage of its samples derived through the tandem tracker's influence to achieve maximum accuracy. The experiment was performed under large motion conditions (between 6–8.4 pixels/frame). Under these circumstances the tandem tracker's influence provided the best approximation of the current state for the periorbital tracker.

3.2. Thermal Templates

Sequential Monte Carlo methods provide frameworks for propagating conditional densities over time, but omit the vehicle with which to measure the posterior density. The open ended property of these methods

facilitates their adaptation to a myriad of different applications (Doucet et al., 2001), but necessitates the selection of an appropriate representation and measurement criteria for the phenomena in question. Tracking research in the thermal infrared spectrum has utilized several different types of representations, from shape contours (Eveland et al., 2003) to blobs. We have a different end use for our tracking than these other research efforts; we use the tracker to extract physiological information. This poses the extra constraint that we need to be able to track any sub region of the face or other exposed skin. If we utilized the outline of the object to be tracked as the representation of our tracking model, measurements could be compromised anytime the outline of the tracked object remained the same while the inside of the object changed, such as a person rotating his/her head out of plane (see Fig. 8).

We opted to adapt the well known technique of template matching (Barker and Matthews, 2001) to the thermal modality. The fundamental idea behind template matching is to sub-sample the object of interest, the template, from the initial frame of video and then to find the region in subsequent frames that most closely resembles the template. The underlying assumption when using templates to represent the object of interest is that the appearance of the object will remain relatively constant throughout the course of the video. In our framework, we work with a thermal video $V(\mathbf{c}, t)$, where $\mathbf{c} = (x, y)^T$ are pixel coordinate pairs and t is the time of the frame. Templates have been traditionally used in visible spectrum images, but they can be easily adapted to the thermal modality by letting each pixel $\mathbf{c} = (x, y)^T$ represent a temperature instead of a visible light intensity. The thermal template $T(\mathbf{c})$ is extracted from a sub-region R_0 of the initial video frame $V(\mathbf{c}, 0)$, which contains the object of interest.




Figure 8. The problem with outline tracking. The subject's head has the same outline (drawn in blue) in both frames *a* and *b*, even though the orientation of the subject's head has changed. If the center of the outline ellipse (the blue dot) was used to extract a temperature reading from the subject in both frames *a* and *b*, the measurements would be taken from different locations.




Figure 9. Creation of the periorbital template from the combination of a uniform grid of points and a variance maximized set of points.

To speed up computation a subset of the pixels within the initial region of interest is used for the template. We chose to compose the set of points in our template $T(\mathbf{c})$ from the union of two subsets, $U(\mathbf{c})$ and $M(\mathbf{c})$ (see Fig 9).

The set $U(\mathbf{c})$ is composed of points uniformly distributed over the region R_0 . The set $M(\mathbf{c})$ is composed of the coldest and hottest pixels from the region R_0 , which effectively maximizes their variance. The motivation for composing $T(\mathbf{c})$ as the combination of $U(\mathbf{c})$ and $M(\mathbf{c})$ is to have both spatial and variance information encoded within the template. Each of $U(\mathbf{c})$ and $M(\mathbf{c})$ contain half of the total number of points designated for $T(\mathbf{c})$.

The set of allowable warps $\mathbf{W}(\mathbf{c}, t; P)$, where $P = (p_1, \dots, p_k)^T$ is a state vector of parameters, that the template can experience depends on the expected motion characteristics of the tracked object. A simple parameterization, assuming an approximately planar object, is described in Eq. (3).

$$\mathbf{W}(\mathbf{c}, t; P) = \begin{pmatrix} p_1 & -p_2 & p_3 \\ p_2 & p_1 & p_4 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \\ 1 \end{pmatrix} \quad (3)$$

The warp $\mathbf{W}(\mathbf{c}, t; P)$ maps the pixel c from the initial reference frame of the template $T(\mathbf{c})$ to a sub-pixel location in the reference frame $V(\mathbf{c}, t)$. The formulation for $\mathbf{W}(\mathbf{c}, t; P)$ can be further condensed to Eq. (4).

$$\mathbf{W}(\mathbf{c}, t; P) = \begin{pmatrix} p_1 \cdot x - p_2 \cdot y + p_3 \\ p_2 \cdot x + p_1 \cdot y + p_4 \end{pmatrix} \quad (4)$$

For the given model we chose the parameterization $P = (p_1, p_2, p_3, p_4)^T$ to allow for translation, scaling, and in-plane rotation. To perform a measurement of the posterior distribution using the previously described thermal template, a triplet containing the state of the template (as described in Eq. (1)) is needed. This necessitates the mapping from the CONDENSATION sample parameters to the warp parameters:

$$\begin{aligned} p_1 &= \cos(\theta) \\ p_2 &= \sin(\theta) \\ p_3 &= x' \\ p_4 &= y' \end{aligned} \quad (5)$$

