

Face Recognition Beyond the Visible Spectrum

Pradeep Buddharaju, Ioannis Pavlidis, and Chinmay Manohar

Abstract. The facial vascular network is highly characteristic of the individual, much the way his fingerprint is. An unobtrusive way to capture this information is through thermal imaging. The convective heat transfer effect from the flow of “hot” arterial blood in superficial vessels creates characteristic thermal imprints, which are at a gradient with the surrounding tissue. This casts sigmoid edges on the human tissue where major blood vessels are present. We present an algorithmic methodology to extract and represent the facial vasculature. The methodology combines image morphology and probabilistic inference. The morphology captures the overall structure of the vascular network and the probabilistic part reflects the positional uncertainty for the vessel walls, due to the phenomenon of thermal diffusion. The accuracy of the methodology is tested through extensive experimentation and meticulous ground-truthing. Furthermore, the efficacy of this information for identity recognition is tested on substantial databases.

9.1 Introduction

Biometrics has received a lot of attention during the last few years both from the academic and business communities. It has emerged as a preferred alternative to traditional forms of identification, such as card IDs, which are not embedded into one’s physical characteristics. Research into several biometric modalities including face, fingerprint, iris, and retina recognition has produced varying degrees of success (Jain et al., 1999). Face recognition stands as the most appealing modality, in as much as 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). Research into face recognition has been biased towards the visible spectrum for a variety of reasons. Among those is the availability and low cost of visible band cameras and the undeniable fact that face recognition is one of the primary activities of the human visual system. Machine recognition of human faces, however,

has proven more problematic than the seemingly effortless face recognition performed by humans. The major culprit is light variability, which is prevalent in the visible spectrum due to the reflective nature of incident light in this band. Secondary problems are associated with the difficulty of detecting facial disguises (Pavlidis and Symosek, 2000).

As a solution to the aforementioned problems, researchers have started investigating the use of thermal infrared for face recognition purposes (Prokoski, 2000; Socolinsky and Selinger, 2002; Wilder et al., 1996). However, many of these research efforts in thermal face recognition use the thermal infrared band only as a way to see in the dark or reduce the deleterious effect of light variability (Socolinsky et al., 2001; Selinger and Socolinsky, 2004). Methodologically, they do not differ very much from face recognition algorithms in the visible band, which can be classified as appearance-based (Cutler, 1996; Chen et al., 2003) and feature-based approaches (Srivastava and Liu, 2003; Buddharaju et al., 2004). The only difference between the two modalities is that each pixel in a visible image contains intensity values whereas the infrared images have temperature values. Hence all the visible face recognition algorithms can be readily applied to infrared facial images.

In this chapter, we propose a novel approach to the problem of thermal facial recognition by extracting the superficial blood vessels on the face (see Figure 9.1). Our goal is to promote a different way of thinking in the area of face recognition in thermal infrared, which can be approached in a distinct manner when compared with other modalities. It consists of a statistical face segmentation and a physiological feature extraction algorithm

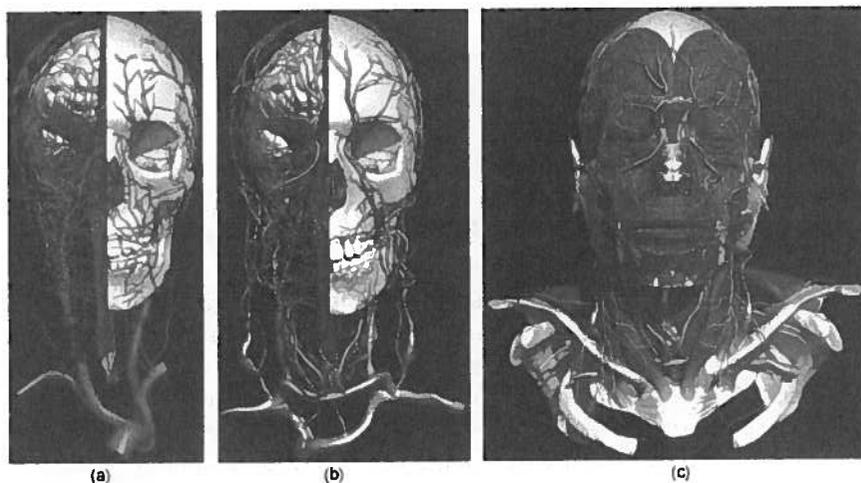


Fig. 9.1. Generic map of superficial blood vessels on the face. Courtesy of Primal Pictures (Moxham et al., 2002). (a) Overview of arterial network; (b) overview of venous network; (c) arteries and veins together underneath the surface of the facial skin.

tailored to thermal phenomenology. Prokoski et al. anticipated the possibility of extracting the vascular network from thermal facial images and using it as a feature space for face recognition (Prokoski and Riedel, 1998). However, they did not present an algorithmic approach for achieving this. We present a full abstraction methodology to extract the vascular network from infrared facial imagery (Buddharaju et al., 2005, 2006).

This is functional imaging at its best, as the vessel network is imaged through its function (blood flow). The biometric advantage is that this information is very difficult to be altered purposefully. Therefore, it has the potential to become a foolproof ID method for high-end security applications. This chapter also presents an in-depth study of uniqueness and repeatability characteristics of facial vasculature, which are very important for any physical feature to be a biometric technology.

9.2 Facial Vasculature Extraction

In thermal imagery of human tissue the major blood vessels have weak sigmoid edges. This is due to the natural phenomenon of heat diffusion, which entails that when two objects with different temperatures are in contact (e.g., vessel and surrounding tissue), heat conduction creates a smooth temperature gradient at the common boundary (Garbey et al., 2004). This phenomenon is strong in some major vessels such as the common carotid artery, and hence can be clearly seen just by visualizing temperature values around them as shown in Figure 9.2. The abstraction methodology to extract these edges is carried out in two stages. Firstly, we segment the facial tissue from the background, which ensures that any further processing is applied to the face alone. Then we segment all regions on the face that exhibit sigmoid edges, which give superficial blood vessels on the face.

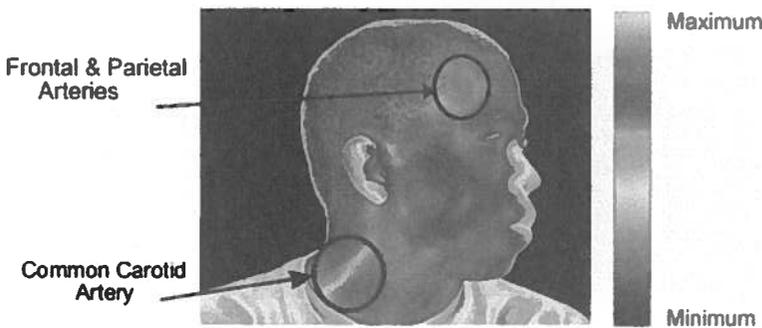


Fig. 9.2. Visualization of the temperature values of a thermal facial image. The smooth gradient around the common carotid artery can be clearly seen.

9.2.1 Face Segmentation

Due to 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 clothes being the “hot” object. The consistency of bimodality across subjects and image backgrounds is striking. We approach the problem of delineating facial tissue from the background using a Bayesian framework because we have a priori knowledge of the bimodal nature of the scene. Figure 9.3b shows the temperature distributions of the facial skin and the background from a typical infrared facial image. We approach the problem of delineating facial tissue from the background using a Bayesian framework (Buddharaju et al., 2005; Pavlidis et al., 2006) because we have a priori knowledge of the bimodal nature of the scene.

