# 4

# Face Recognition under the Skin

Pradeep Buddharaju and Ioannis Pavlidis

## 4.1 Introduction

Face recognition stands as the most appealing biometric modality, since it is the natural mode of identification among humans and is totally unobtrusive. At the same time, however, it is one of the most challenging modalities (Zhao et al. 2003). Several face recognition algorithms have been developed in recent years, mostly in the visible and a few in the infrared domains. A serious problem in visible face recognition is light variability, due to the reflective nature of incident light in this band. This can clearly be seen in Figure 4.1. The visible image of the same person in Figure 4.1(a) acquired in the presence of normal light appears totally different from that in Figure 4.1(b), which was acquired in low light.

Many of the research efforts in thermal face recognition were narrowly aiming to see in the dark or reduce the deleterious effect of light variability (Figure 4.1) (Socolinsky et al. 2001; Selinger and Socolinsky 2004). Methodologically, such approaches did not differ very much from face recognition algorithms in the visible band, which can be classified as appearance-based (Chen et al. 2003) and feature-based (Buddharaju et al. 2004). Recently, attempts have been made to fuse the visible and infrared modalities to increase the performance of face recognition (Socolinsky and Selinger 2004a; Wang et al. 2004; Chen et al. 2005; Kong et al. 2005).

The authors have presented a physiological facial recognition method that promotes a different way of thinking about face recognition in the thermal infrared (Buddharaju et al. 2006, 2007; Buddharaju and Pavlidis 2008). This work has shown that facial physiological information, extracted in the form of a superficial vascular network, can serve as a good and time-invariant feature vector for face recognition. However, the methodology in that prior work had some weak points. The recognition performance reported in past experiments

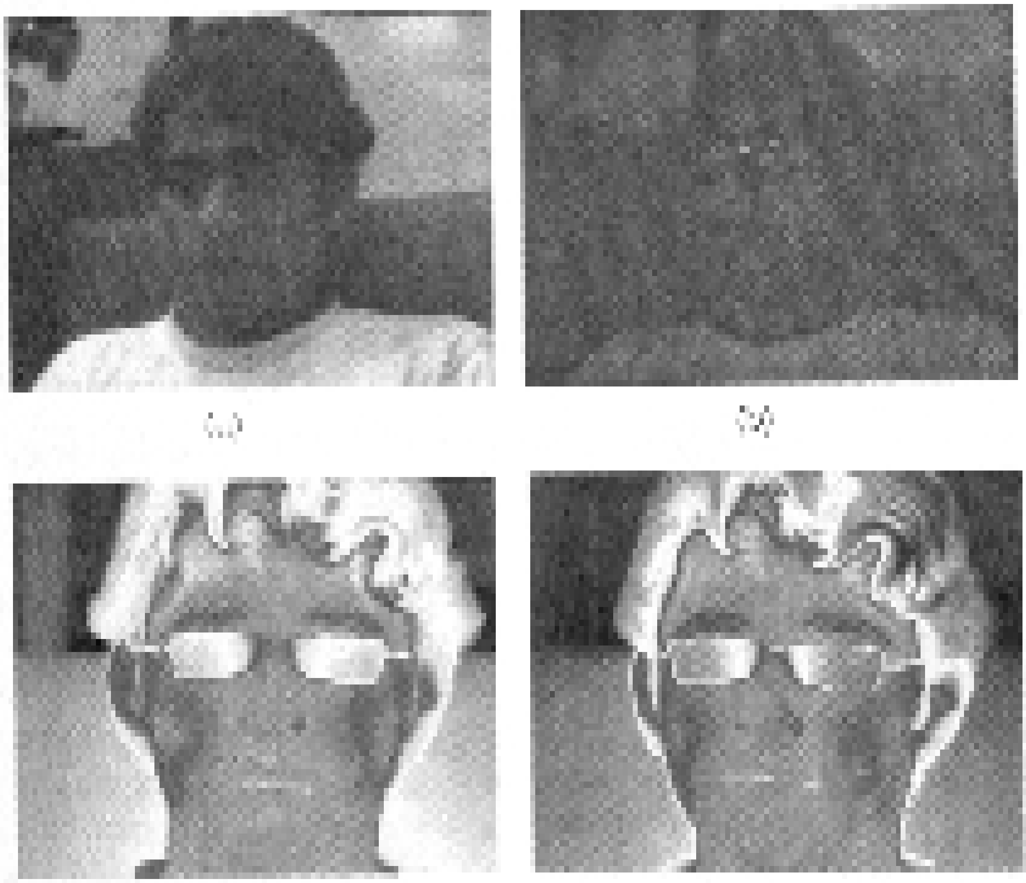

(a) (b) (c) (d)

Figure 4.1. Example showing illumination effect on visible and thermal infrared images. All the images were acquired from the same subject at the same time. (a) Visible image in normal light. (b) Visible image in low light. (c) Thermal infrared image in normal light. (d) Thermal infrared image in low light.

(Buddharaju et al. 2007) can be substantially improved by curing these weaknesses. This chapter presents an advanced methodological framework that aspires to transform physiological face recognition from feasible to viable. Specifically, the main contributions in the chapter are:

- A new vessel segmentation post processing algorithm that removes fake vascular contours detected by the top hat vessel segmentation method.
- A new vascular network matching algorithm that is robust to nonlinear deformations due to facial pose and expression variations.
- Extensive comparative experiments to evaluate the performance of the new method with respect to previous methods.

The rest of the chapter is organized as follows. Section 4.2 presents an overview of the new methodology. Section 4.3 presents in detail the vessel segmentation post-processing algorithm. Section 4.4 discusses the new vascular

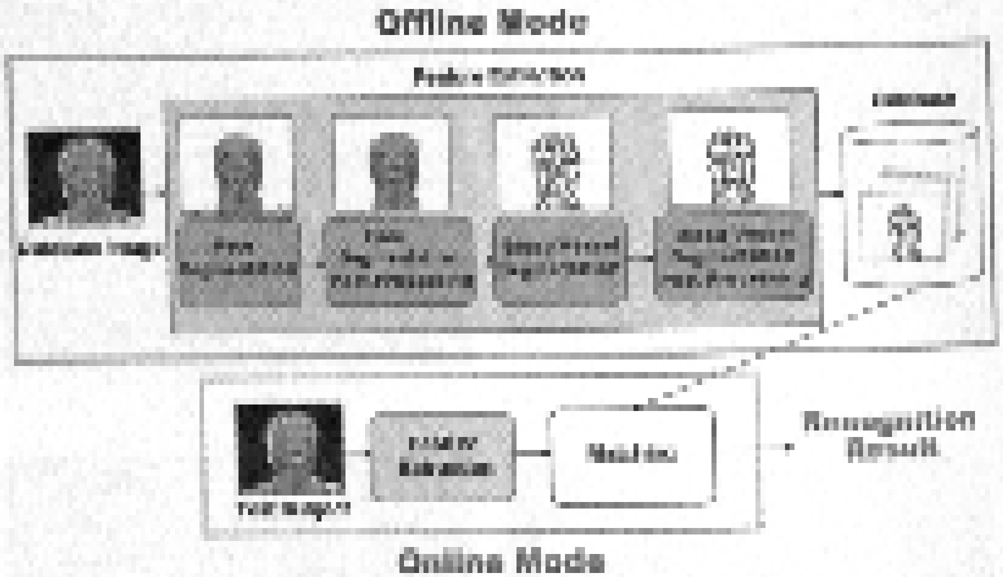


Figure 4.2. Methodological architecture.