Equation set (5) includes the parameters of a CONDENSATION sample; x' and y' are the center of mass coordinates for the tracker and θ is the angle of the tracked region with respect to the x -axis of the image coordinate system. Each sample can be measured by using its x , y , and rotation information (i.e., its position on the posterior density) to warp the original template onto the current image and perform a point-wise thermal differencing with the current image. To find the state vector P_t , that represents the optimal template match for the current frame we perform an optimization on the discrete set of state vectors generated from the set of CONDENSATION samples (see Eq. (6)).

$$P_t = \arg \min_P \sum_c [\mathbf{W}(\mathbf{c}, t; P) - T(\mathbf{c})]^2 \quad (6)$$

Initially, we used the estimated posterior mean to generate the state vector P_t that represents the optimal template match for the current frame. In that approach the target's estimated motion over time tended to smooth




Figure 10. A sample of two consecutive frames of the simulated tracking environment. We extract a sub-image of a subject's face from a real thermal image and then project it around the scene according to the parameters of the experiment. The same thermal image is flipped and is used to provide a challenging background for the simulation. A blank background could unfairly bias the tracker to always stay with the facial image.

out. Typically, this is a desirable characteristic because it reduces high frequency noise in the signal, but due to the CONDENSATION process there is a propagation delay whenever there is an abrupt change in the target object's dynamics (e.g., the target was moving left then suddenly moved right). This is because most of the CONDENSATION samples expect the target to continue on its previous motion dynamic, so there might only be one or two samples that, by virtue of the stochastic component of the prediction step, happen to catch the target's correct state. This could still drive the estimated posterior mean close to the true state of the target, if the few samples that correctly located were weighted much heavier than the rest of the samples. To realize this we use the maximum likelihood sample within the tandem tracking framework, where one tracker (the central facial tracker) is acting as a global search tracker and the second tracker (the periorbital tracker) is acting as a localizing tracker.

3.3. Tracking Tests

To ensure that our tandem tracking configuration was capable of performing the tracking required by our data set we performed a series of experiments. The experiments were used to quantify the limits of the periorbital tracker using an artificial environment. Specifically, we took a sample facial image from an actual thermal video frame and applied controlled displacement transformations (see Fig. 10).

By having the tracker track this facial image we were able to quantify its capabilities. We tested the tracker along each type of state vector component (translational and rotational) as well as every combination thereof.




Figure 11. The periorbital tracker's mean error versus increased translational speed along the x dimension. Each experimental run corresponding to a specific target speed was carried out over 200 frames.

1. X Position:

For this experiment we initialized the tracker to the correct initial state vector of the facial image to be tracked and then allowed the facial image to move along the x dimension. We performed multiple experimental runs with different constant target speeds every time. When the target reached a border of the image, it then "bounced off" and continued its motion in the opposite direction. Figure 11 graphs the results of this experiment.

2. Rotation:

For this experiment we initialized the tracker to the correct initial state vector of the facial image to be tracked and then allowed the facial image to rotate at its initial position (in the x - y plane). We performed multiple experimental runs with different rotational target speeds every time. Figure 12 graphs the results of this experiment.




Figure 12. The periorbital tracker's mean error versus increased rotational speed. Each experimental run corresponding to a specific target speed was carried out over 200 frames.




Figure 13. The periorbital tracker's mean error versus increased translational and rotational speeds. Each experimental run corresponding to a specific target speed pair (translational, rotational) was carried out over 200 frames.

5. X position, Y position, and Rotation:

For this experiment we initialized the tracker to the correct initial state vector of the facial image to be tracked and then allowed the facial image to move along the x and y dimensions and simultaneously rotate in the x - y plane. We performed multiple experimental runs with different translational and rotational target speeds every time. Figure 13 graphs the results of this experiment.

The conclusion from Figs. 11–13 is that the tracker performs acceptably (± 2 pixel accuracy with respect to the true target) when the translational speed $u \leq 3$ pixels/frame and the rotational speed $\omega \leq 0.05$ radians/frame. Fortunately, these are within the observed subject's dynamics in our data set. Above this level of motion, the tracker experiences decreasing accuracy at a greater than linear rate.

Another important issue is what happens when the tracker loses its target in the presence of occlusion. To quantify the ability of the tracker to recover after a tracking failure we again employed the artificial environment. This time we modified the environment to allow control over the “visibility” of the facial image. This meant that the facial image could be made to disappear and reappear while following its prescribed transformations. We performed the following experiment:

- Set the translational speed of the target to $u = 3$ pixels/frame. This was chosen because it is the upper movement bound observed in our data set.
- Allowed the target to move around the image and bounce off its borders in the opposite direction.
- Occluded the target (by making it disappear) for a set time interval, and then measured how long it took for the tracker to recover. This was repeated 20 times as the target moved around the image.
- Repeated the experiment 19 times, starting with an occlusion time interval of 1 frame and gradually increasing it to a time interval of 20 frames, in increments of 1 frame at a time.