We call θ 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:

$$p^{(t)}(\theta|x_t) = \begin{cases} p^{(t)}(s|x_t), & \text{when } \theta = s, \\ p^{(t)}(b|x_t) = 1 - p^{(t)}(s|x_t), & \text{when } \theta = b. \end{cases} \quad (9.1)$$

We develop the statistics only for skin, and then the statistics for the background can easily be inferred from Equation (9.1).

According to Bayes’ theorem:

$$p^{(t)}(s|x_t) = \frac{\pi^{(t)}(s)f(x_t|s)}{\pi^{(t)}(s)f(x_t|s) + \pi^{(t)}(b)f(x_t|b)}. \quad (9.2)$$

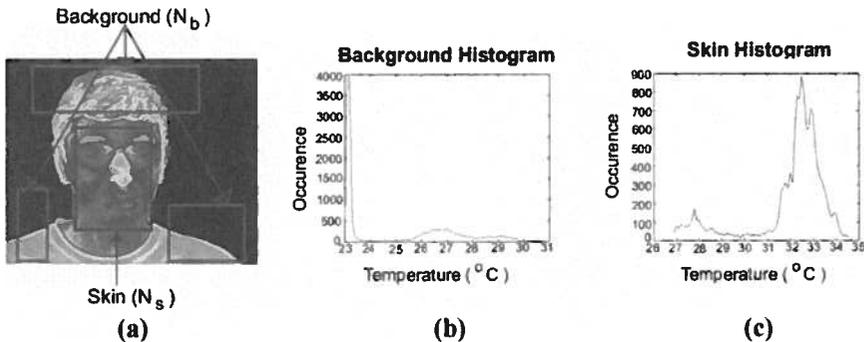


Fig. 9.3. Skin and background: (a) Selection of samples for EM algorithm; (b) corresponding bimodal temperature distribution of background region; (c) corresponding bimodal temperature distribution of skin region.

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:

$$\pi^{(1)}(s) = \frac{1}{2} = \pi^{(1)}(b). \quad (9.3)$$

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

$$\pi^{(t)}(s) = p^{(t-1)}(s|x_{t-1}). \quad (9.4)$$

The likelihood $f(x_t|s)$ of pixel x representing skin at time $t \geq 1$ is given by

$$f(x_t|s) = \sum_{i=1}^2 w_{s_i} N(\mu_{s_i}^{(t)}, \sigma_{s_i}^{2(t)}), \quad (9.5)$$

where the mixture parameters w_{s_i} (weight), μ_{s_i} (mean), $\sigma_{s_i}^2$ (variance): $i = 1, 2$ and $w_{s_2} = 1 - w_{s_1}$ of the bimodal skin distribution can be initialized and updated using the EM algorithm. For that, we select N representative facial frames (offline) from a variety of subjects that we call the training set. Then, we manually segment, for each of the N frames, skin (and background) areas, which yields N_s skin (and N_b background) pixels as shown in Figure 9.3a.

Figure 9.4b visualizes the result of our Bayesian segmentation scheme on the subject shown in Figure 9.4a. Part of the subject's nose has been erroneously classified as background and a couple of cloth patches from the subject's shirt have been erroneously marked as facial skin. This is due to occasional overlapping between portions of the skin and background distributions. The isolated nature of these mislabeled patches makes them easily correctable through postprocessing. We apply a three-step postprocessing algorithm on the binary segmented image. Using foreground (and background) correction, we find the mislabeled pixels in the foreground (and background) and reassign them. Figure 9.4c visualizes the result of postprocessing where all the segmentation imperfections have been eliminated.

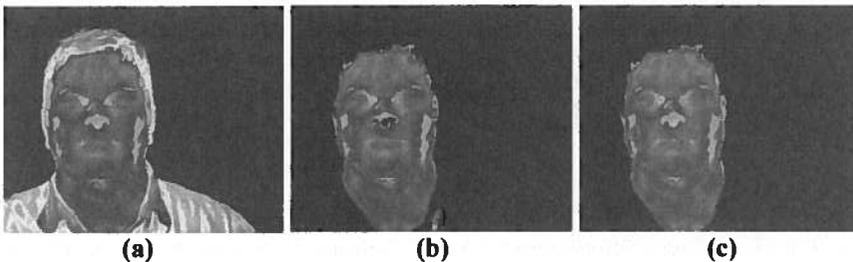


Fig. 9.4. Segmentation of facial skin region: (a) Original thermal facial image; (b) result of Bayesian segmentation where background is depicted in black; (c) result of postprocessing.

9.2.2 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 two steps (Pavlidis et al., 2006; Manohar, 2004).

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

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

The weak sigmoid edges formed due to heat diffusion at blood vessels can be handled effectively using anisotropic diffusion. The anisotropic diffusion filter is formulated as a process that enhances object boundaries by performing intraregion as opposed to interregion smoothing. One can visualize this clearer in an area with sparser vasculature than that of the face. Figure 9.5 shows vividly how the application of anisotropic diffusion on the thermal image of a wrist enhanced the sigmoid edges around the vessel and at the same time helped to remove noise formed due to hair.

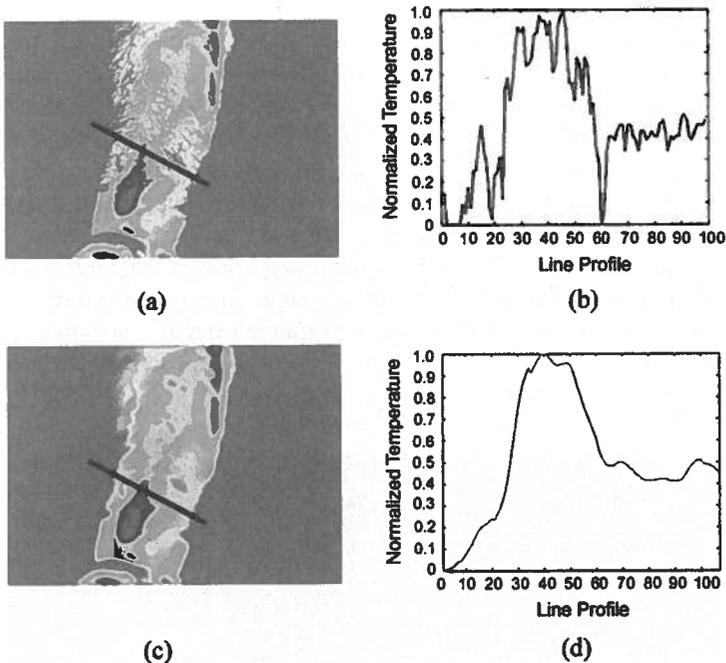


Fig. 9.5. Anisotropic diffusion on the thermal image of a human wrist: (a) segmented wrist image; (b) profile of the line drawn across segmented image (shown in black in (a)); (c) result of applying anisotropic diffusion on (a); (d) profile of the same line drawn across diffused image (shown in black in (b)).