network matching algorithm. Section 4.5 presents the experimental results and attempts a critical evaluation. The chapter concludes in Section 4.6.

## 4.2 Methodology

Figure 4.2 shows the methodological architecture. The method operates in the following two modes:

*Off-line Mode*: The thermal facial images are captured by a mid-wave infrared (MWIR) camera. For each subject to be registered in the database, the feature extraction algorithm extracts the feature vector from the facial image and links it to the subject's record. The feature extraction algorithm has four steps.

First, a Bayesian *face segmentation* separates facial tissue from background. Second, *face segmentation post-processing* corrects face segmentation errors, which are due to occasional overlapping between portions of the tissue and background distributions. Third, a top-hat *vessel segmentation algorithm* extracts the vascular network from the facial segment after an anisotropic diffusion clarifies fuzzy edges due to heat diffusion. These three steps have been adopted from Buddharaju et al. (2007) and are briefly presented in this section. Fourth, a new *vessel segmentation post-processing* algorithm, which is one of the chapter's contributions, corrects vessel segmentation. The vessel segmenter occasionally is fooled by areas of high contrast (e.g., hair-line and skin edges) and reports them as vascular contours. These fake vascular contours participate in the matching process with deleterious effects.

*On-line Mode*: Given a query image, its vascular network is extracted using the feature extraction algorithm outlined in the off-line mode, and it is matched against vascular networks stored in the database. The new matching *algorithm*, which is another of this chapter's contributions, has two stages.

First, a *face pose estimation* algorithm estimates the pose of the incoming test image, which is required to calculate the vascular network deformation between it and database images. Second, the *dual bootstrap ICP* matching algorithm registers the test and database vascular networks. The matching score between the two depends on the *amount of overlapping*.

### 4.2.1 Face Segmentation

Because of its physiology, a human face consists of "hot" parts that correspond to tissue areas that are rich in vasculature and "cold" parts that correspond to tissue areas with sparse vasculature. This casts the human face as a bimodal temperature distribution entity, which can be modeled using a mixture of two Normal distributions. Similarly, the background can be described by a bimodal temperature distribution with walls being the "cold" objects and the upper part of the subject's body dressed in clothing being the "hot" object. The consistency of bimodality across subjects and image backgrounds is striking. We approach the problem of delineating facial tissue from background using a Bayesian framework since we have a priori knowledge of the bimodal nature of the scene.

We call $\theta$ the *parameter of interest*, which takes two possible values (skin $s$ or background $b$) with some probability. For each pixel $x$ in the image at time $t$, we draw our inference of whether it represents skin (i.e., $\theta = s$) or background (i.e., $\theta = b$) based on the posterior distribution $p^{(t)}(\theta|x_t)$ given by

$$p^{(t)}(\theta|x_t) = \begin{cases} p^{(t)}(s|x_t), & \text{when } \theta = s, \\ p^{(t)}(b|x_t) = 1 - p^{(t)}(s|x_t), & \text{when } \theta = b. \end{cases} \tag{4.1}$$

We develop the statistics only for skin, and then the statistics for the background can easily be inferred from Eq. (4.1).

According to Bayes' theorem,

$$p^{(t)}(s|x_t) = \frac{\pi^{(t)}(s)\, f(x_t|s)}{\pi^{(t)}(s)\, f(x_t|s) + \pi^{(t)}(b)\, f(x_t|b)}. \tag{4.2}$$

Here $\pi^{(t)}(s)$ is the prior skin distribution and $f(x_t|s)$ is the likelihood for pixel $x$ representing skin at time $t$. In the first frame ($t = 1$) the prior distributions for skin and background are considered equiprobable:

$$\pi^{(1)}(s) = \frac{1}{2} = \pi^{(1)}(b). \tag{4.3}$$

For $t > 1$, the prior skin distribution $\pi^{(t)}(x)$ at time $t$ is equal to the posterior skin distribution at time $t-1$:

$$\pi^{(t)}(x) = p^{(t-1)}(s|x, \theta). \tag{4.4}$$

The likelihood $f(x|\theta)$ of pixel $x$ representing skin or background is given by

$$f(x|\theta) = \sum_{i=1}^{2} \alpha_i [\text{illegible}]. \tag{4.5}$$

where the vector of parameters [illegible] $\mu_i$ [illegible] and [illegible] are initialized [illegible] using the EM algorithm. [illegible] represents the total number of [illegible] of subjects that we call the training set. [illegible] segments for each of the $N$ frames (skin and background), which yields [illegible] skin and [illegible] background pixels.

Sometimes, part of the subject's skin can be erroneously classified as background, and a few background pixels can be erroneously classified as skin. This is due to [illegible] overlapping between portions of the skin and background distributions. The isolated nature of these mislabeled pixels makes them easily correctable through postprocessing. We apply a three-step postprocessing algorithm to the binary segmented image. Using [illegible] background correction, we find the mislabeled pixels in foreground and background and remove them.

### 4.2.2 Blood Vessel Segmentation

Once a face is delineated from the rest of the scene, the segmentation of superficial blood vessels from the facial tissue is carried out in the following steps:

*Step 1:* Process the image to reduce noise and enhance the edges.

*Step 2:* Apply morphological operations to localize the superficial vasculature.

The weak signal edges formed due to heat diffusion at the blood vessels can be handled effectively using anisotropic diffusion. The anisotropic diffusion filter is formulated as a process that enhances object boundaries by performing intra-region as opposed to inter-region smoothing.