The results show that the tracker has the ability to recover even from lengthy occlusion incidents, but expectedly the potential for deterioration in tracking performance increases as the duration of the occlusion increases (see the blue curve in Fig. 14).

Typical tracking failures observed during the manual validation lasted between 15–500 frames for normal tracking and between 2–10 frames for tandem tracking. This supports the claim that the tandem tracking configuration is not only preventing tracking failures, but that it is also shortening the time it takes for the tracker to recover after a tracking failure.




Figure 14. The mean number of frames needed for the tandem tracker to recover after an occlusion of the target.

3.4. Tracking Results

To validate the benefit of using our proposed tandem tracking configuration we performed an experiment over the entire set of data selected for deception analysis. This consisted of the interrogation videos of 39 subjects that ranged between 3–5 min in duration each. At the acquisition rate of 60 frames per second this translated to 10,800–18,000 frames per video clip. This resulted in a massive undertaking, necessitating the tracking of over a half a million frames of video in the experiment. Each of these frames was tracked twice, once with the tandem assisted tracker and once with the normal periorbital region tracker.

A common method for measuring the performance of tracking systems is to detect and enumerate tracking failures by manually observing whether the tracker has deviated from the true state of the target object (Kolsch and Turk, 2004; Verma et al., 2003). Alternatively, one can manually set the correct state for the target object in each frame to be used as a ground truth (Eveland et al., 2003). Due to the large size of our dataset, enumeration of tracking failures with manual means was impractical. Consequently, we automated tracking failure enumeration by capitalizing on the natural continuity constraints imposed by human thermo-physiology. For our application, we defined a tracking failure as any fluctuation in the extracted temperature signal greater than the maximum physiologically allowable fluctuation for the given period (Garbey et al., 2004). Figure 15 shows the results of the tracking experiment as quantified by

our automated tracking failure enumeration scheme. The superior performance of the tandem tracking configuration versus normal periorbital tracking is evident.

To verify that the data from Fig. 15 correlates with the true number of tracking errors we manually verified the tracking failures for subjects 4, 22, and 34. In manual verification we defined as tracking failure a deviation from the periorbital region, whereby the tracker's rectangle did not contain both the right and left inner eye corners. In Fig. 15 we reported a small, medium, and large number of tracking failures for these three subjects respectively, which is why we chose these examples as representative. In Table 1 we juxtapose the manually extracted tracking failure data with the automated ones reported in Fig. 15.

Based on the results reported in Table 1, we conclude that our method of automatically reporting tracking failures by using the continuity constraint of human thermo-physiology is valid. In fact, Table 1 signifies that the automated failure enumeration method may introduce a small error, but this is within acceptable limits. In the interest of complete transparency we posted all of the tracking videos on the web at: <http://www.cpl.uh.edu/html/users/jdowdall/html/IJCV/results.htm>. The web page also contains a more detailed analysis of each of the manual validations reported in Table 1.

It should be noted that in our manual validation of the tracking failures we considered a tracking error any deviation of the tracker from the periorbital region for more than 1 frame. This meant that no error was




Figure 15. Tracking failures. The red (dark) bars indicate the number of tracking failures observed when normal tracking is in use and the green (light) bars indicate the number of tracking failures observed when the tandem assisted tracking is in use.

Table 1. Manual validation of tracking errors.

		Number of tracking failures reported in Figure 15		Number of tracking failures recorded manually		Discrepancy		Number of single frame tracking failures ignored in manual validation	
		Normal	Tandem	Normal	Tandem	Normal	Tandem	Normal	Tandem
		Subject 4	0	0	0	0	0	0	0
Subject 22	3	2	8	2	5	0	1	0	
Subject 34	12	6	14	5	2	-1	2	6	

recorded if the tracker recovered to the correct orientation the subsequent frame after a tracking failure (see the last column of Table 1). We defined tracking failures this way because the high frequency noise in temperature extraction resulting from single frame tracking failures was easily removed by the noise suppression algorithm.

We found three main causes for the tracking failures that we observed within our data set:

1. Out-of-plane rotation of the target
2. Occlusion of the target
3. Frames dropped in the video sequence during recording

Surprisingly, by far the most common cause of the tracking failures was due to dropped frames (1 or more frames) in the video. This issue occurred in a few of the video clips because of a hard drive fragmentation problem during the recording of the interrogations.