The mathematical equation that describes this process is:

$$\frac{\partial I(\bar{x}, t)}{\partial t} = \nabla(c(\bar{x}, t)\nabla I(\bar{x}, t)). \quad (9.6)$$

In our case $I(\bar{x}, t)$ is the thermal infrared image, \bar{x} refers to the spatial dimensions, and t to time. $c(\bar{x}, t)$ is called the diffusion function. The discrete version of the anisotropic diffusion filter of Equation (9.6) is as follows.

$$\begin{aligned} I_{t+1}(x, y) = & I_t + \frac{1}{4} * [c_{N,t}(x, y)\nabla I_{N,t}(x, y) \\ & + c_{S,t}(x, y)\nabla I_{S,t}(x, y) + c_{E,t}(x, y)\nabla I_{E,t}(x, y) \\ & + c_{W,t}(x, y)\nabla I_{W,t}(x, y)]. \end{aligned} \quad (9.7)$$

The four diffusion coefficients and four gradients in Equation (9.7) correspond to four directions (i.e., 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,

$$c_{N,t}(x, y) = \exp\left(\frac{-\nabla I_{N,t}^2(x, y)}{k^2}\right), \quad (9.8)$$

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

Image morphology is then applied on the diffused image to extract the blood vessels that are at a relatively low contrast compared to that of the surrounding tissue. We employ for this purpose a top hat segmentation method, which is a combination of erosion and dilation operations. Top hat segmentation takes one of 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") ridgelike 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:

$$\begin{aligned} I_{open} &= (I \ominus S) \oplus S, \\ I_{top} &= I - I_{open}, \end{aligned} \quad (9.9)$$

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. Figure 9.6b depicts the result of applying anisotropic diffusion to the segmented facial tissue shown in Figure 9.6a and Figure 9.6c shows the corresponding vascular network extracted via white top hat segmentation.

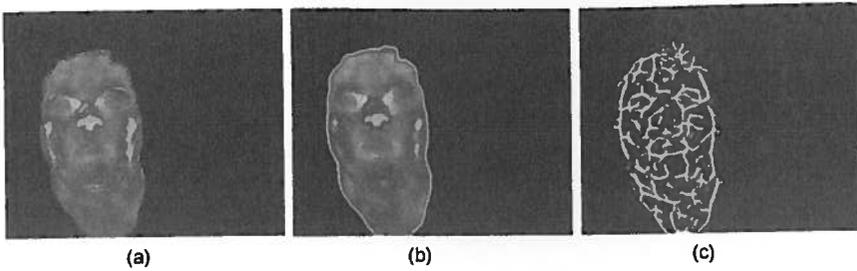


Fig. 9.6. Vascular network extraction: (a) original segmented image; (b) anisotropically diffused image; (c) blood vessels extracted using white top hat segmentation

9.3 Performance Metrics for Segmentation

To quantify the performance of our segmentation algorithms, we performed validation studies on the superficial vessels on the forearm. These vessels are more pronounced and hence they are easy to manually segment for ground truth purposes. Although the validation was performed at a different tissue area than the face, because of the project's objective, it gives a quantifiable indication of how accurately the method localizes superficial vasculature in general.

In the validation experiments subjects were sitting at a distance of 6 to 8 ft from the thermal imaging system. We captured 21 images from 15 subjects. Therefore, some subjects were imaged more than once in imaging sessions that were held days apart.

The ground truth was created by manual delineation at a variable super-resolution level (up to pixel by pixel) by medical experts. The ground-truth results from the two experts were reconciled and a composite ground-truth set was formed as a result. Figure 9.7 depicts some samples from the thermal image set used in the validation experiments along with the corresponding segmentation results. Figure 9.8 depicts the generic vascular map of the forearm to facilitate interpretation. In the forearm, vasculature is sparser and grander with respect to the face. This facilitates ground-truthing and visualization. Indeed, even a layman can identify in the images of Figure 9.7 the radial arterio-venous complex, which runs hotter (brighter) across the length of the arm and it is successfully segmented by our algorithm.

We perform quantitative analysis in terms of two measures: overlap ratio and Hausdorff distance. To avoid any confusion some definitions are in order. The building block that we need in the definitions of the measures is the notion of a confusion matrix. The confusion matrix tells us about the extent of the overlap between the segmented and ground-truthed images. Figure 9.9 shows the confusion matrix for a two-class classifier:

- TP is the number of correct predictions that a pixel belongs to a vessel (positive).

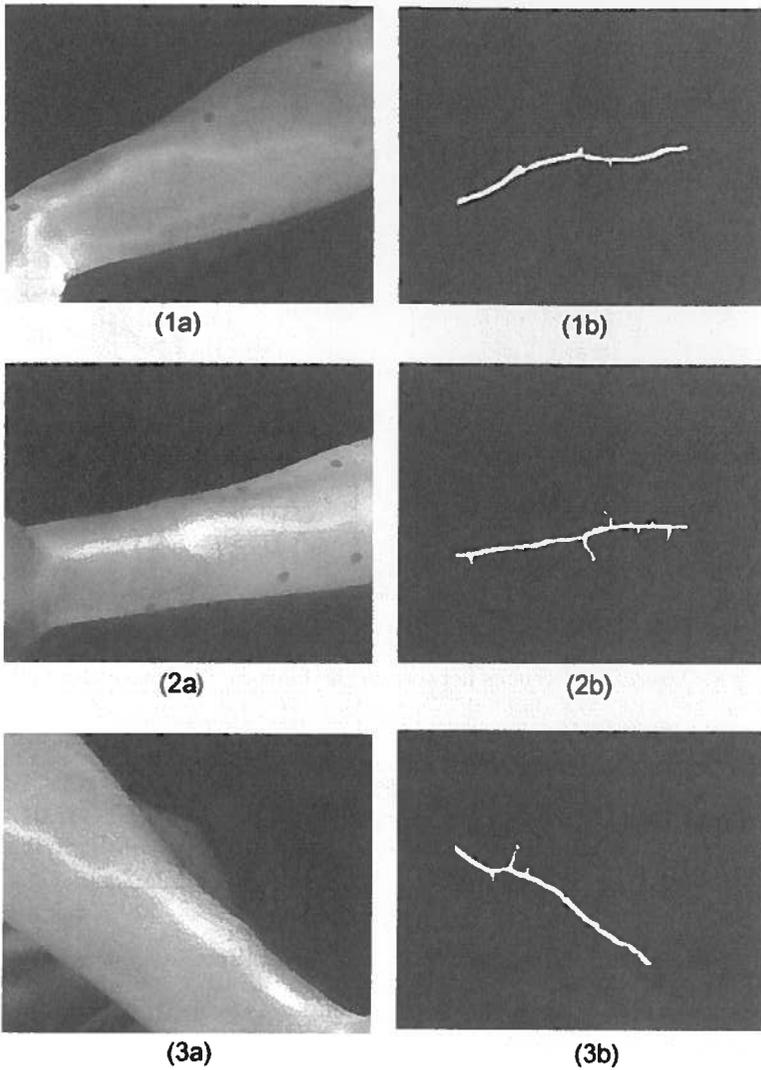


Fig. 9.7. (a) samples from the thermal imaging set used to validate the vasculature segmentation algorithm; (b) corresponding segmentation results.

- FP is the number of incorrect predictions that a pixel belongs to a vessel (positive).
- FN is the number of incorrect predictions that a pixel belongs to surrounding tissue (negative).
- TN is the number of correct predictions that a pixel belongs to surrounding tissue (negative).