One mathematical representation describes this process as:

$$\frac{\partial I(\bar{x},t)}{\partial t} = \nabla(c(\bar{x},t)\nabla I(\bar{x},t)). \tag{4.6}$$

In this case, $I(x, t)$ is the thermal infrared image, $x$ refers to the spatial dimension, $t$ is time, and $c(x, t)$ is called the diffusion function. The discrete version of the anisotropic diffusion filter of Eq. (4.6) is as follows:

$$\begin{aligned} I_{t+1}(x, y) = I_t + \frac{1}{4} \times [&c_{N,t}(x, y)\nabla I_{N,t}(x, y) \\ &+ c_{S,t}(x, y)\nabla I_{S,t}(x, y) + c_{E,t}(x, y)\nabla I_{E,t}(x, y) \\ &+ c_{W,t}(x, y)\nabla I_{W,t}(x, y)] \end{aligned} \quad (4.7)$$

The four diffusion coefficients and four gradients in Eq. (4.7) correspond to four directions (North, South, East, and West) with respect to the location $(x, y)$. Each diffusion coefficient and the corresponding gradient are calculated in the same manner. For example, the coefficient along the north direction is calculated as follows:

$$c_{N,t}(x, y) = \exp\left(\frac{-\nabla I_{N,t}^2(x, y)}{k^2}\right), \quad (4.8)$$

where $\nabla I_{N,t} = I_t(x, y - 1) - I_t(x, y)$.

Image morphology is then applied to the image to [illegible] areas that are relatively brighter than compared with their surrounding areas. We [illegible] top-hat segmentation, which is a [illegible]. Top-hat segmentation takes two forms: white top-hat segmentation that enhances the bright objects in the image or black top-hat segmentation that enhances dark objects. In our case, we are interested in the white top-hat segmentation because it helps to enhance the bright ("hot") ridge like structures corresponding to the blood vessels. In this method the original image is first opened and then this opened image is subtracted from the original image.

$$\begin{aligned} I_{open} &= (I \ominus S) \oplus S, \\ I_{top} &= I - I_{open} \end{aligned} \quad (4.9)$$

where $I$, $I_{open}$, $I_{top}$ are the original, opened, and white top-hat segmented images, respectively, $S$ is the structuring element, and $\ominus$, $\oplus$ are morphological erosion and dilation operations, respectively.

## 4.3 Vessel Segmentation Post-Processing

A superficial vessel has higher temperature than surrounding [illegible] due to [illegible]

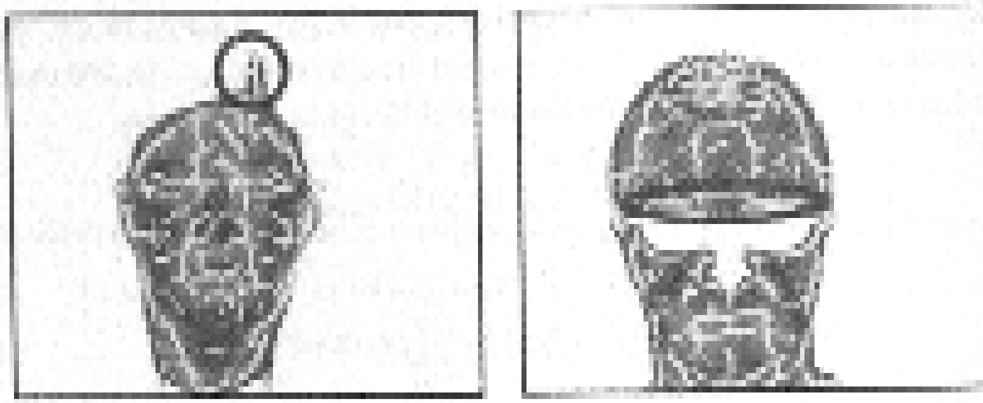

(a) (b)

Figure 4.3: Errors from the top-hat segmentation algorithm for vascular network extraction. (a) outliers due to facial hair; (b) outliers due to glasses.

The top-hat segmentation algorithm is successful in localizing vessels because it targets transitions from hot to cold to hot. In some instances, such transitions are not due to the presence of vessels. Examples include tissue between hair lines or tissue between glasses and eyebrows (see Figure 4.3).

It is essential to detect and remove these outliers from the vascular network before applying a matching algorithm. To study in depth the properties of these outliers, the authors selected 25 representative subjects from the University of Houston dataset and analyzed the segmentation errors. Specifically, the authors identified the locations of both true vessels and outliers. Then they drew measurement lines across each of the vessels and outliers and plotted the corresponding temperature profiles. They noticed that the difference between minimum and maximum temperature values was much larger in outliers than in true vessels. Indeed, the gradient in outliers is quite steep (several degrees C), as it is formed between facial hair or glasses and tissue. By contrast, in true vessels, the gradient is quite small (only tenths of degrees C), as it is formed between the projection of the vessel's lumen and surrounding tissue (Figure 4.4). Figure 4.5 shows the difference between minimum and maximum temperatures ($T_{max} - T_{min}$) across all the selected line profiles from the 25 representative subjects.

The new segmentation post-processing algorithm removes outliers based on the above findings. Specifically, it carries out the following steps:

*Step 1:* Skeletonize the vascular network to one pixel thickness.

*Step 2:* Draw a normal parallelogram across each skeleton pixel and gather all the pixels covered by this parallelogram.

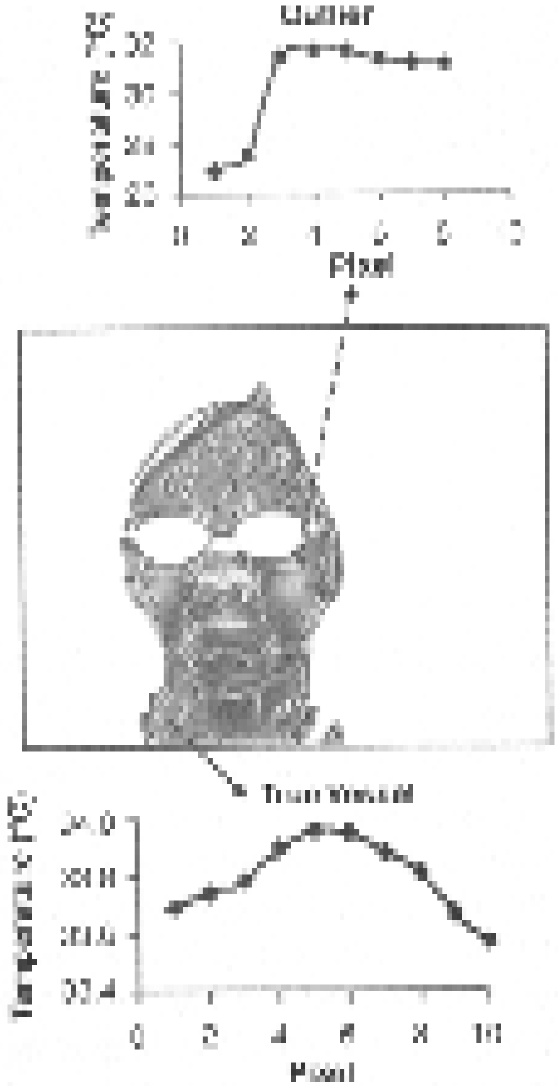


Figure 4.6. Temperature profiles of lines crossing an outlier and true vessel.