Subjects wearing glasses were excluded from our data set because temperature extraction around the periorbital region would be impaired by the lenses of

the glasses, which are opaque in thermal infrared. The tracking performance, however, would actually be improved by the presence of glasses, because they would provide a high contrast feature (with respect to the subject's face) for the tracker to use.

4. Physiological Measurement and Noise Suppression

For every frame we compute the mean temperature of the 10% hottest pixels from within the periorbital region of interest. We have found experimentally that this represents the mean temperature on the vasculature in the inner corners of the eyes (see Fig. 16). This correlates with the portion of the periorbital region that has the highest blood flow rate (that's why it is the hottest) and is minimally affected in the imagery by blinking. Therefore, assuming that tracking is accurate, it provides a temperature signal that is indicative of blood flow perfusion in the eye musculature. We consider that the periorbital temperature signal, as defined above, consists of several components:




Figure 16. (a) Facial thermal frame of a subject with the periorbital region of interest superimposed. (b) Blow-up of the selected region of interest with the 10% hottest pixels marked in pink. (c) The periorbital region of interest with its 10% hottest subset superimposed on the facial and ophthalmic arteriovenous complex. The heat convected by the blood flow in this complex is responsible for the elevated temperature with respect to the rest of the periorbital region. Supply of additional blood to the eye muscle is realized through this complex. Therefore, monitoring the conduit in the eye corners is sufficient to detect the ‘fight or flight’ syndrome during stress.

1. A low varying component indicative of the long term trend of blood flow levels, which is of high information value.
2. A mid frequency component, which is associated with temporary disturbances in blood flow caused by stress in specific Question and Answer (Q&A) sessions. This is also of high information value.
3. A high frequency component caused primarily by tracker instability and systemic noise.

Our noise suppression effort is focused on eliminating the high frequency component. We define the mean periorbital temperature signal as $S[n]$, $1 \leq n \leq N$, where N is the number of frames in the thermal video clip. We remove the high frequency component of signal $S[n]$ through a novel application of Fourier decomposition. Specifically, we use the following algorithmic steps:

Step 1. We apply the Fast Fourier Transform (FFT) on signal $S[n]$ to get the power spectrum:

$$S(\omega) = F(S[n]) \quad (7)$$

The FFT method was first introduced by Cooley and Tukey (1965). Thereafter, it was used widely in signal analysis due to its high efficiency in comparison to other methods, such as the solution of linear equations or the correlation method (Smith, 1999). To apply the FFT on the signal $S[n]$, first we use a low order trigonometric polynomial as follows (Dupros et al., 1992):

$$u[n] = S[n] = (\alpha \cos(n) + \beta),$$

$$\text{with } \alpha = \frac{1}{2}(S[0] - S[N-1]),$$

$$\beta = \frac{1}{2}(S[0] + S[N-1]) \quad (8)$$

This ensures that the shift will not affect the stability of the scheme by minimizing the Gibbs phenomenon. Then, we extend $u[n]$ to a $2N$ periodic function as follows. First, we apply the symmetry (Press et al., 1992):

$$\forall n \in [0, N], \quad u[N-n] = -u[n] \quad (9)$$

and second the periodic extension:

$$\forall n \in [0, 2N], \quad \forall k \in \mathbb{Z}, u[n+k2N] = u[n] \quad (10)$$

We apply a classical decimation-in-time (Cooley and Tukey, 1965) 1D base-2 FFT method given in Press et al. (1992).

Step 2. We apply a low-pass filter on the power spectrum $S(\omega)$ to suppress high frequency noise. The filter is constructed using an exponential function:

$$H(e^{j\omega}) = \sum_{n=0}^{\infty} u[n] - e^{j\omega \cdot 1n}. \quad (11)$$

Step 3. We move back to the time domain by applying the inverse FFT on the filtered spectrum $S'(\omega)$:

$$\tilde{u}[n] = F^{-1}(S'(\omega)), \quad (12)$$




Figure 17. The yellow (light) curve shows the raw periorbital temperature signal. The blue (dark) curve shows the low-pass filtered signal after the application of the Fourier decomposition.

followed by re-normalization:

$$\bar{s}[n] = \tilde{u}[n] + (\alpha \cos[n] + \beta). \quad (13)$$

The low-pass filtered signal $\bar{s}[n]$, $1 \leq n \leq N$ (see Figure 17) is fed to the pattern recognition module.

Additionally, an explicit mechanism within our signal processing framework could be employed to ignore the temperature measurements whenever a serious tracking failure occurs. Such a failure could be ascertained by monitoring an appropriately defined fitness function. Because the recovery time of tracking failures in our present data set was generally very small, we chose not to employ such a method. In the future, however, it may prove advantageous as a first level denoising step to directly segment the temperature signal into “correctly tracked” and “incorrectly tracked” sections. Depending on the extent and criticality of the incorrectly tracked sections, they can either be ignored, or the entire signal declared corrupt and no prediction is being attempted.