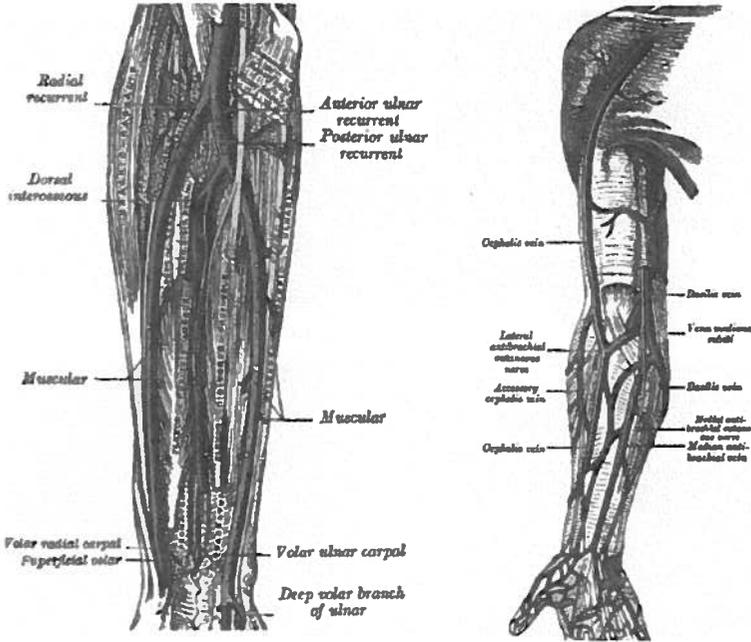


Fig. 9.8. Arterial and venous network in the forearm. (Courtesy Gray (1977).

		Predicted	
		Positives	Negatives
Actual	Positives	TP	FP
	Negatives	FN	TN

Fig. 9.9. Confusion matrix.

Based on the confusion matrix definition let us define the measures we use to quantify the overlap ratio.

Accuracy: Accuracy of a system is defined as the ratio of correctly classified pixels (true positives and true negatives) to the total number of pixels available at hand. It is usually expressed as a percentage. Accuracy is then calculated as

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \tag{9.1}$$

Specificity (or Precision): Specificity is the ability to correctly identify the background pixels. It is the ratio of the number of number of true negatives to the sum of true negatives and false positives, and it is given as

$$Specificity = \frac{TN}{TN + FP} \tag{9.2}$$

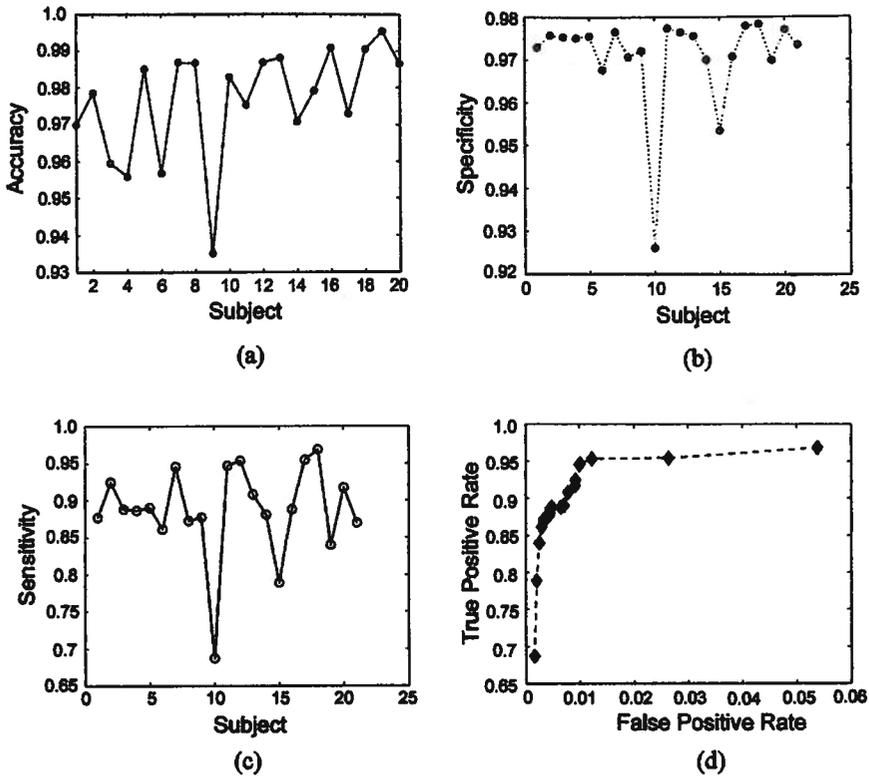


Fig. 9.10. Performance measures for overlap ratio: (a) Accuracy; (b) specificity; (c) sensitivity; (d) receiver operating characteristic (ROC) curve.

Sensitivity (or Recall): Sensitivity is the ability to correctly identify the segmented pixel. It is the ratio of the number of true positives to the sum of true positives and false negatives,

$$Sensitivity = \frac{TP}{TP + FN} \tag{9.12}$$

The plots for accuracy, precision, recall, and ROC are shown in Figure 9.10. One can observe that the method segments vascular thermal imprints with accuracy that is consistently above 95%. Precision and recall also feature very high values. All these performance measures indicate that there is very good overlapping between the segmented and expertly handdrawn vessels.

Quantifying the extent of overlap between segmented and expertly delineated vessel imprints is not enough to fully illuminate the qualities of the segmentation algorithm. For this reason we also used the Hausdorff distance, which is a measure of the closeness of two contours. Specifically, the Hausdorff distance is the maximum distance of a set to the nearest point in the other set.

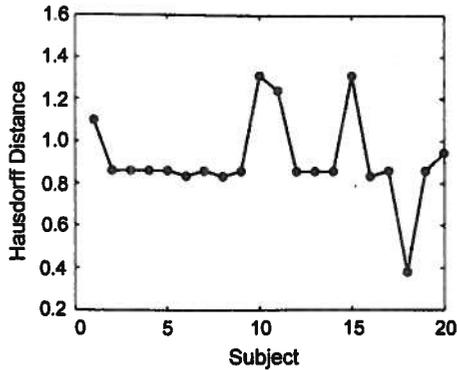


Fig. 9.11. Hausdorff distance plot between the automatically segmented and expertly delineated images.

More formally, Hausdorff distance from set A to set B is a max-min function, defined as

$$h(A, B) = \max(\min(d(A, B))) \quad (9.13)$$

where a and b are points of sets A and B , respectively, and $d(a, b)$ is any metric between these points; for simplicity, we take $d(a, b)$ as the Euclidean distance between a and b . The algorithm to find out the Hausdorff distance is:

1. $h = 0$
2. for every point a_i of A ,
 - 2.1 shortest = Inf ;
 - 2.2 for every point b_j of B
 - $d_{ij} = d(a_i, b_j)$
 - if $d_{ij} < \text{shortest}$ then
 - shortest = d_{ij}
 - 2.3 if shortest $> h$ then
 - $h = \text{shortest}$

Figure 9.11 shows the Hausdorff distance calculated between our manual segmented data and the original data. The low Hausdorff distance values indicate that the two sets in our case are close to each other. In other words, the automatically segmented vessel curves are very close to the expertly drawn ones.

9.4 Facial Vasculature for Biometric Identification

If a human physiological or behavioral characteristic has to be considered as a biometric feature, it should satisfy certain desirable characteristics such as universality, uniqueness, repeatability, collectability, performance, acceptability,

and circumvention (Jain et al., 1999). Every living and nonliving object at a finite temperature emits radiations, which are captured by infrared cameras. The temperature data can be universally extracted by applying Planck's equation on the radiations captured from the face, which on further analysis yields vascular structure. The main advantage of face recognition among other biometric technologies is that it is completely noncontact and allows for on-the-fly identification. Minimal or no cooperation is demanded from a person in order to extract his or her facial vasculature. Hence, this technology is easily collectable and is highly acceptable. Because the vascular network lies below the skin and is imaged through its function (blood flow), it is almost impossible to be forged making it very hard to circumvent. Buddhharaju et al. (2005,2006) showed that the performance of the biometric identification system based on facial vasculature is very promising. This leaves the following characteristics of facial vasculature to be addressed in order to be considered as a good biometric technology.