*Step 5:* Apply K-Means (with $K = 2$) on the pixels covered by the parallelogram. If the difference between the centers of each cluster is greater than 1.5, then mark it as an outlier pixel.

*Step 6:* Remove all the branches from the vascular network that have more than 50% of their pixels marked as outliers.

After deleting the outliers from the vascular network, the remaining vascular map can be stored in the database.

## 4.4 Vascular Network Matching

The matching method presented in Buddharaju et al. (2007) could not cope with nonlinearities in the deformation of the vascular network, because of variations

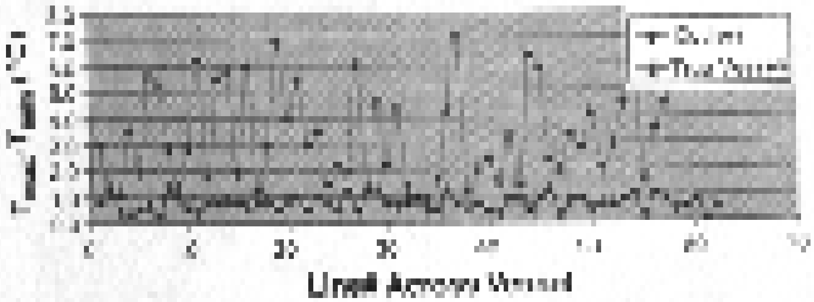


Figure 4.5 Difference between minimum and maximum temperatures for all the selected line profiles.

in facial pose and expressions. This chapter presents a new vascular network matching algorithm that is robust to nonlinear deformations.

### 4.4.1 Registration of Vascular Networks

The aim is to register the vascular network of the test image with that of the database image, so that they can be aligned. The Iterative Closest Point (ICP) algorithm has appealing properties for point-based registration (Besl and McKay 1992). ICP requires proper initialization, as different instantiations of the algorithm use different combinations of image points, distance metrics, and transformation models. In Stewart et al. (2003), developed a variation of the ICP algorithm, called dual bootstrap ICP, that works well when initialization provides just a "toe hold" on the correct estimate of the transformation and successfully registers elongated structures such as vasculature. Specifically, they reported good results from registration of retinal vascular images in the visible band. Since superficial vasculature extracted in the thermal infrared has morphological resemblance to retinal vasculature in the visible, the authors adopted the dual bootstrap algorithm for the registration task at hand.

After successful registration of the test and database vascular images, the matching score is computed based on the number of overlapping vessel pixels. If $I_{test}$ represents the test vascular image with $N_{test}$ vessel pixels, and $I_{db}$ represents the database vascular image with $N_{db}$ vessel pixels, the matching score is

$$\text{Score} = \frac{N_{overlap}}{\max(N_{test}, N_{db})} \times 100, \qquad (4.7)$$

where $N_{overlap}$ represents the number of vessel pixels in $I_{test}$ with a corresponding vessel pixel in $I_{db}$ within a certain distance.

Figure 4.6 shows some examples of the performance of the dual bootstrap ICP algorithm in registering vascular images in the thermal infrared. The example at the bottom row of the figure features substantial pose variation between

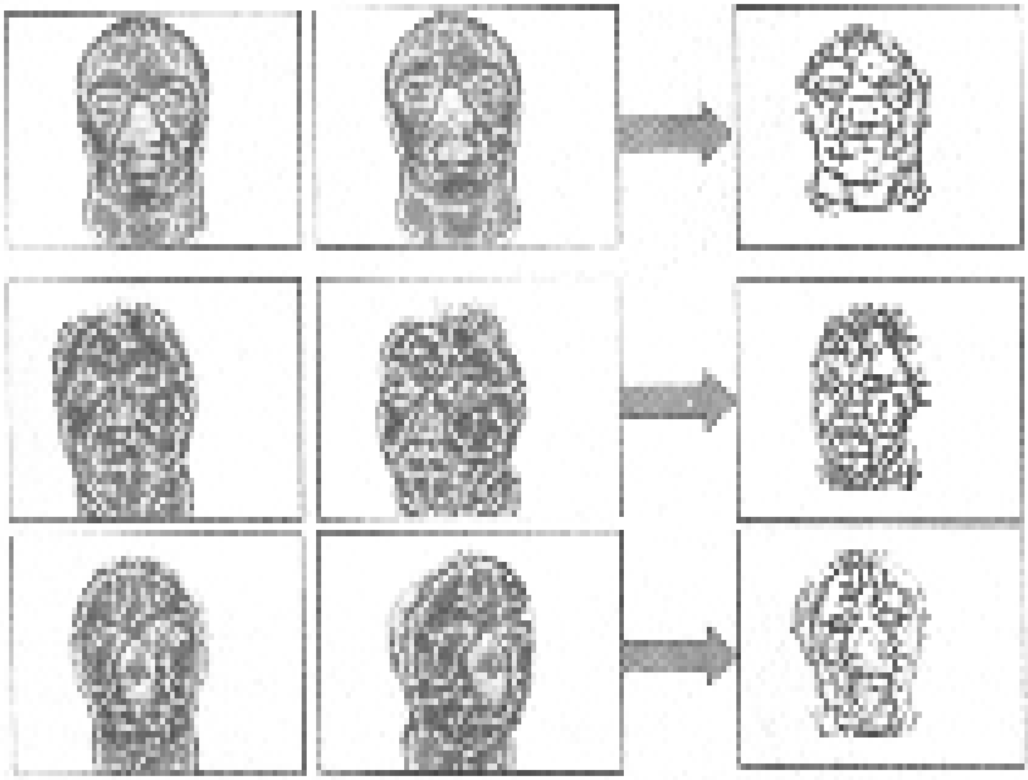

(a) (b) (c)

Figure 4.6. Representative positive matching examples: (a) test vascular networks; (b) database vascular networks; (c) registration results.

the test and database images. The larger the pose variation, the more difficult it becomes for the dual bootstrap ICP to cope successfully. This can be improved by estimating the pose of the test and database images and setting accordingly the threshold value ($Th$) used to accept or reject the test image. The pose estimation algorithm that was developed for this purpose is presented in the next section.