5. Pattern Recognition

We were interested in deriving a decision scheme, which could separate the Deceptive (D) subjects from the Non-Deceptive (ND) subjects, based on the noise-cleaned temperature signals. Our previous research indicated that elevated levels of stress are associated with elevated periorbital perfusion and temperatures (Pavlidis et al., 2000; Pavlidis et al., 2002). However, the degree of temperature elevation depends not only on the intensity of stress but also on the psycho-physiology of the subject. Different subjects react with different intensity to the same stress stimuli (Pavlidis and Levine, 2002). Therefore, an effective decision scheme has to normalize inter-individual variability through an intra-individual measure. We compute the slope D_i of the filtered temperature signal $\bar{s}[n]$ that corresponds to the entire length of the interrogation for each subject i . This represents the trend or baseline response of the subject to the experiment. We also compute the slope d_{i4} of the portion of the signal that corresponds to the Q&A session with the highest *impact factor*. We define as the *impact factor*, the level of perceived psychological stress a question has per unit time. In other words, for a question to have a high *impact factor* it is not only necessary to be ‘tough,’ but also short and feature an experimentally confirmable lie. It should also be preceded by a series of questions that slowly build pressure

on a subject up to a culmination point. In our interrogation scheme, question 4 qualifies as such a question. In the computation of d_{i4} we consider the signal from the beginning of question 4 to the beginning of the respective answer. Per the psycho-physiological theory this is the interval of interest in a Q&A session.

Our decision scheme is based on the comparison of slope d_{i4} with D_i for subject i :

$$d_{i4} - D_i \rightarrow \begin{cases} > 0 \text{ subject } i \text{ is D,} \\ \leq 0 \text{ subject } i \text{ is ND.} \end{cases} \quad (14)$$

In other words, if the physiological change in the critical question is greater than the baseline change, then the subject is classified as deceptive. This binary classification rule can be relaxed by establishing a third ‘indecision’ class when $d_{i4} - D_i$ is very close to 0.

6. Experimental Setup and Results

The current psychological experiment builds upon the methodology outlined in Frank and Ekman (1997) by Dr. M.G. Frank et al. It is based on the psychological theory of lying developed by Dr. P. Ekman over the last thirty years (Ekman, 1986; Ekman and Maureen, 1991; Ekman, 1996; Ekman, 2001). The experiment was designed by the Communication Lab of Professor M.G. Frank in Rutgers University (currently at University of Buffalo, SUNY). Subjects were a sub-sample of 39 (19 males, 20 females) taken randomly from a larger population of 250 developed by Frank & Ekman. They were all members of politically active groups both on and off campus. The subjects responded to either face to face or e-mail enquiries to participate in a communication skills experiment. When they appeared in the laboratory, they were told that deception is a part of communication, and that is the skill under examination for this experiment. They then completed a number of demographic and personality inventories. Afterward, they were presented with their specific instructions for the deception situation. They were told that there was a check in an envelope hidden down the hall. This check was made out to a group that the subject adamantly opposed (e.g., pro-life subject, check made out to pro-choice group, and vice versa). The subject was told to find the check, and then decide whether or not to ‘steal’ it. If they stole it, they were to put it in their pocket. Whether or not they took the check, they then returned to the experimenter. The subject was then interrogated by a person whom they believe belongs to (or is a strong supporter of) a group they oppose (e.g., the subject who

is pro-life will be interrogated by a person whom they believe is pro-choice). This means that a subject who stole the check lied about it under interrogation. If they fooled the interrogator, they were able to tear up the check (thus depriving their hated group of money), and they obtained \$75 for themselves and this \$100 would go to their group instead. If they failed to convince the interrogator of their innocence (i.e., caught lying), the \$100 check would go to the group that the subject despised, the subject and their group would get no money, and the subject must withstand 30 min of loud, startling blasts of noise delivered via headphones. If the subject did not take the money, and they were judged truthful, they received \$25 for themselves and \$25 for their group, and the despised group would not get \$100 but a smaller sum instead. If the subject was mistakenly judged as lying, then neither they nor their group got any money, the despised group got less money, and the subject was to face the loud startling blasts of noise. Note that no subjects were actually punished with the loud startling blasts of noise and all received at least some money for participation. What was important was for subjects to feel there was a punishment for not being seen as truthful.