1. *Uniqueness*: Is it possible for two persons to have the same vascular structure on the face?
2. *Repeatability*: Is facial vasculature invariant with time?

Pankanti et al. (2002) studied intraclass and interclass variations among fingerprints probabilistically using the minutia points extracted from fingerprint ridges. Recently Zhu et al. (2006) developed a stochastic model to capture variability among fingerprint minutia datasets. Similar techniques can be applied to study the variability among facial vasculatures of different individuals. Minutia points can be extracted from branching points of vessels similar to the way fingerprint minutia points are extracted at the bifurcations and endings of fingerprint ridges as shown in Figure 9.12.

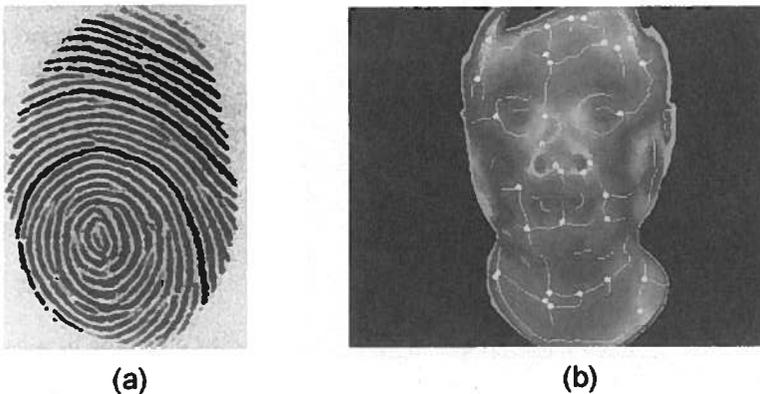


Fig. 9.12. (a) Fingerprint minutia points (Pankanti et al., 2002); (b) minutia points extracted from branches of facial vasculature.

9.4.1 Uniqueness

The pattern of the underlying blood vessels of the face (and the corresponding thermal imprints) is quite complex (see Figure 9.1). The question is whether this complex pattern is characteristic of each individual and can serve as a useful biometric signature.

In the area of medicine some very interesting work was conducted regarding the uniqueness of the facial vascular network. The primary motivation behind this line of research was the localization of anatomical features for reconstructive surgery purposes. For example, Pinar and Govsa (2006) conducted extensive research on the anatomy of the superficial temporal artery (STA) and its branches. They studied the STA anatomy in 27 subjects. Among other things they found that the bifurcation point of STA (see Figure 9.13) was above the zygomatic arch in only 20 out of the 27 samples. In 6 samples the bifurcation was exactly over the arch and in one sample there was no bifurcation at all. Further variability was observed in the STA branches. Specifically, in one sample double parietal branches were observed. In 21 samples zygomatico-orbital arteries ran towards the face, parallel to the zygomatic arch and distributed in the orbicularis oculi muscle. One has to take into account that STA is only one major facial vessel among many. Assuming that such variability is typical of other facial vessels and branches, their combination is bound to produce a very characteristic pattern for each individual.

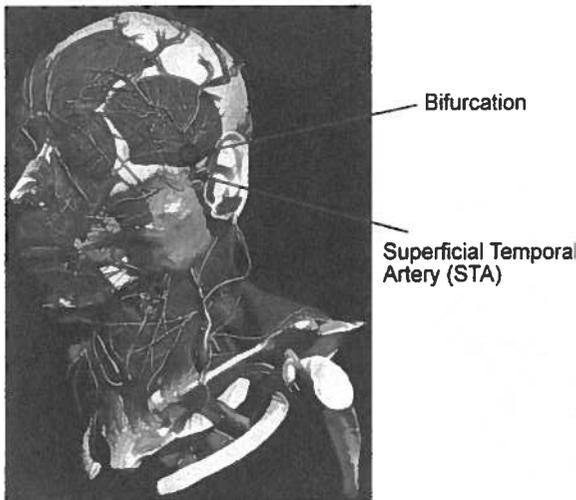


Fig. 9.13. Example of the superficial temporal artery (STA) and its bifurcation around the zygomatic arch. Courtesy of Primal Pictures (Moxham et al., 2002). Clinical studies have established its highly variable topology across individuals.

In another study, medical researchers found implicit evidence of uniqueness of the cutaneous vasculature in the high variability of reflex drives (Rowell, 1977).

In addition, one has to take into account that the proposed face recognition method does not depend only on the topology of the facial vascular network but also on the fat depositions and skin complexion. The reason is that imagery is formed by the thermal imprints of the vessels and not the vessels directly. Even if the vessel topology were absolutely the same across individuals, still the thermal imprints would differ due to variable absorption from different fat padding (skinny faces versus puffy faces) (De Geef et al., 2006) and variable heat conductance from different skin complexions (dark skin is less conductive).

In addition to the medical evidence that appears to be strong and the supporting heat transfer principles, “uniqueness” of the facial vascular network is also reinforced by the experimental investigation we presented in our previous efforts (Buddharaju et al., 2005, 2006). Such experimental investigations constitute the main “proof of uniqueness” in other biometric modalities (e.g., fingerprint recognition; Pankanti et al., 2002) and of course they gain more weight as the size of the databases increases. In the case of thermal facial vessel imprints, the size of the databases is still relatively small, yet statistically significant (several hundred samples). One particular example that makes a very strong case for “uniqueness” is the discovery of different thermal facial vessel imprints even in identical twins (Prokoski and Riedel, 1998).

In the last few years one relevant biometric that has gained acceptance is the venous structure at the back of the hand. It is imaged typically with active near-infrared light and the image is formed due to backscattering. The claim of “uniqueness” is based primarily on experimental evidence from database classification efforts. No substantial medical research was pursued on the uniqueness of the hand’s venous structure, as reconstructive hand surgery is not as prevalent as facial surgery. In addition, the venous network at the back of the hand is not nearly as complicated as the facial vessel network (see Figure 9.1). Yet, it is increasingly accepted as a legitimate biometric (Zhuang et al., 2005) and it is used in practice (Snowflake Technologies) based mainly on experimental evidence from database classification efforts. Hence, evidence from medical research and reasoning based on heat transfer principles suggest that the facial vessel network is characteristic of each individual.

9.4.2 Repeatability

As shown in Figure 9.12, minutia points can be extracted from the branches of blood vessel contours in ways similar to those used for fingerprint minutia extraction. Numerous methods have been proposed for matching fingerprint minutiae, most of which try to simulate the way forensic experts compare fingerprints (Maltoni et al., 2003). Similar techniques can be employed to match thermal minutia points of two subjects. We first reported the complete

matching algorithm and experimental results on the University of Houston database in Buddharaju et al. (2005, 2006), where the interested reader may find more details.

A major challenge associated with thermal face recognition is the recognition performance over time (Socolinsky and Selinger, 2004). Facial thermograms may change depending on the physical condition of the subject. This renders the task of acquiring similar features for the same person over time difficult. Previous face recognition methods in thermal infrared that use direct temperature data reported degraded performance over time (Chen et al., 2003, 2005). However, our method attempts to solve this problem by extracting facial physiological information to build its feature space. This information is not only characteristic of each person but also remains relatively invariant to physical conditions. Although the thermal facial maps of the same subject appear to shift, the vascular network is more resistant to change. In imaging terms, the contrast between the temperatures in the vascular pixels and the surrounding pixels is relatively invariant, albeit the absolute temperature values shift appreciably. This is a direct consequence of the thermoregulatory mechanism of the human body. Our morphological image processing simply capitalizes upon this phenomenon and extracts the invariant vascular contours from the variable facial thermal maps.

