

The Segmentation of the Supraorbital Vessels in Thermal Imagery

Zhen Zhu
The Department of Computer Science
University of Houston
zhu@cs.uh.edu

Panagiotis Tsiamyrtzis
Department of Statistics
Athens University of Economics and Business, Greece
pt@aueb.gr

Ioannis Pavlidis
The Department of Computer Science
University of Houston
ipavlidis@uh.edu

Abstract

Thermal imaging techniques have been applied to detect and measure mental stress in polygraph screening and other applications. Mental stress is highly correlated with the activation of the corrugator muscle on the forehead. The vessels that supply blood to the corrugator muscle, proportionally to its degree of activation, are the supraorbital vessels. The rate of blood flow in these vessels can be indirectly measured via the intensity of heat emission from their segments. However, segmenting the thermal imprints of the supraorbital vessels is challenging because (1) they are fuzzy due to thermal diffusion, and (2) exhibit significant inter-individual and intra-individual variation. In this paper, a new segmentation method is proposed to extract the supraorbital vessels in thermal imagery. The new method features three steps: (1) automatic initialization of vessels; (2) automatic localization of the central lines of vessels; and (3) fast determination of vessel boundaries. The results show that the new method achieves high quality segmentation in both a simulated and a real dataset. The proposed method is expected to further increase the accuracy of stress measurements via thermal imaging.

1 Introduction

Thermal video has been widely used in surveillance and screening for various purposes. Recently, a series of thermal imaging methods have been developed to measure human physiological signs [5, 13, 16, 4], including pulse, blood perfusion, and breath. Stress measurement is another important application of thermal imaging. Previous work has shown that mental stress can be quantified through a thermal imaging system called “Stress-cam” [14]. There are

potentially many applications for this passive imaging technology; for example, a “Stress-cam” in lie detection could help identify interviewees’ critical physiological facial variations. In another example, the monitoring of computer users’ affective states could be used in software usability studies. A “Stress-cam” can also be used to monitor job-related stress.

Researchers reported that the activation of the corrugator muscle on the forehead is highly related to mental stress [9, 2]. The vessels that supply with blood the corrugator muscle are the supraorbital vessels. The activation of the corrugator muscle draws more blood. In this case and due to the heat diffusion effect, the thermal imprints of the supraorbital vessels are intensified in the thermal imagery. Therefore, the intensity of the imprints could be used to measure the activation of the corrugator muscle, which can be correlated with the level of mental stress.

A methodological problem that is still open is the segmentation of the thermal imprints of the supraorbital vessels. In [14], Puri *et al.* reported a heuristic segmentation method that certainly introduces non-trivial noise in the measurement process. Segmentation of the supraorbital imprints is challenging because:

1. The clarity of the supraorbital vessel imprints is low on average. Many factors lead to this. For example, the heat diffusion effect leaves vague boundaries between vessel and non-vessel regions; subcutaneous tissue may block thermal emissions and curtail part of the vessels’ thermal imprints.
2. The pattern of the vessel imprints on the forehead varies significantly both inter-individually and intra-individually. The overall thermal signature of the forehead is naturally unique to each individual (inter-individual variability). Furthermore, intra-individual

variability is largely related to physiological status, e.g., the vessel imprints expand when a person is under stress.

A number of approaches for segmenting blood vessels in various imaging modalities have been reported in the literature [8, 11]. These approaches can be divided into two categories, *bottom-up* and *top-down*. The bottom-up approach is based on criteria that are used to classify pixels in terms of salient image features; e.g., edges, intensity, and texture. The problem of this approach is that the vessels could be segmented into many small regions and some non-vessel areas could be erroneously labeled as vessels. The top-down approach, on the other hand, uses prior knowledge about vessels to guide the segmentation. The top-down method can find the overall vessel structure but fails to accurately delineate details, e.g., boundaries. Considering these factors, we developed a new segmentation method that combines merits of the bottom-up and the top-down approaches. This idea has its underpinnings in [1].