## 4.4.2 Face Pose Estimation

At a neutral pose ($\theta_{face} = 0°$) the nose is at the center of the face, that is, the position of the nose is halfway between the left and right ends of the face. From the face segmentation algorithm presented in Buddharaju et al. (2007), one can find the left and right ends of the face. Hence, if one localizes the nose, he or she can estimate the facial pose.

In a thermal infrared image, the nose is typically at a gradient with its surroundings, as shown in Figure 4.7. This is because of the tissue's shape (tubular cavity), its composition (cartilage), and the forced circulation of air, due to breathing. The combination of all three creates a nasal thermal signature that is different from that of the surrounding solid, soft tissue.

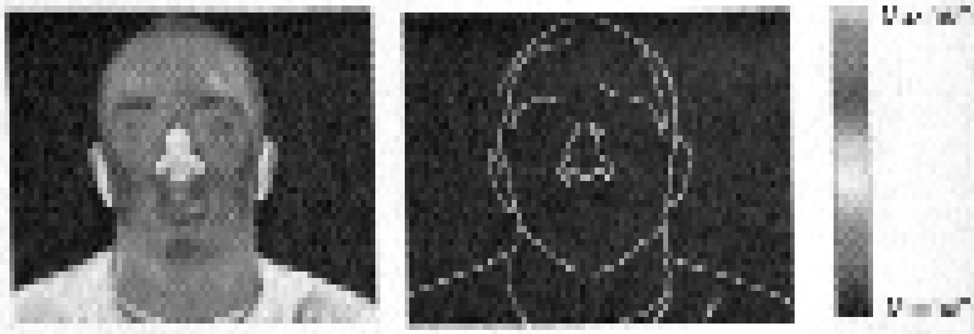

(a) (b) (c)

Figure 4.7. Nose edge extraction from thermal facial image: (a) thermal facial image; (b) edges extracted using the Canny edge detection algorithm; (c) color map.

By using a standard edge detection algorithm, one can extract the nose edge from the facial image. The next step is to search for the nose edge model in the facial edge map. The authors used a Hausdorff-based matching algorithm to localize the nose model in the face edge image (Huttenlocher et al. 1993). Figure 4.8 shows performance examples of the face pose estimation algorithm.

## 4.5 Experiments

The authors conducted several experiments to validate the performance of the new physiological face recognition method. This section presents the experimental setup and results in detail.

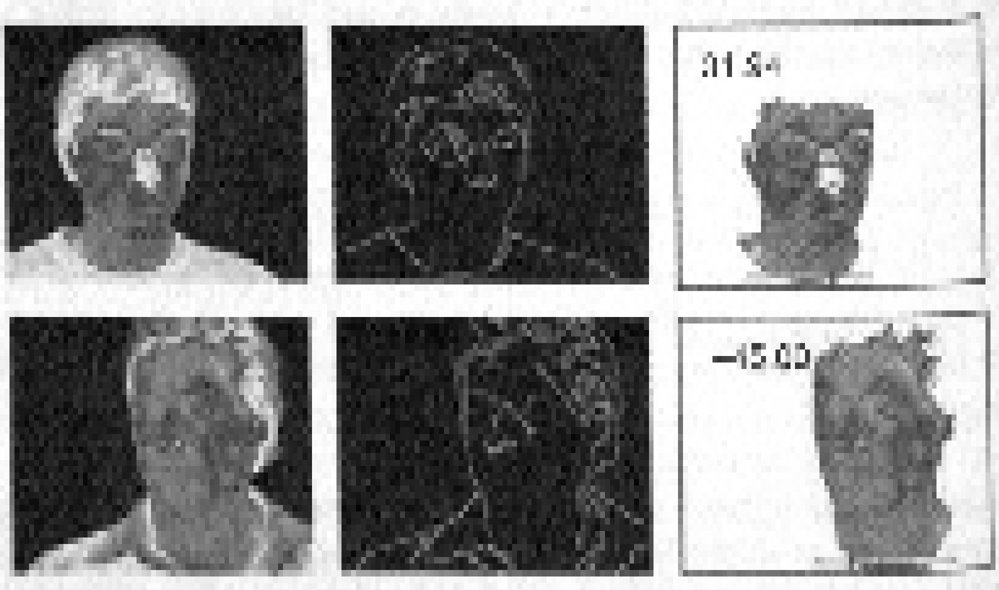


(a) (b) (c)

Figure 4.8. Pose estimation performance examples: (a) thermal facial images; (b) nose detection using Hausdorff-based matching; (c) pose estimates.

(a)

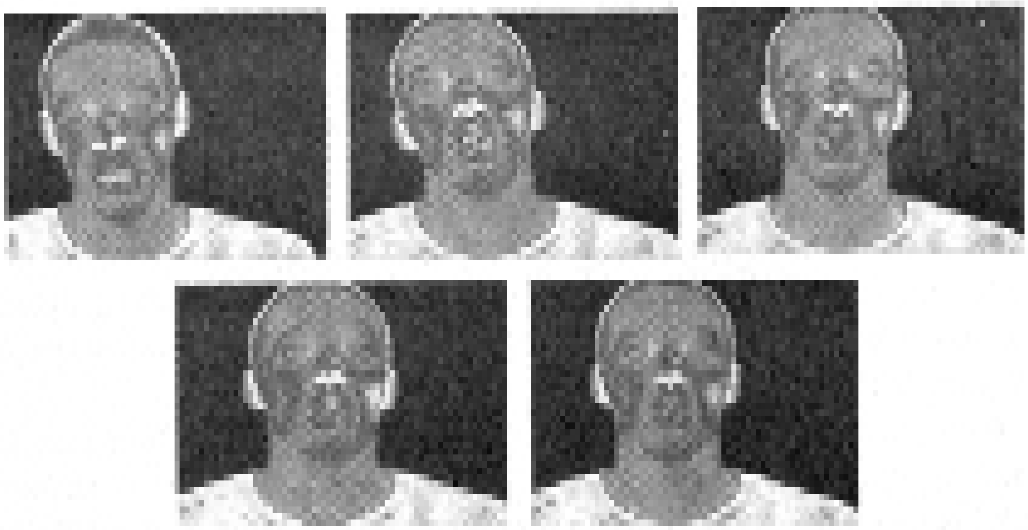

(b)

Figure 4.9. Sample subject from FEDS dataset: (a) one gallery image at neutral pose and expression; (b) five probe images at varying facial expressions.

### 4.5.1 Experiments on the University of Houston Database

The authors collected a substantial thermal facial dataset for the purposes of this evaluation. This set, known as the University of Houston (UH) database, has thermal facial images of varying expressions and poses from 300 subjects. The images were captured using a high-quality mid-wave infrared (MWIR) camera.