Due to the small number of subjects in the University of Houston database for whom we had images spread over several months, no statistically significant quantification of the low permanence problem was possible. For this reason, we obtained permission to apply the method on the database of the University of Notre Dame (Computer Vision Lab). This database has a large collection of facial images acquired from both visible and long-wave infrared cameras. They held acquisitions weekly and most of the subjects in the database participated multiple times.

In more detail, the database consists of 2294 images acquired from 63 subjects during nine different sessions under specific lighting and expression conditions. The spatial resolution of the images is 312×239 pixels (about half of that featured in the UH database). They used three lights during data collection, one located in the center approximately 8 ft in front of the subject, one located 4 ft to the right, and the other 4 ft to the left of the subject. The subjects were asked to provide two expressions during acquisition, "neutral" and "smiling". The database is divided into four different gallery and probe sets using the FERET-style naming convention (Phillips et al., 2000).

1. LF (central light turned off) + FA (neutral expression)
2. LF (central light turned off) + FB (smiling expression)
3. LM (all three lights on) + FA (neutral expression)
4. LM (all three lights on) + FB (smiling expression)

The database also contains an exclusive training set (different from the gallery and probe sets) with samples collected from several subjects, from which a face space can be constructed for the PCA recognition algorithm.

We did not use this training set because our algorithm is feature-based and hence does not require any explicit training. However, each of the gallery sets (say LF—FA) can be tested against the other three probe sets (say LF—FB, LM—FA, and LM—FB). This way we tested our algorithm on 12 different pairs of gallery and probe sets. In each of these experiments, the gallery set had one image per subject, and the probe set had several disjoint images per subject depending on how many different acquisition sessions the subject attended. Figure 9.14 shows a sample of the gallery and probe images of a subject from the University of Notre Dame database.

The authors of the University of Notre Dame (UND) database compared the performance of face recognition in visible and IR modalities from both same-session and time-gap datasets (Chen et al., 2003, 2005). They used a PCA-based face recognition algorithm for these studies. They found that both visible and IR modalities performed well on same-session experiments, and that none of them was significantly better than the other. However, in time-lapse experiments they found that the PCA-based recognition using IR images had poor performance. This is an expected outcome because PCA is an appearance-based face recognition algorithm that directly uses temperature values to project the query image onto face space. The thermal facial map may be different between gallery and probe images depending on the ambient and physical conditions, which may cause the PCA algorithm to fail.

We compared the performance of our method with a PCA-based recognition algorithm to test the robustness of features extracted from the facial vascular network. Table 9.1 summarizes the rank 1 recognition results using our algorithm versus the PCA algorithm on each of the 12 possible experiments. Each entry of the left column in the table corresponds to a gallery set, and each entry in the top row corresponds to a probe set. From the results, it can be clearly seen that our method yields better recognition results despite the presence of time and temperature variations inherent in this database. This is clear indication that by abstracting away the thermal facial map to a physiological feature vector the low permanence problem can be addressed more adequately.

9.5 Operational Limitations

Major operational limitations to the vascular feature extraction method from the face fall into following categories.

1. Glasses are opaque in the infrared domain and hence block important vascular information around the eyes. Also, too much facial hair usually blocks the radiations emitted from the surface of the skin, and hence causes the part of the face with (bulky) hair to get segmented out. Figure 9.15 shows examples of face segmentation where parts of the face containing glasses and hair are segmented out.

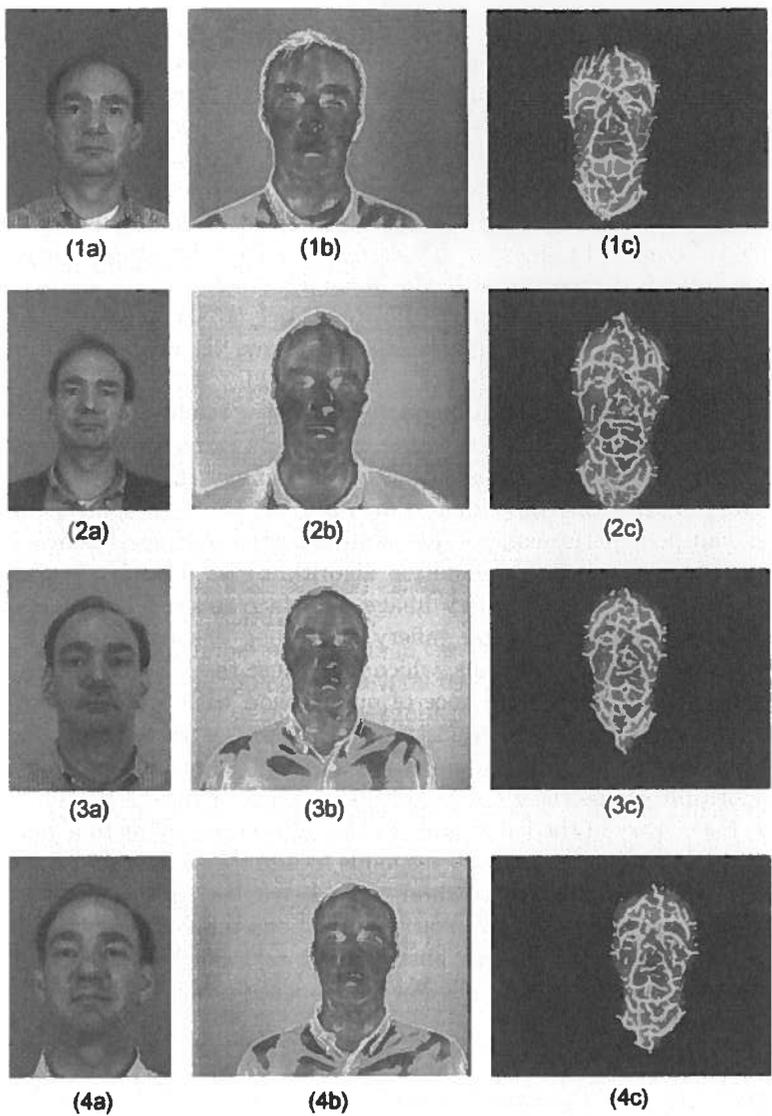
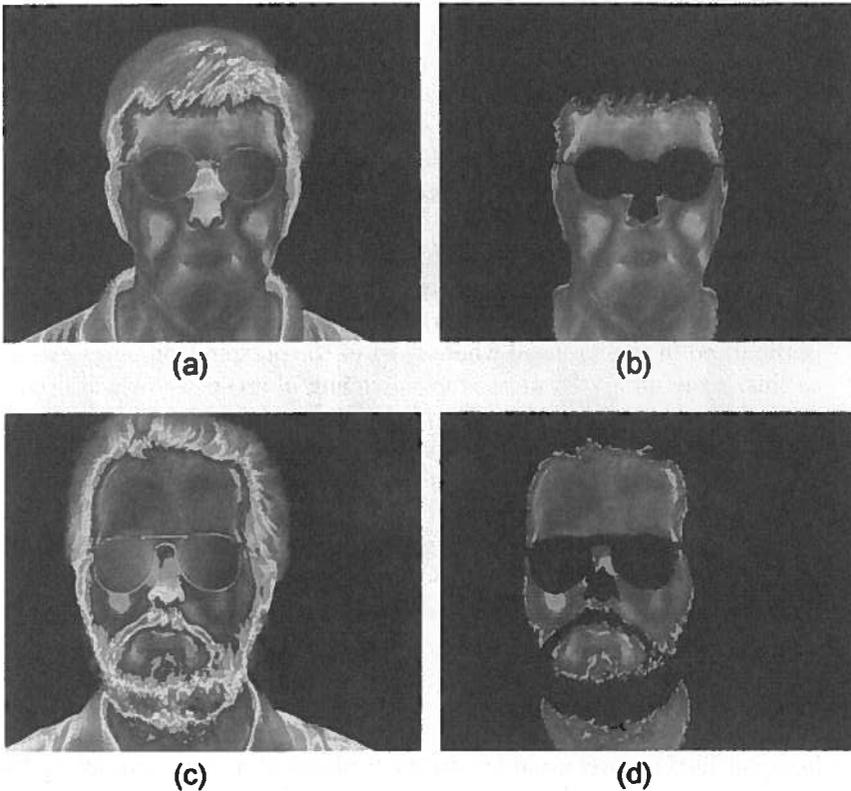