The core component of the proposed method is an active contour model (snake), which is a top-down approach. A traditional snake requires user interaction to initialize the contours. To overcome this problem, we introduce an automatic initialization method for the active contour model. In addition to that, we apply a bottom-up approach to find the boundaries. By combining top-down and bottom-up approaches, we not only localize the supraorbital vessels in the forehead region without being distracted by artifacts, but also maintain a detail finding capability, such as finding a vessel's boundaries. Specifically, the new method, called Open End Snake (OES), features three steps: (1) automatically initializes vessel location; (2) automatically localizes the central lines of a vessel; and (3) determines the vessel boundaries via a specialized operator. Experiments have demonstrated that the method has solid performance in segmenting forehead vessels, even in the presence of image noise. In fact, it achieves higher quality segmentation with respect to other methods.

The remainder of the paper is structured as follows: In Section 2, we discuss the underlying physiological mechanism. In Section 3, we present the details of the segmentation methodology. In Section 4, we validate the segmentation method and compare it with other methods to prove the merits. Finally, we discuss the open issues and conclude the paper in Section 5.

2 Forehead Physiology and Stress Level

Every human being undergoes certain physiological changes under stress. The facial region is heavily innervated with neuronal pathways and it is not surprising that stress affects its physiology. Previous work has demonstrated that

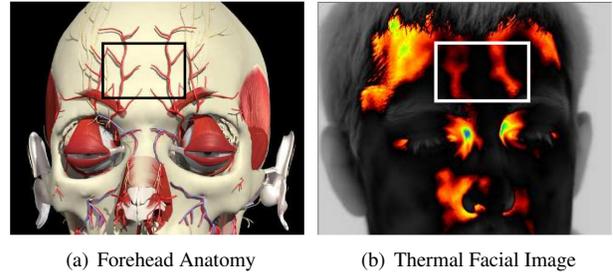


Figure 1. The anatomical and thermal images for the facial region. (a) is a courtesy of Primal Pictures [12].

both the periorbital [10] and supraorbital [14] regions manifest stress signs.

One of the most salient manifestations of sustained stress is that the eyebrows frown more frequently [3]. The frowning is caused by the contraction of the corrugator muscles. Activation of the corrugator muscles requires more blood, which is drawn from the supraorbital vessels. The supraorbital vessels are found in the middle of the forehead and consist of two major branches and several subbranches (see Figure 1(a)). Increased blood flow in supraorbital vessels, directly increases the cutaneous temperature on the forehead. In Figure 1(b), the imprints of the supraorbital vessels are apparent within the rectangular annotation of the thermal image. Our task is to delineate such supraorbital vessels in thermal imagery. Thereafter, accurate measurement of these vessels' temperature becomes feasible, which in turn provides a strong clue about the stress level of subjects (e.g., in polygraph screening).

3 Supraorbital Vessel Segmentation

3.1 Overview

The OES algorithm proceeds in three steps:

1. It detects the approximate location of the vessel imprints via the Hough transform. Examples of initial vessel locations are given by the straight lines shown in Figure 2(c).
2. It finds the exact central lines of the vessels via an active contour method (see Figure 2(d)).
3. It determines the boundaries of the vessels via a special operator (see Figure 2(d)).

These three steps ensure full vessel localization. In the following subsections, we will present more details for each of these steps. For simplicity, we interchangeably use "thermal imprints of the supraorbital vessels" and "vessels."

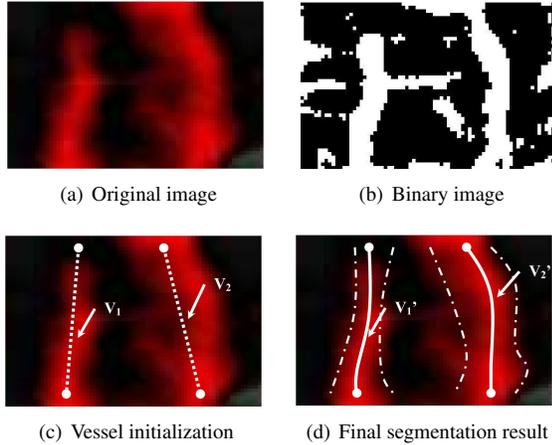


Figure 2. Example of vessel segmentation.