The interrogation questions were derived by a panel of homicide and counter-terrorism professionals in conjunction with behavioral scientists. Two of the questions specifically tested the ideas around a popular interrogation technique called the Reid technique (Inbau et al., 2001), whereas the others were designed to elicit maximum useful nonverbal behaviors. In particular, the interrogation was composed of the following questions:

1. Describe in detail everything you saw down the hall.
2. Describe your actions while you were down the hall.
3. What was going through your mind regarding whether to take or leave the check?
4. Did you remove the check from the envelope?
5. Is there anything else you wish to tell me about the check or what you did?
6. Is everything you have told me about the check the truth?
7. What would you say if later I determined you lied to me about the check?
8. Where did you put the check?
9. What should happen to a person who took a check like this?
10. Have you ever told a lie to get out of serious trouble?

11. I know who that check was made out to. It is very understandable that you would not want to see money going to that organization. If I were in your shoes, I probably would have done anything in my power to avoid having money go to this organization. Now tell me the truth, did you take that check or not?
12. If you were to take a polygraph (lie detector machine) test, what do you think the results of this test would be?
13. Why do you think that someone would take this check?

The experiments were conducted through the course of 19 months in a standard room. Three different environmental conditions were realized through the course of the experiments, each spanning several months in duration. Initially, the subjects were illuminated with a strong photographic light, which caused significant elevation of the facial temperature. The setting simulated to a certain degree outdoor conditions, where the subject would have been exposed to the sun. This initial batch of subjects is marked as Tape I in Table 2. Midway in the course of the experiments the photographic light was removed and the room was heated (winter time) to a typical indoor temperature (25°C) while the subjects were dressed moderately. This is marked as Tape II in Table 2. Towards the end of the experiment the room was maintained to a typical indoor temperature (27°C) without heating (spring time) and the subjects were lightly dressed. This is marked as Tape III in Table 2.

In essence, our experiment created a “high stakes” situation for the subjects. Other mock-crime scenarios are much milder and it is questionable how strongly they motivate the subjects (Pavlidis and Levine, 2002). For example, a typical scenario is for the subject to get \$10 if he/she gets caught lying and \$20 if he/she is successfully evasive. This is void of any ideological content and with only a small amount of money at stake. The difference between the two psychological designs is obvious even to a layperson. The other characteristic that set our experiment apart is that it allowed subjects to respond freely; questioning happened only once, and proceeded at a natural speed. In traditional mock crime scenarios binary answers are enforced (‘yes’ or ‘no’) and the exact same questioning is repeated three times in fixed intervals. Therefore, “high stakes” coupled with a much more natural interrogation scheme made our experiment quite realistic. This removes to a large degree the perennial criticism

Table 2. Experimental Results. Successful prediction is marked as '1', while unsuccessful prediction as '0'.

	Subject number	Ground-truth	Original signal	Low-pass filtered signal
TAPE I	1	ND	1	1
	2	D	0	1
	3	D	0	1
	4	D	0	0
	5	D	0	1
	6	D	1	1
	7	D	1	1
	8	D	1	1
	9	D	1	1
TAPE II	10	ND	1	1
	11	ND	1	1
	12	D	1	0
	13	ND	0	1
	14	D	1	1
	15	D	0	1
	16	D	1	1
	17	D	0	1
	18	D	0	1
	19	ND	0	1
	20	D	0	1
	21	ND	1	1
	22	D	1	1
	23	D	0	1
	24	ND	0	0
	25	D	1	1
	26	D	1	1
	27	ND	0	1
	28	D	1	1
TAPE III	29	ND	1	1
	30	ND	1	1
	31	D	0	0
	32	D	1	1
	33	D	1	1
	34	ND	1	1
	35	ND	0	0
	36	ND	0	1
	37	D	0	1
	38	ND	1	1
	39	ND	1	1
Overall Success Rate for D			54.2%	87.5%
Overall Success Rate for ND			60.0%	86.7%
Overall Success Rate			56.4%	87.2%

that such experiments have very little to do with deceptive behavior in the real world.

The interrogations were recorded with the University of Houston's specialized thermal imaging system along with a separate synchronized voice channel. The recorded thermal clips were processed using the methodology outlined in Sections 3–5. The voice channel was used to delineate the interrogation and the critical question 4 interval in the frame sequences. The decisions supported by our scheme were compared against the ground-truth released by the Communication Lab of Professor M.G. Frank. The detailed classification results are shown in Table 2. The results clearly demonstrate the beneficial role of the low-pass filtering algorithm. It is also interesting to note that the method is more successful in detecting Deceptive (D) rather than Non-Deceptive (ND) subjects.

Based on the trend exhibited by the current experimental results, it is reasonable to expect that the performance of the method can improve simply by improving the accuracy of the tracking algorithm and the quality of the thermal imaging system.