Fig. 9.14. Sample images of a subject in the University of Notre Dame database. The images were acquired over the span of several months: (a) visible images (not used here); (b) corresponding thermal infrared images; differences in the thermal facial maps can be visually appreciated; (c) vascular annotation after the application of our feature extraction algorithm.

Table 9.1. Rank 1 recognition performance of our algorithm (UH) versus the PCA algorithm on each of the 12 experiments on the UND database

Gallery	Probe			
	FA—LF(%)	FA—LM (%)	FB—LF (%)	FB—LM (%)
FA—LF	—	82.65 (UH)	80.77 (UH)	81.33 (UH)
	—	78.74 (PCA)	76.83 (PCA)	75.77 (PCA)
FA—LM	81.46 (UH)	—	79.38 (UH)	80.25 (UH)
	79.23 (PCA)	—	75.22 (PCA)	73.56 (PCA)
FB—LF	80.27 (UH)	81.92 (UH)	—	80.56 (UH)
	74.88 (PCA)	76.57 (PCA)	—	74.23 (PCA)
FB—LM	80.67 (UH)	82.25 (UH)	79.46 (UH)	—
	69.56 (PCA)	74.58 (PCA)	78.33 (PCA)	—

**Fig. 9.15.** (a) Thermal facial image with glasses and (b) result of segmentation. (c) Thermal facial image with facial hair and glasses and (d) result of segmentation.

2. The robustness of the method degrades when there is substantial perspiration. This results in a highly nonlinear shift of the thermal map that radically alters the radiation profile of the face. For the moment, this should be considered as the operational limit of the method. A practical scenario where such a case may arise is when a subject is imaged after strenuous exercise that lasted several minutes. Another possible breakdown may arise when the subject remains in a very hot environment, heavily dressed, for a substantial amount of time.

We have performed an experiment whereby a subject is imaged at the following instances.

- In a baseline condition (Figure 9.16, image 1a)
- After 1 min of rigorous walking (Figure 9.16, image 2a)
- After 5 min of rigorous walking (Figure 9.16, image 3a)
- After 5 min of rigorous jogging (Figure 9.16, image 4a)

Column b of Figure 9.16 shows the corresponding vessel extraction results. In the case of image 2a, the metabolic rate of the subject shifted to higher gear, but perspiration is still not a major problem. One can find evidence of the higher metabolic rate by looking at the left temporal area, where the region around the rich vasculature has become deeper cyan (hotter) in image 2a with respect to image 1a. This is an example of a positive linear shift (warming up), which the vessel extraction algorithm handles quite well (see image 2b versus image 1b). As the exercise become more strenuous and lasts longer, perspiration increases and introduces a negative nonlinear shift (cooling down) in the thermal map. This is especially pronounced in the forehead where most of the perspiration pores are. Due to this, some unwanted noise starts creeping in image 3b, which becomes more dramatic in image 4b. The performance of the vessel extraction algorithm deteriorates but not uniformly. For example, the vessel extraction algorithm continues to perform quite well in the cheeks where perspiration pores are sparse and the cooling down effect is not heavily nonlinear. In contrast, performance is a lot worse in the forehead area, where some spurious vessel contours are introduced due to severe nonlinearity in the thermal map shift.

9.6 Conclusions

We have outlined a novel approach to the problem of face recognition in thermal infrared. The cornerstone of the approach is the use of characteristic and time-invariant physiological information to construct the feature space. We presented a two-stage segmentation algorithm to extract superficial vasculature from the thermal facial image. The facial tissue is first separated from the background using a Bayesian segmentation method. The vascular network on the surface of the skin is then extracted based on a white

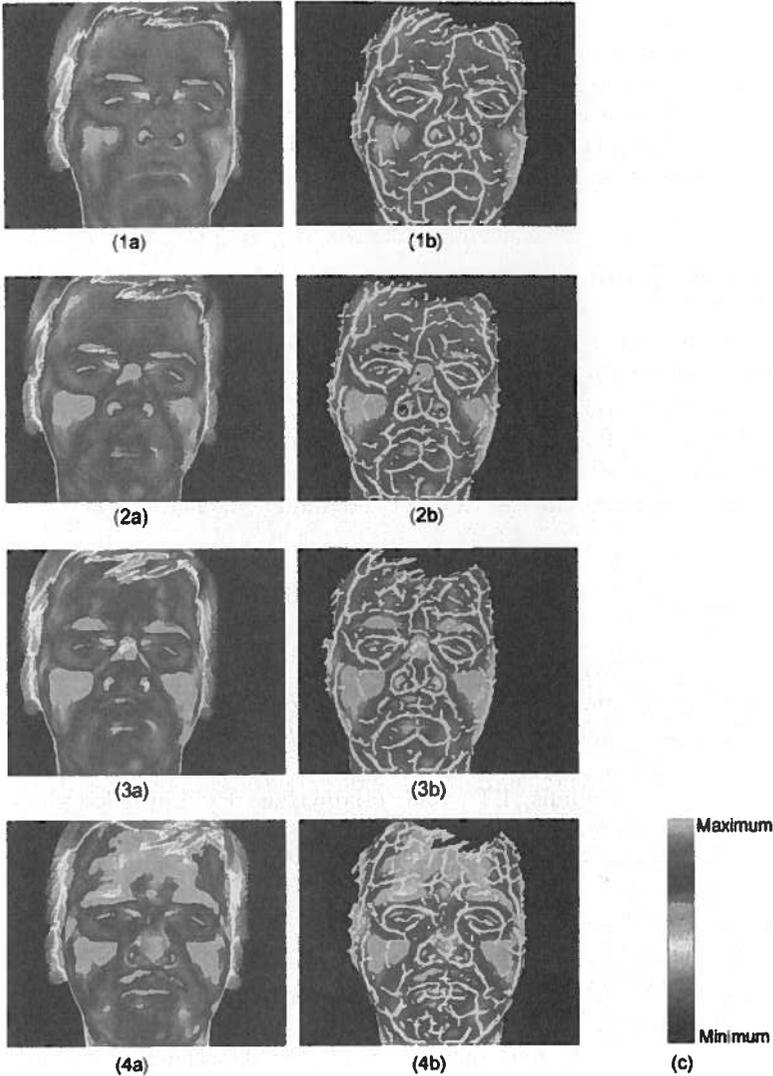


Fig. 9.16. Effect of perspiration on feature extraction. Thermal facial image of a subject (1a) at rest; (2a) after 1 minute of rigorous walking; (3a) after 5 min of rigorous walking; (4a) after 5 min of rigorous jogging, and (1b,2b,3b,4b) corresponding vascular network maps; and (c) color map used to visualize temperature values.

top-hat segmentation preceded by anisotropic diffusion. The good performance measures confirm the validity of segmentation algorithms.