3.2 Step I: Vessel Initialization

The first step is to approximate the location of the vessels. This initialization process consists of two steps: producing a binary vessel image via the Tophat method and finding initial vessel lines via the Hough transform.

Tophat is a morphological method that enhances the bright (“hot”) ridge-like structures corresponding to the blood vessels. It aids the identification of vessel structures in a relatively low contrast image. The structuring element we use in the Tophat operation is a disk with diameter roughly equal to the width of the vessels. One example binary image produced by the Tophat method is seen in Figure 2(b).

The Hough transform is chosen to detect vessel lines in images. We apply the Hough transform on the binary image resulted from the Tophat operation, to find all major vessel lines. The line detection process has been tailored to the specific case. Since all of the supraorbital vessels are oriented vertically or nearly so, we only search the corresponding domain inside the transformed space. An example result of the Hough transform is illustrated in Figure 2(c).

3.3 Step II: Central Lines of Vessels

Having initialized the vessel locations, we want to further localize the central lines of the vessels. The actual vessel is shaped as a curved instead of a straight line, which was the outcome of the Hough Transform. A top-down approach, like the active contour, has the ability to remedy this and accurately capture the vessel’s shape. The active contour converges on the vessel’s central line, along which we have locally maximum temperature values. The proposed active contour model features two open ends instead of a closed loop in the traditional model [7]. Also, note that

there could be more than one active contours in the region of interest (ROI) and each of them evolves independently of the others.

3.3.1 The Active Contour Model

A traditional active contour [7] is a curve $X(s) = [x(s), y(s)]$, $s \in [0, 1]$, that minimizes the energy functional

$$E = \int_0^1 \frac{1}{2} [\alpha |X'(s)|^2 + \beta |X''(s)|^2] + E_{ext}(X(s)) ds, \quad (1)$$

where α and β are parameters controlling the curve’s tension and rigidity. The external energy, E_{ext} , is defined in terms of image features. In our case, this is the image intensity: $E_{ext} = I(x, y)$. Thus, the active contour has minimum energy when it matches the central line of the vessel. The curve $X(s)$ that minimizes Equation (1) must satisfy the Euler equation and can be written as follows [15]:

$$\frac{\partial X(s, t)}{\partial t} = \alpha \frac{\partial^2 X(s)}{\partial s^2} - \beta \frac{\partial^4 X(s)}{\partial s^4} - \nabla E_{ext}. \quad (2)$$

When $X(s, t)$ converges ($t \rightarrow \infty$), the left hand side of Equation (2) equals 0 and then the Euler equation has a valid solution.

3.3.2 Boundary and Initial Conditions of the OES

Both boundary and initial conditions are needed to solve Equation (2). In a traditional active contour, where we have a closed curve, no boundary conditions are required. These are necessary though in our case, where the active contour is employed as an open-end curve. Fortunately, these conditions can be determined by refining the starting and ending points of the initialized lines from Step I of the algorithm (Section 3.2). We outline this refinement process in the following steps and we depict it in Figure 3:

1. Refine the positions of the two end points (A and B in Figure 3(a)) of the initial contour line. Refinement is accomplished by searching for pixels with the maximum intensity along a horizontal line (Figure 3(b)). The searching radius should be approximately equal to the width of the vessel. The refined end points (A' and B' in Figure 3(a)) reside in the middle line of the vessel.
2. The boundary conditions are given in terms of the locations of the newly found central line points (A' and B'). For example, at the end A', the boundary condition consists of points A' and x'_1 . x'_1 is one of the linearly interpolated points between A' and B' that is the closest one to A'. Similarly, B' and x'_N are the boundary conditions for the other end of the line (see

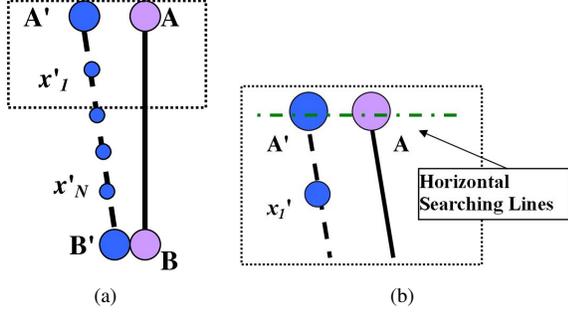


Figure 3. The determination of the boundary and initial conditions of the OES.