7. Conclusions

We have presented a new methodology for recovering the periorbital signal from thermal facial sequences. It uses tandem tracking and noise suppression. Tandem tracking ensures that the periorbital region of interest is tracked accurately by weighing in feedback from the central facial region tracker. Noise suppression removes the high frequency components from the periorbital signal, which are due to tracking imperfections and electronic interference. The filtered signal is classified as Deceptive (D) or Non-Deceptive (ND) based on the slope differential between the entire curve and portions thereof (question 4 in our case). Subjects who exhibit higher change in critical moments with respect to their interrogation baseline are classified as D. Our method scored an 87.2% accurate prediction rate on 39 subjects interrogated by the Communication Lab at Rutgers University per the designed mock-crime protocol. This is on par with the performance exhibited by seasoned interrogation experts.

This work follows up on the authors' previously reported methodology and results by the way of the DOD PI experiments (Pavlidis and Levine, 2002). This time the sample was double in size (39 versus 20 subjects) and the conditions were quite realistic. The old method of amplifying the signal through transformation from

the temperature to the blood flow domain was abandoned in favor of a noise suppression method. The periorbital region of interest has been more clearly defined as the inner corner of the eyes, where the facial and ophthalmic arteriovenous complex is supplying blood flow to the eye musculature. The importance of tissue tracking to cope with motor and fidgeting motion has been brought to the forefront.

Our research further demonstrates the efficacy of automated lie detection by shedding more light into the associated physiological mechanism and coping with increasingly more realistic experiments.

A number of extensions to this work could be made. To improve the performance of the tracking algorithm, the information fusion between the central region tracker and the periorbital region tracker could be extended to be bi-directional. Also, the thermal template could be updated over time (Matthews et al., 2004) to conform more closely to the dynamic perfusion of the subject, resulting in better tracking. Additionally, we could extend the parameterization of the tracker to full affine transformations to better represent out-of-plane rotations. In the future, we plan on processing more subjects to ensure that the method continues to scale up in larger data sets and is adaptable to diverse interrogation scenarios.

There are still a number of limitations to this technology. For example, eyeglasses were an impermeable obstacle to the thermal signal and thus subjects with glasses could not be analyzed. However, this study clearly shows that our system and method have a great deal of promise in applied lie catching situations.

Acknowledgments

This research was supported by DARPA/ONR grant #N00014-03-1-0622 ('Thermal Imaging Experimentation and Analysis for Deception Detection') and by the University of Houston start-up funds of Professor Ioannis T. Pavlidis. Additional support was provided by DARPA/ONR grant #N00014-02-1-0709 ('Robust Face Recognition from a High Stakes Paradigm') to Professors Mark G. Frank and Paul Ekman, DARPA/ONR grant #N00014-03-1-0847 ('Remote Measures in the Detection of Deception in a Counter-Terror Situation'), and HSARPA grant # ITT-180120 ('Examining Deception in a Counter-Terror Paradigm') to Professor Mark G. Frank. The views expressed in this paper do not necessarily reflect the views of the funding Agencies. We would like to express our gratitude

to Dr. Ralph Chatham (DARPA Program Manager), Dr. Jonathon Philips (DARPA Program Manager), Dr. Thomas McKenna (ONR Program Manager), Larry Willis (HSARPA Program Manager), and Dr. Andrew Ryan (DOD PI) for their support. We would also like to thank Dr. Arcangelo Merla (University G. D'Annunzio), Amy Marie Keller, Sinuk Kang (Rutgers University), and Professor Marc Garbey (University of Houston) for their valuable help.