#### 4.5.1.1 Facial Expression Dataset

To test the performance of the new method in the presence of varying facial expressions between gallery and probe images, the authors created the facial expression dataset (FEDS) out of the UH database as follows. From each of the 300 subjects one frontal facial image at 0° pose and neutral expression was used as a gallery image. Then five different facial images at ~0° pose but with varying facial expressions were used as probe images for each subject. Hence, FEDS has a total of 300 gallery images and 1500 probe images from 300 subjects. Figure 4.9 shows a sample subject set from FEDS.

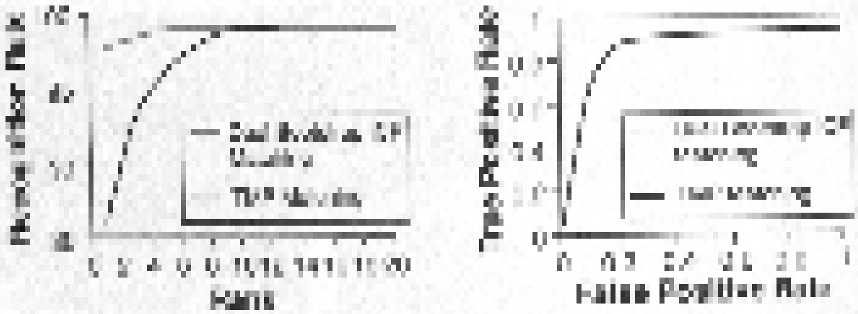


(a) (b)

Figure 4.10. Experimental results on the FEDS dataset using dual bootstrap ICP versus TMP matching: (a) CMC curves; (b) ROC curves.

Figures 4.10(a) and 4.10(b) show, respectively, the Cumulative Match Characteristic (CMC) and Receiver Operating Characteristic (ROC) curves of the FEDS experiments using the thermal minutia point (TMP) matching algorithm reported in Buddharaju et al. (2007) (a fingerprinting variety) versus the dual bootstrap ICP matching algorithm.

The results demonstrate that the dual bootstrap ICP outperforms the TMP matching algorithm. In the case of TMP matching, the CMC curve shows that rank 1 recognition is 86%, whereas for the for dual bootstrap ICP is 97%. Also the dual bootstrap ICP matching method achieves a high positive detection rate at very low false detection rates, as shown in Figure 4.10(b). This indicates that the ICP matching algorithm is highly robust to deformations caused in the vascular network by facial expression variations.

#### 4.5.1.2 Facial Pose Dataset

To test the performance of the new method in the presence of varying pose between gallery and probe images, the authors created the facial pose dataset (FPDS) out of the UH database as follows. From each of the 300 subjects one frontal facial image at 0° pose and neutral expression was used as a gallery image. Then four different facial images at neutral facial expression but at varying poses between −30° and 30° were used as probe images. Hence, FPDS has a total of 300 gallery images and 1200 probe images from 300 subjects. Figure 4.11 shows a sample subject set from FPDS.

Figures 4.12(a) and 4.12(b) show the results of the FPDS experiments using the dual bootstrap ICP matching algorithm. The results demonstrate that the algorithm copes well with facial pose variations between gallery and probe images. Specifically, the CMC curve shows that rank 1 recognition is 99% and the ROC curve shows that it requires a false acceptance rate over 2% to reach a positive acceptance rate above the 90% range. One can notice that the false acceptance rate on FPDS experiments is a bit higher than on FEDS

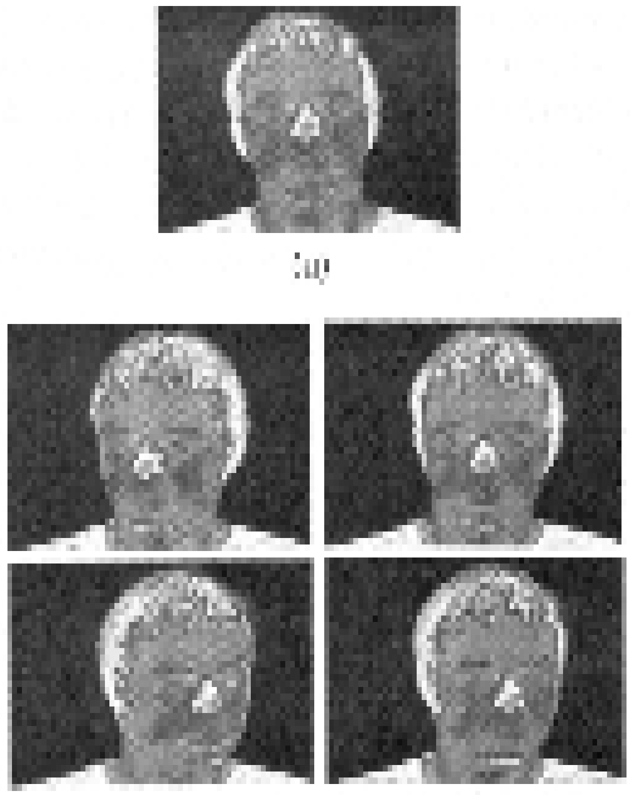

(a)

(b)

Figure 4.11 Sample subject from FEDS dataset (a) one gallery image at neutral pose and expression, (b) four probe images at varying poses.

experiments. This is to be expected, as variations in pose typically cause more nonlinear deformations in the vascular network than those caused by variations in facial expressions.

### 4.5.2 Experiments on the University of Notre Dame Database

A major challenge associated with thermal face recognition is recognition performance over time (Socolinsky and Selinger 2004b). Facial thermograms

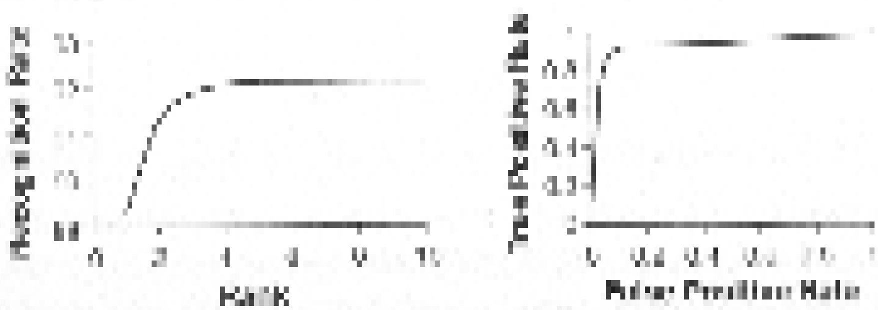


(a) (b)

Figure 4.12 Experimental results on the FEDS dataset using the dual bootstrap ICP matching algorithm: (a) CMC curve, (b) ROC curve.

may change depending on the physical condition of the subject, making it difficult to acquire similar features for the same person over time. Previous face recognition methods in the thermal infrared [illegible] reported degraded performance over time (Chen et al. 2003, 2005). Physiological information, however, remains invariant to physical [illegible] [illegible] the [illegible] vascular structure and [illegible] maintained [illegible] constant.