Figure 3(a)). The boundary points are already either on the middle line of the vessel or very close to it and remain anchored during the curve evolution process.

3. The initial conditions are points A' and B' and the full series of interpolated points between them, x'_1, x'_2, \dots, x'_N , as shown by the blue dotted line in Figure 3(a).

3.4 Step III: Vessel Boundaries

After locating the vessel's central line, we need to find its outer boundaries. It is rather challenging to obtain the exact vessel boundaries in thermal imagery due to the heat-diffusion effect. The fuzzy boundaries disqualify standard edge detection methods. Instead, we use a special operator to find the vessel edges. The operator searches the maximum temperature gradient along the vessel's radial direction at each point of the active contour. This operator is basically a bottom-up approach. The benefit is obvious: given that the OES has found the correct central lines of vessels, the boundary searching avoids any other sharp gradients that are found in the ROI and do not belong to the vessels.

The operator is implemented as follows (see Figure 4):

1. Evenly select $N + 1$ points on the active contour S_0, S_1, \dots, S_N . Then, at each point draw a line normal to the tangential direction (see Figure 4(a)).
2. Along each of the normal lines, extract a $3 \times M$ rectangular image matrix, which centers at the current active contour point and is aligned with the normal line (see Figure 4(b)).
3. Convolve a 3×3 mask with the $3 \times M$ rectangular matrix. A typical mask is $\begin{pmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{pmatrix}$. If the

vessel density profile is of a bell shape in the radial direction, the convolution function should be similar to the function displayed in Figure 4(c).

4. The width of the vessel ($w_- + w_+$) at each point S_n is determined by the location of the extreme values of the convolution operation. These extreme values indicate the maximum gradient on each side of the vessel. (see Figure 4(c)).
5. The final vessel boundaries are obtained by connecting the corresponding $N + 1$ boundary points on each side of the central line (see Figure 4(a)).

4 Experimental Results

Two experiments were carried out to validate the performance of the new segmentation tool. The first experiment used simulated vessels. The simulated image enabled us to find a relation between the algorithm's performance and the image quality. The second experiment involved real vessels in 65 forehead thermal images.

Two different performance measures were used to evaluate the results:

- *Accuracy*, which is based on the confusion matrix: $Accuracy = (TP + TN) / (TP + TN + FP + FN)$, where TP is *True Positive*, TN is *True Negative*, FP is *False Positive*, and FN is *False Negative*. In other words, Accuracy is simply the percentage of the correctly classified pixels. The higher the Accuracy is, the better the segmenter.
- *Hausdorff distance* [6], which measures the closeness of two vessels' contours. The Hausdorff distance H is defined as $H(P_1, P_2) = \max(\min(d(P_1, P_2)))$, where the $d(P_1, P_2)$ is the Euclidean distance between the point set P_1 and the point set P_2 , i.e., it is the maximum distance of a set P_1 from the nearest point in the other set P_2 . In our problem, P_1 is the segmented vessel and P_2 is the ground truth. The lower the Hausdorff distance is, the better the segmenter.