The most important conclusion of our research so far, is that physiology-based face recognition appears to be feasible and have potential, especially as a way of addressing the issue of uniqueness and low permanence over time.

Although thermal facial maps shift over time, the contrast between superficial vasculature and surrounding tissue remains invariant. This physiological feature has permanence and is very difficult to alter (under the skin). Therefore, it gives a potential advantage to any face recognition method that may use it. It is an indication that the method is aided by the natural uniqueness and constancy of the feature space.

Acknowledgments

We would like to thank the University of Notre Dame for kindly providing us IR images acquired during different sessions. This research was supported mainly by NSF grant # DUE-0313880 in “Collaborative Research: Capacity Expansion in Information Assurance” and NSF grant # IIS-0414754 in “Interacting with Human Physiology.” The views expressed in this chapter do not necessarily represent the views of the funding agencies.

References

- Buddharaju, P., Pavlidis, I., and Kakadiaris, I. 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., USA, June 2004.
- Buddharaju, P., Pavlidis, I.T., and Tsiamyrtzis, P. Physiology-based face recognition. In *Proceedings of the IEEE Conference on Advanced Video and Signal Based Surveillance*, pages 354–359, Lake Como, Italy, September 15–16 2005.
- Buddharaju, P., Pavlidis, I.T., and Tsiamyrtzis, P. 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, June 17 2006.
- Chen, X., Flynn, P.J., and Bowyer, K.W. 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, October 17 2003.
- Chen, X., Flynn, P., and Bowyer, K. Ir and visible light face recognition. *Computer Vision and Image Understanding*, 99(3):332–358, September 2005.
- The Computer Vision Lab at the University of Notre Dame. Biometrics database distribution. <http://www.nd.edu/cvrl/>.
- Cutler, R. Face recognition using infrared images and eigenfaces. cs.umd.edu/rgc/face/face.htm, 1996.

- De Geef, S., Claes, P., Vandermeulen, D., Mollemans, W., and Willems, P.G. Large-scale in-vivo caucasian facial soft tissue thickness database for craniofacial reconstruction. *Forensic Science*, 159(1):S126–S146, May 2006.
- Garbey, M., Merla, A., and Pavlidis, I. Estimation of blood flow speed and vessel location from thermal video. In *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 1, pages 356–363, Washington D.C., USA, June 2004.
- Gray, H. *Gray's Anatomy: The Classic Collector's Edition*. Bounty Books, New York, 1977.
- Jain, A., Bolle, R., and Pankanti, S. *Biometrics: Personal Identification in Networked Society*. Kluwer Academic, 1999.
- Maltoni, D., Maio, D., Jain, A.K., and Prabhakar, S. *Handbook of Fingerprint Recognition*. Springer Verlag, June 2003.
- Manohar, C. Extraction of superficial vasculature in thermal imaging. Master's thesis, Department of Electrical Engineering, University of Houston, Houston, Texas, December 2004.
- Moxham, B. J., Kirsh, C., Berkovitz, B., Alusi, G., and Cheeseman, T. *Interactive Head and Neck (CD-ROM)*. Primal Pictures, December 2002.
- Pankanti, S., Prabhakar, S., and Jain, A.K. On the individuality of fingerprints. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(8):1010–1025, August 2002.
- Pavlidis, I. and Symosek, P. The imaging issue in an automatic face/disguise detection system. In *Proceedings of IEEE Workshop on Computer Vision Beyond the Visible Spectrum: Methods and Applications*, pages 15–24, Hilton Head Island, South Carolina, USA, June 2000.
- Pavlidis, I., Tsiamyrtzis, P., Manohar, C., and Buddharaju, P. *Biomedical Engineering Handbook*, Chapter Biometrics: Face recognition in thermal infrared. CRC Press, February 2006.
- Phillips, P.J., Moon, H., Rizvi, S.A., and Rauss, P.J. The feret evaluation methodology for face-recognition algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(10):1090–1104, 2000.
- Pinar, Y.A., and Govsa, F. Anatomy of the superficial temporal artery and its branches: its importance for surgery. *Surgical and Radiologic Anatomy (SRA)*, 28(3):248–253, June 2006.
- Prokoski, F. History, current status, and future of infrared identification. In *Proceedings of IEEE Workshop on Computer Vision Beyond the Visible Spectrum: Methods and Applications*, pages 5–14, Hilton Head Island, South Carolina, USA, June 2000.
- Prokoski, F.J. and Riedel, R. Infrared identification of faces and body parts. In A.K. Jain, R. Bolle, and S. Pankati, editors, *BIOMETRICS: Personal Identification in Networked Society*, Chapter 9. Kluwer Academic, 1998.
- Rowell, L.B. Reflex control of cutaneous vasculature. *Journal of Investigative Dermatology*, 69(1):154–166, July 1977.

- Selinger, A. and Socolinsky, D.A. Face recognition in the dark. In *Proceedings of the Joint IEEE Workshop on Object Tracking and Classification Beyond the Visible Spectrum*, Washington D.C., USA, June 2004.
- Snowflake Technologies. www.luminetx.com.
- Socolinsky, D.A. and Selinger, A. Thermal face recognition over time. In *Proceedings of the 17th International Conference on Pattern Recognition*, volume 4, pages 23–26, August 2004.
- Socolinsky, D.A. and Selinger, A. A comparative analysis of face recognition performance with visible and thermal infrared imagery. In *Proceedings of 16th International Conference on Pattern Recognition*, volume 4, pages 217–222, Quebec, Canada, 2002.
- Socolinsky, D.A., Wolff, L.B., Neuheisel, J.D., and Eveland, C.K. Illumination invariant face recognition using thermal infrared imagery. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2001)*, volume 1, pages 527–534, Kauai, Hawaii, USA, 2001.
- Srivastava, A. and Liu, X. Statistical hypothesis pruning for recognizing faces from infrared images. *Journal of Image and Vision Computing*, 21 (7):651–661, 2003.
- Wilder, J., Phillips, P.J., Jiang, C., and Wiener, S. Comparison of visible and infrared imagery for face recognition. In *Proceedings of the Second International Conference on Automatic Face and Gesture Recognition*, pages 182–187, Killington, Vermont, October 1996.
- Zhao, W., Chellappa, R., Phillips, P.J., and Rosenfeld, A. Face recognition: A literature survey. *ACM Computing Surveys (CSUR)*, 35(4):399–458, December 2003.
- Zhu, Y., Dass, S.C., and Jain, A.K. Compound stochastic models for fingerprint individuality. In *Proc. of International Conference on Pattern Recognition (ICPR)*, volume 3, pages 532–535, Hong Kong, August 2006.
- Zhuang, D., Ding, Y., and Wang, K. A study of hand vein recognition method. In *Proceedings of the 2005 IEEE International Conference on Mechatronics and Automation*, volume 4, pages 2106–2110, July 2005.

Part II

Algorithms