Since most of the subjects in the UA database had images collected during the same session, no statistically significant quantification of the time-lapse performance problem was possible. For this reason, the authors decided to apply the proposed approach on a data set, that of the University of Notre Dame (UND) [illegible] the University of Notre Dame. This database has a large collection of visible and thermal facial images acquired with a time gap. The database consists of [illegible] images acquired from [illegible] subjects during [illegible] sessions under specific lighting (LF: central light turned off, LM: all lights on) and facial expression conditions (FA: neutral expression, FB: smiling expression).

The database is divided into four different gallery and probe sets (Chen et al. 2003): LF–FA, LF–FB, LM–FA, and LM–FB. Each of the gallery sets (e.g., LF–FA) can be tested against the other three probe sets (e.g., LF–FB, LM–FA, and LM–FB). Thus, 12 different pairs of gallery and probe sets were used for testing. The performance of the new thermal face recognition algorithm was compared to that of TSCP matching (Buddharaju et al. 2006) and PCA matching (Chen et al. 2003) algorithms. Table 4.2 summarizes the rank 1 recognition results of these algorithms on each of the 12 experiments. Each entry in the left column of the table corresponds to the gallery set, and each entry in the top row corresponds to a probe set. From the table, it can be clearly seen that the new matching algorithm yields better recognition results even in the presence of time and temperature variations, thus outperforming the TSCP (physiology-based) and PCA (appearance-based) recognition algorithms.

### 4.5.3 Experiments on the University of Arizona Database

The authors have also used data from stress experiments carried out at the University of Arizona (UA). These data, although acquired within a few minutes for each subject, feature dramatic changes in the facial thermal map due to the onset of stress. In fact, in a few minutes, much more variability [illegible] than in the UND, which [illegible] to [illegible] the [illegible] experiment is approximately 20 minutes. For each subject, the authors [illegible] about five equally time-spaced frames from the interview. [illegible]

*Table 4.1 Rank 1 Recognition Performance of Dual Bootstrap Iterative Closest Point (DBICP) Vascular Network Matching Algorithm, TMP Matching Algorithm (Buddharaju et al. 2007), and PCA Algorithm (Chen et al. 2005) for Each of the 12 Experiments on the [illegible] Database*

| | Probe | | | |
|---|---|---|---|---|
| Gallery | [illegible] | [illegible] | [illegible] | [illegible] |
| [illegible] | – | [illegible] (DBICP)<br>[illegible] (TMP)<br>[illegible] (PCA) | [illegible] (DBICP)<br>[illegible] (TMP)<br>[illegible] (PCA) | [illegible] (DBICP)<br>[illegible] (TMP)<br>[illegible] (PCA) |
| [illegible] | [illegible] (DBICP)<br>[illegible] (TMP)<br>[illegible] (PCA) | – | [illegible] (DBICP)<br>[illegible] (TMP)<br>[illegible] (PCA) | [illegible] (DBICP)<br>[illegible] (TMP)<br>[illegible] (PCA) |
| [illegible] | [illegible] (DBICP)<br>[illegible] (TMP)<br>[illegible] (PCA) | [illegible] (DBICP)<br>[illegible] (TMP)<br>[illegible] (PCA) | – | [illegible] (DBICP)<br>[illegible] (TMP)<br>[illegible] (PCA) |
| [illegible] | [illegible] (DBICP)<br>[illegible] (TMP)<br>[illegible] (PCA) | [illegible] (DBICP)<br>[illegible] (TMP)<br>[illegible] (PCA) | [illegible] (DBICP)<br>[illegible] (TMP)<br>[illegible] (PCA) | – |

from the database is shown in Figure 4.13. It can clearly be seen [illegible] that the vascular map changed significantly between [illegible] (Figure 4.13[a]) and [illegible] (Figure 4.13[b]).

Figure 4.14(a) shows the CMC and Figure 4.14(b) the ROC curves of the [illegible] experiments using the dual bootstrap ICP versus the TMP matching algorithms. The results demonstrate that the dual bootstrap ICP outperforms the TMP matching algorithm (Buddharaju et al. 2007). In the case of TMP matching, the CMC curve shows that rank 1 recognition is [illegible], whereas the [illegible] is [illegible]. Also, the dual bootstrap ICP matching method achieves a high positive detection rate at very low false detection rates as is shown in Figure 4.14(b).

## 4.6 Conclusions

This chapter presents new algorithms that substantially improve the performance of physiology-based face recognition in the thermal infrared. Specifically, a vascular network postprocessing algorithm removes fake contours detected by the top-hat segmentation algorithm. A new vascular network matching algorithm can cope with nonlinear deformations between test and database vascular network [illegible].

The experimental results obtained with the [illegible] databases showcase [illegible] the superiority of the new method.