4.1 Simulated Vessels

In the first experiment, we generate a series of simulated vessel images to evaluate the segmentation method. The size of the simulated ROI is 100×80 pixels. Within the ROI, two sigmoid functions are used to simulate the supraorbital vessels. The average width of the vessels is equal to 7 pixels. One vessel resides on the left and one on the right half of the ROI image respectively. We add 5% to 20% uncertainty on the vertical orientation, the curvature degree, and the width of the vessel. Hence, the appearance

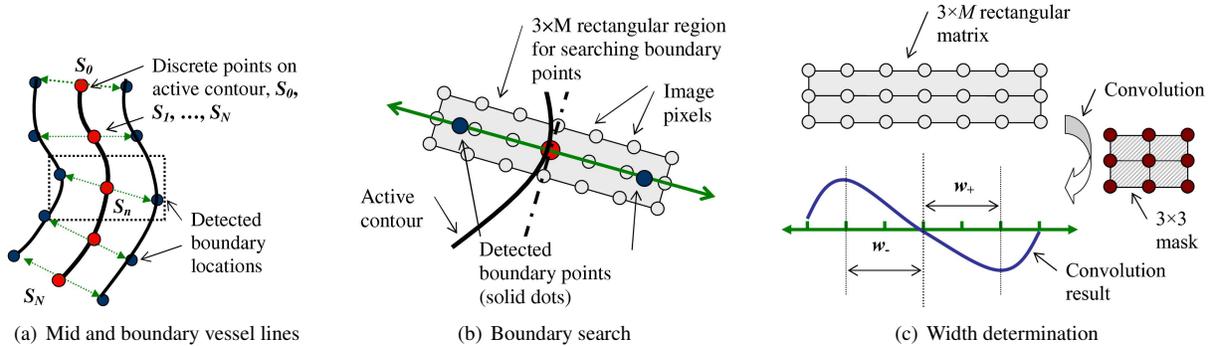


Figure 4. Localization of vessel boundaries.

of simulated vessels varies in accordance with numerous observations of forehead thermal images. The objective is to determine the degree of noise that the method can tolerate before its performance breaks down.

First we define the noise level of the simulated image as follows:

Definition 1: *If the image pixels representing the vessels have intensity value 1, the image pixels representing the non-vessels (background) have random intensity values that are uniformly distributed. This uniform distribution has lower bound 0, and upper bound nl . Then the image noise level is defined as nl .*

After having generated vessel configurations, we generate different noise levels for the background pixels. The noise levels we select range from 0.1 to 2.1. Thereafter, a 15×15 Gaussian filter is used to smooth the whole ROI image. Ten simulated vessel images are shown in Figure 5, where the vessels' central lines and boundaries are annotated. It is evident that the segmentation performance deteriorates as the noise level increases. The increasingly noisy background gradually erodes the vessel imprints. The algorithm fails to find the vessels when $nl = 1.9$.

The simulation dataset consists of 100 different vessel configuration images for each noise level. We applied the OES algorithm to these simulated vessel images and evaluated the performance by two measures: Accuracy and Hausdorff distance. In Figure 6, we see that the Accuracy of the segmentation maintains high values (>0.9) until nl reaches 1.7. The Accuracy decreases dramatically when nl increases from 1.7 to 2.1. On the other hand, the Hausdorff distance increases sharply after $nl = 1.7$, which also indicates deterioration in performance.

From this experiment, we conclude that the new algorithm can locate the vessels even with a fair amount of noise ($nl \leq 1.7$). The algorithm fails to segment the vessels only when the intensity of the background noise approaches

the intensity of the vessels. In that case, the vessel outline blends into the background.

4.2 Real Vessels

In the second experiment, we chose 65 thermal forehead images from 65 different subjects. The ground truth vessel images were created by manual delineation at a super-resolution level based on our best anatomical knowledge. The ground-truth results from two different experts were reconciled and the composite ground-truth set was used in our validation. In this experiment, we compared the results of the proposed method with two other segmentation methods, *10% hottest thresholding* (10P) and *Tophat morphological method* (Tophat). The 10P is a thresholding method that selects the top 10% values of the ROI. The Tophat is a morphological method, similar to the one we used in vessel initialization (see Section 3.2).

We divided the set of 65 real images into two groups according to their quality: normal group (34) and abnormal group (31). So far, we have not found a strict way to quantitatively define the quality of vessel images. The criteria of grouping are qualitative. The differences between the two groups are in most cases clear, but sometimes are subtle. Quality-based grouping facilitates more targeted evaluation.