References

- Baker, S. and Matthews, I. 2001. Equivalence and efficiency of image alignment algorithms. In *Proceedings of the 2001 IEEE Conference on Computer Vision and Pattern Recognition*, Vol. 1, pp. 1090–1097.
- Cooley, J.W. and Tukey, J.W. 1965. An algorithm for the machine calculation of complex Fourier series. *Mathematics of Computation*, 19:297–301.
- Committee to Review the Scientific Evidence on the Polygraph—Board on Behavioral, Cognitive, and Sensory Sciences and Committee on National Statistics—Division of Behavioral and Social Sciences and Education – National Research Council of the National Academies. 2003. *The Polygraph and Lie Detection*, National Academies Press, Washington D.C.,
- Doucet, A., DeFreitas, N., and Gordon, N. (Eds.) 2001. *Sequential Monte Carlo Methods in Practice*. Springer-Verlag.
- Drummond, P. and Lance, J. 1987. Facial flushing and sweating mediated by the sympathetic nervous system. *Brain*, 110:793–803.
- Dupros, F., Fitzgibbon, W.E., and Garbey, M. 2004. *A filtering technique for system of reaction diffusion equations*. Technical Report UH-CS-04-02, Department of Computer Science, University of Houston.
- Ekman, P. 1986. *Telling Lies*. Berkeley Publishing Group.
- Ekman, P. and Maureen, O. 1991. Who can catch a liar? *American Psychologist*, 46(9):913–20.
- Ekman, P. 1996. Why don't we catch liars? *Social Research*, 63(3):801–17.
- Ekman, P. 2001. *Telling Lies: Clues to Deceit in the Marketplace, Politics, and Marriage*, 2nd ed, W.W. Norton and Company.
- Eveland, C., Socolinsky, D., and Wolff, L. July 2003. Tracking human faces in infrared video. *Image and Vision Computing*, 21:578–590.
- Frank, M.G. and Ekman, P. 1997. The ability to detect deceit generalizes across different types of high-stake lies. *Personality and Social Psychology*, 72:1429–1439.
- Garbey, M., Merla, A., and Pavlidis, I. 2004. Estimation of blood flow speed and vessel location from thermal video. In *Proceedings of the 2004 IEEE Computer Vision and Pattern Recognition*, Vol. 1, pp. 356–63.
- Inbau, F.E., Reid, J.E., Buckley, J.P., and Jayne, B.C. 2001. *Criminal Interrogations and Confessions*. Aspen Publishers, 4th edn.
- Isard, M. and Black, A. 1998a. Condensation—conditional density propagation for visual tracking. *International Journal of Computer Vision*, 19(1):5–28.
- Isard, M. and Blake, A. June 1998b. ICONDENSATION: Unifying low-level and high-level tracking in a stochastic framework. In *Proceedings of the 5th European Conference on Computer Vision*, Freiburg, Germany, Vol. 1, pp. 893–908.
- Isard, M. and MacCormick, J. 2001. BraMBLe: A Bayesian multiple-blob tracker. In *Proceedings of the 8th International Conference on Computer Vision*, Vol. 2, pp. 34–41.
- Kleinmuntz, B. and Szucko, J.J. 1984. Lie detection in ancient and modern times. *American Psychologist*, 39:766–776.
- Knight, J. April 15 2004. The truth about lying. *Nature*, 428:692–94.
- Kolsch, M. and Turk, M. 2004. Fast 2D hand tracking with flocks of features and multi-cue integration. In *Proceedings of the 2004 IEEE Conference on Computer Vision and Pattern Recognition*, Vol. 10, pp. 158–164.
- MacCormick, J. and Isard, M. 2000. Partitioned sampling, articulated objects, and interface-quality hand tracking. In *Lecture Notes in Computer Science*, Springer-Verlag, London, U.K., Vol. 1843, pp. 3–19.
- MacCormick, J. and Blake, A. 2000. A probabilistic exclusion principle for tracking multiple objects. *International Journal of Computer Vision*, 39(1):57–71.
- Martini, F.H., Ober, W.C., Garrison C.W., Welch, K., and Hutchings, R.T. 2001. *Fundamentals of Anatomy and Physiology*, Prentice Hall, 5th edn, Upper Saddle River, New Jersey.
- Matthews, I., Ishikawa, T., and Baker, S. June 2004. The template update problem. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(6):810–815.
- Pavlidis, I., Levine, J., and Baukol, P. June 16 2000. Thermal imaging for anxiety detection. In *Proceedings of the 2000 IEEE Workshop on Computer Vision Beyond the Visible Spectrum: Methods and Applications*, Hilton Head Island, South Carolina, pp. 104–109.
- Pavlidis, I., Eberhardt, N.L., and Levine, J. January 3 2002. Human behavior: Seeing through the face of deception. *Nature*, 415(6867):35.
- Pavlidis, I. and Levine, J. November/December 2002. Thermal image analysis for polygraph testing. *IEEE Engineering in Medicine and Biology Magazine*, 21(6):56–64.
- Press, W.H., Teukolsky, S.A., Vetterling, W.T., and Flannery, B.F. 1992. *Numerical Recipes in C*. 2nd edn. Cambridge University Press, New York, New York.
- Smith, S.W. 1999. *The Scientist's and Engineer's Guide to Digital Signal Processing*. 2nd edn. California Technical Publishing, San Diego, California.
- Sokolov, E.N. and Cacioppo, J.T. 1997. Orienting and defense reflexes: Vector coding the cardiac response. In *Attention and Orienting: Sensory and Motivational Processes*, P.J. Lang, and R.F. Simons, (Ed.) Lawrence Erlbaum Associates, Mahwah, New Jersey, pp. 1–22.
- Verma, R., Schmid, C., and Mikołajczyk, K. October 2003. Face detection and tracking in a video by propagating detection probabilities. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(10):1215–28.
- Yankee, W.J. 1965. An investigation of sphygmomanometer discomfort thresholds in polygraph examinations. *Police*, 9:12–18.

Copyright of *International Journal of Computer Vision* is the property of Springer Science & Business Media B.V. and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.