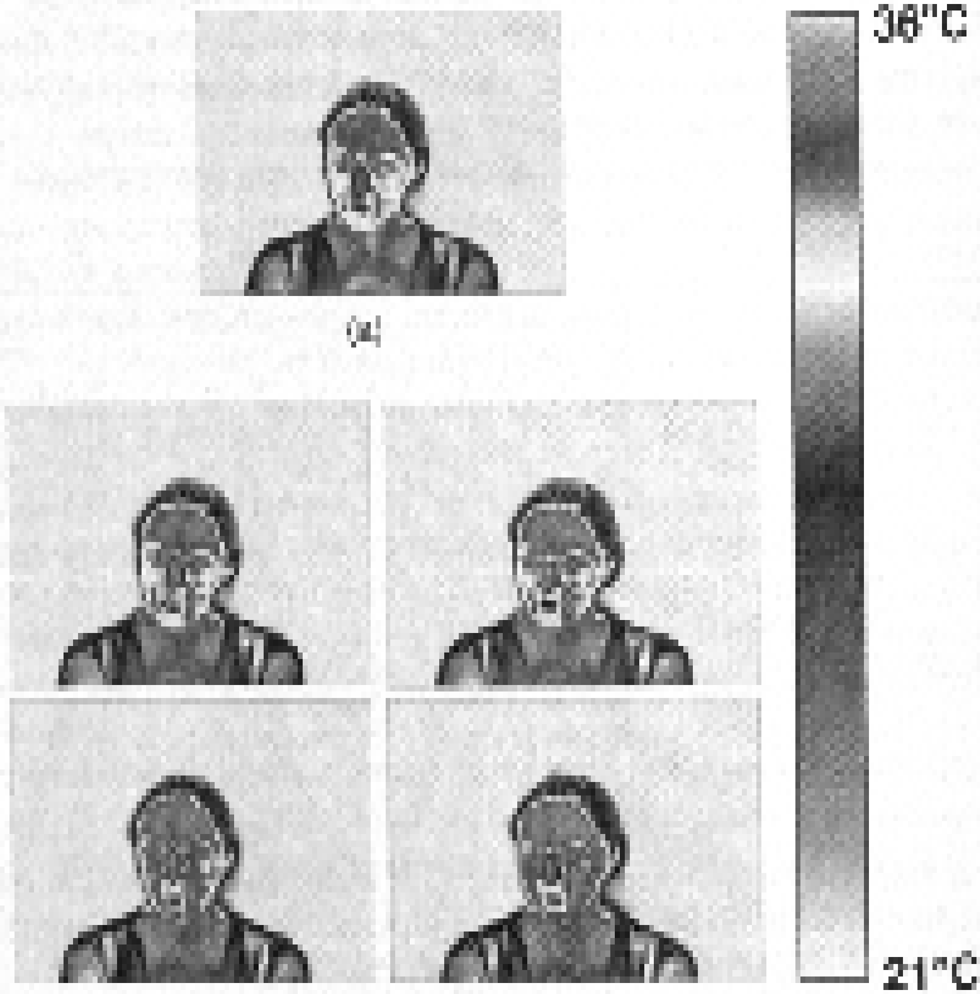


(a) (b) (c)

Figure 4.13. Sample subject from UA dataset: (a) one database image captured at 2 minutes 6 seconds of the interview; (b) four test images extracted at 6 minutes, 9 minutes 18 seconds, 13 minutes 28 seconds, and 17 minutes 40 seconds of the interview, respectively; (c) Thermal color map used for visualization.

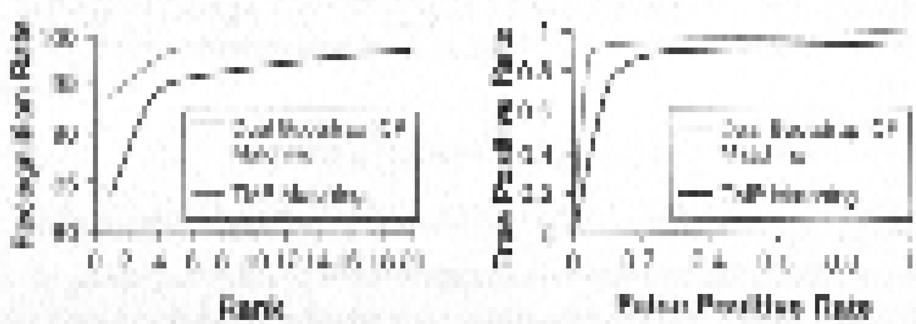


(a) (b)

Figure 4.14. Results of the experiments on the UA database using dual bayesian versus TMP matching: (a) CMC curves; (b) ROC curves.

## Acknowledgements

The authors would like to thank the University of Notre Dame for kindly providing its facial database. They also acknowledge Dr. Judee Burgoon at the University of Arizona for designing and conducting the stress experiments, as part of a collaborative DOD effort. This research was supported in part by a contract from the Defense Academy for Credibility Assessment (DACA). The views expressed in the chapter do not necessarily represent the views of the funding agencies.


## References


CVRL Biometrics at the University of Notre Dame. Biometrics database distribution. http://www.nd.edu/~cvrl/.

Besl, P., and N. McKay. 1992. A method for registration of 3-d shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2), 239–256.

Buddharaju, P., and I. Pavlidis. 2009. Face recognition beyond the visible spectrum. In N. Ratha and V. Govindaraju, editors, *Advances in Biometrics: Sensors, Algorithms and Systems*, chapter 9, pages 157–180. Springer.

Buddharaju, P., I. Pavlidis, and I. Kakadiaris. 2004. Face recognition in the thermal infrared spectrum. In *Proceedings of the Joint IEEE Workshop on Object Tracking and Classification Beyond the Visible Spectrum*, Washington, D.C.

Buddharaju, P., I. Pavlidis, and P. Tsiamyrtzis. 2005. Pose-invariant physiological face recognition in the thermal infrared spectrum. In *Proceedings of the 2006 IEEE Conference on Computer Vision and Pattern Recognition*, pages 53–60, New York.

Buddharaju, P., I. Pavlidis, P. Tsiamyrtzis, and M. Bazakos. 2007. Physiology-based face recognition in the thermal infrared spectrum. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(4), 613–626.

Chen, X., P. Flynn, and K. Bowyer. 2003. PCA-based face recognition in infrared imagery: baseline and comparative studies. In *Proceedings of the IEEE International Workshop on Analysis and Modeling of Faces and Gestures*, pages 127–134, Nice, France.

Chen, X., P. Flynn, and K. Bowyer. 2005. IR and visible light face recognition. *Computer Vision and Image Understanding*, 99(3), 332–358.

Huttenlocher, D. P., G. A. Klanderman, and W. A. Rucklidge. 1993. Comparing images using the Hausdorff distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(9), 850–863.

Kong, S. G., J. Heo, B. R. Abidi, J. Paik, and M. A. Abidi. 2005. Recent advances in visual and infrared face recognition – a review. *Computer Vision and Image Understanding*, 97(1), 103–135.

Selinger, A., and D. Socolinsky. 2004. Face recognition in the dark. In *Proceedings of the Joint IEEE Workshop on Object Tracking and Classification Beyond the Visible Spectrum*, Washington, D.C.

Socolinsky, D., and A. Selinger. [illegible]. Thermal face recognition in [illegible]. In Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, volume 2, pages [illegible], Washington, DC.

[illegible], and [illegible]. 200[illegible]. [illegible]. In Proceedings of the [illegible] International Conference on Pattern Recognition, volume 4, pages [illegible].

[illegible] recognition using [illegible] images. In Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, volume 1, pages [illegible], [illegible].

[illegible]. IEEE Transactions on Medical Imaging, 23(11), [illegible].

Wang, J., G., E. Sung, and R. Venkateswarlu. 2004. Registration of infrared and visible spectrum imagery for face recognition. In Proceedings of the Sixth IEEE International Conference on Automatic Face and Gesture Recognition, pages [illegible], Seoul, Korea.

Zhao, W., R. Chellappa, P. J. Phillips, and A. Rosenfeld. 2003. Face recognition: A literature survey. ACM Computing Surveys (CSUR), 35(4): 399–458.