4.2.1 The Normal Group

Six segmentation examples from the normal group are displayed in Table 1. Each ROI is processed by three different algorithms, 10P, Tophat, and OES. All of the depicted vessel images belonging to the normal group, have two clear supraorbital vessel imprints. The 10P method finds most of the "hottest" pixels from the vessels; many times, however, it fails to capture the entire vessel because the temperature profile along the vessel's axial direction is not even. The Tophat method shows better performance compared with the 10P, since it is able to find the local ridge-like peaks.

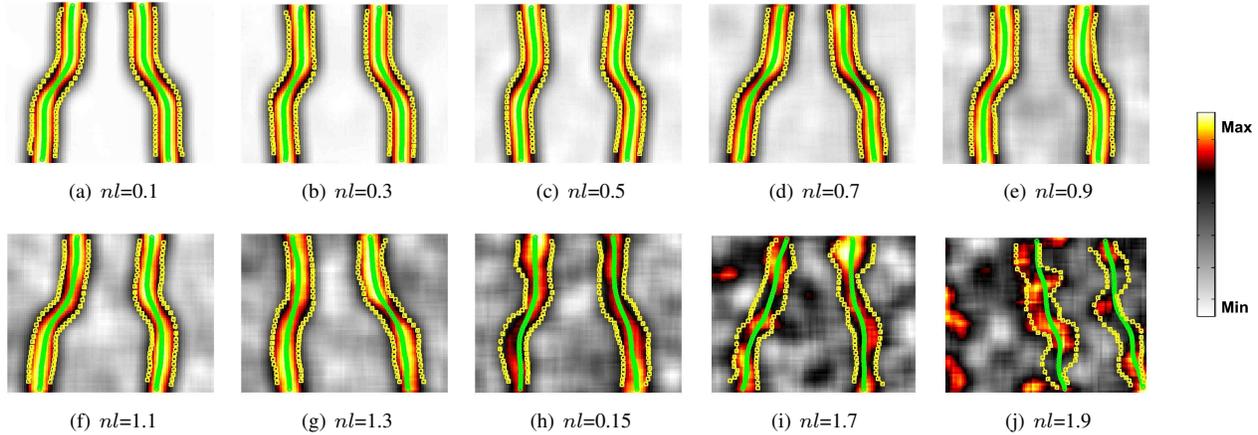


Figure 5. Vessel simulation with different noise levels (nl). The detected central lines of the vessels are marked in green. The detected vessel boundaries are marked in yellow.

We observe that it classifies correctly most of the vessel pixels. But, one problem of the Tophat is that a number of locally hot pixels in non-vessel region are classified as vessel pixels. The OES method fully captures the vessels without being fooled by the fuzzy background, thanks to its mixed bottom-up and top-down approach.

The performance measures, Accuracy and Hausdorff distance, are used to evaluate these algorithms. We see in Figure 7(a) the average Accuracy is 0.859 for the OES, 0.791 for the Tophat, and 0.782 for the 10P. The Hausdorff distance is in Figure 7(b). The average Hausdorff distance is 7.30 for the OES, 12.58 for the Tophat, and 14.98 for the 10P. From these results, out of 34 normal images the OES definitely outperforms other two. Also, the Tophat scores better points than the 10P. The quantitative results correlate well with the visualizations in Table 1.

4.2.2 The Abnormal Group

Six examples from the abnormal group are given in the second part of Table 1. From these “abnormal” examples, we see that the forehead thermal imprints have significant inter-individually differences. Also, some of the real challenges are becoming evident: images feature vessels with unclear imprints (Sub 05), or multiple branches (Sub 20), or broken central lines (Sub 32), or highly asymmetric ones (Sub 39), or distracting objects (e.g., hairs in Sub 62), or only a single branch (Sub 71).

Because of the low quality in the abnormal group images, the performance of all three segmentation methods has been negatively affected. The average Accuracy is 0.785 for the OES, 0.728 for the Tophat, and 0.757 for the 10P. The average Hausdorff distance is 11.95 for the OES, 12.74 for the Tophat, and 15.21 for the 10P. These numbers indi-

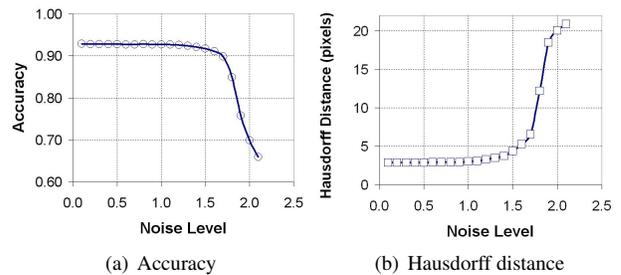


Figure 6. Evolution of segmentation performance for the simulated vessels under different noise levels (nl).

cate that the OES is still relatively better from the other two methods when it comes to abnormal vessel images.

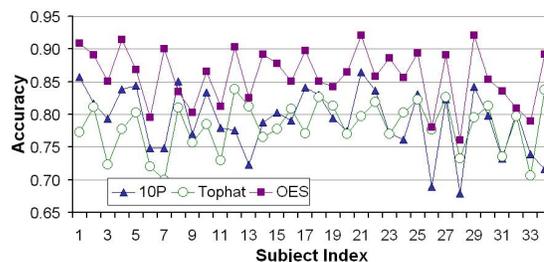
5 Discussion and Conclusions

We have developed a new segmentation method for extracting the thermal imprints of the supraorbital vessels in thermal imagery. The algorithm has three parts. First, it uses an automatic vessel detection method, which is based on the Hough transform, to initialize the vessel locations. Second, an active contour model, the OES, finds the central lines of the major vessels. Third, the boundaries of those vessels are determined by a special operator. The proposed method combines top-down and bottom-up approaches to achieve optimal results.

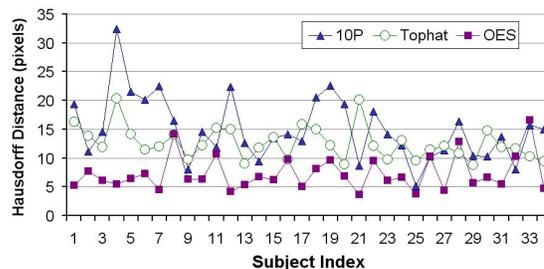
Two experiments have been conducted to evaluate the performance of the new algorithm. In the first experiment, we demonstrated that the OES can segment simulated ves-

Table 1. Visualization of segmentation results from the various methods. (Raw: the original ROI, GT: Ground-Truth, 10P: 10 Percent thresholding, Tophat: Tophat morphological method, OES: Open End Snake)

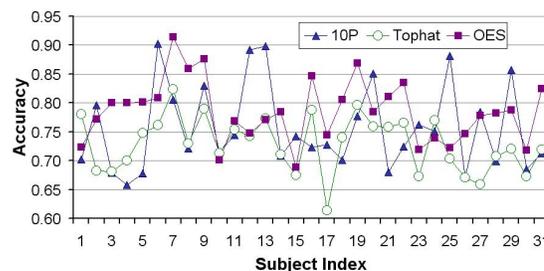
Normal Group					
Sub	Raw	GT	10P	Tophat	OES
Sub03					
Sub04					
Sub26					
Sub28					
Sub61					
Sub81					
Abnormal Group					
Sub	Raw	GT	10P	Tophat	OES
Sub05					
Sub20					
Sub32					
Sub39					
Sub62					
Sub71					



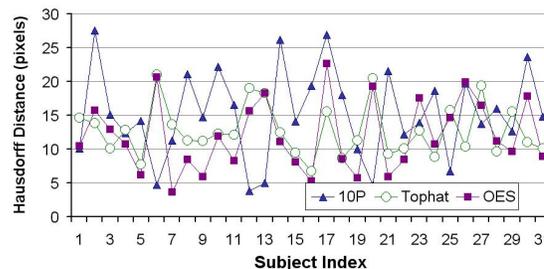
(a) Accuracy, the normal group



(b) Hausdorff Distance, the normal group



(c) Accuracy, the abnormal group



(d) Hausdorff Distance, the abnormal group

Figure 7. Performance comparison of the three different segmentation methods.

sels with high noise level ($n_l=1.7$). In the second experiment, we showed that the OES method outperformed other vessel segmentation methods according to two measures, Accuracy and Hausdorff distance. The new method may find wide applicability in segmenting thermal vessel imprints.

Some limitations of the OES algorithm include:

1. The number of the vessels in the ROI has to be indicated by the operator. The OES can not automatically determine the number of vessels.
2. The OES deteriorates with decreasing image quality.
3. The operator needs to be skillful enough to select the correct ROI on the forehead before the OES is applied. An automatic ROI localization tool may solve this problem in the future.

6 Acknowledgments

This material is based upon work supported in part by the Defense Academy for Credibility Assessment (DACA) and in part by the National Science Foundation (NSF) under Grant No. 0414757 and 0741581 entitled “Interacting with Human Physiology.” Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agencies.

References

- [1] E. Borenstein, E. Sharon, and S. Ullman. Combining top-down and bottom-up segmentation. In *Proceedings of the 2004 Conference on Computer Vision and Pattern Recognition Workshop (CVPRW'04)*, volume 4, page 46, Washington, DC, USA, 2004.
- [2] P. Branco, P. F. L. Encarnacao, and P. Bonato. Faces of emotion in human-computer interaction. In *SIGCHI Conference on Human Factors in Computing Systems*, pages 1236–1239. ACM Press, 2005.
- [3] P. Ekman, W. Friesen, and J. Hager. *The Facial Action Coding System*. London: Weidenfeld & Nicolson (World), second edition, 2002.
- [4] J. Fei, Z. Zhu, and I. Pavlidis. Imaging breathing rate in the CO₂ absorption band. In *Proceedings of the 27th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 700–5, Shanghai, China, September 1-4 2005.
- [5] M. Garbey, N. Sun, A. Merla, and I. Pavlidis. Contact-free measurement of cardiac pulse based on the analysis of thermal imagery. *IEEE Transactions of Biomedical Engineering*, 54(8):1418–1426, 2007.
- [6] D. Huttenlocher, D. Klanderman, and A. Rucklidge. Comparing images using the Hausdorff distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(9):850–863, September 1993.
- [7] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1(4):321–331, 1987.
- [8] C. Kirbas and F. Quek. A review of vessel extraction techniques and algorithms. *ACM Computing Surveys*, 36(2):81–121, June 2004.
- [9] K. Larson, R. Hazlett, B. Chaparro, and R. Picard. Measuring the aesthetics of reading. In *Proceedings of Human Computer Interaction 2006*, pages 41–56, 2006.
- [10] J. Levine, I. Pavlidis, and M. Cooper. The face of fear. *Lancet*, 357:1757, 2001.
- [11] T. McInerney and D. Terzopoulos. Deformable models in medical image analysis: A survey. *Medical Image Analysis*, 1(2):91–108, 1996.
- [12] B. Moxham, C. Kirsh, B. Berkovitz, G. Alusi, and T. Cheeseman. Interactive head and neck (CD-ROM). Primal Pictures, December 2002.
- [13] I. Pavlidis, J. Dowdall, N. Sun, C. Puri, and J. Fei. Interacting with human physiology. *Computer Vision and Image Understanding*, 108(1-2):150–170, 2007.
- [14] C. Puri, L. Olson, I. Pavlidis, J. Levine, and J. Starren. Stresscam: Non-contact measurement of users’ emotional states through thermal imaging. In *Proceedings of the 2005 ACM Conference on Human Factors in Computing Systems*, pages 1725–8, Portland, Oregon, April 2-7 2005.
- [15] C. Xu and J. Prince. Snakes, shapes, and gradient vector flow. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 7(3):359–69, March 1998.
- [16] Z. Zhu, J. Fei, and I. Pavlidis. Tracking human breath in infrared imaging. In *Proceedings of the IEEE 5th Symposium on Bioinformatics and Bioengineering*, pages 227–31, Minneapolis, Minnesota, October 19-21 2005.