Proceedings of the 28th IEEE
EMBS Annual International Conference
New York City, USA, Aug 30-Sept 3, 2006

Segmentation and Landmark Identification in Infrared Images of the Human Body

C.L. Herry, M. Frize, *Senior Member IEEE*, and R.A. Goubran

Abstract—**The segmentation and landmark identification in infrared images of the human body are key steps in a computerized processing of large database of thermal images. The segmentation task is especially challenging due to specific characteristics of thermal images. Few papers deal with segmentation techniques for clinical infrared images and available segmentation methods (e.g. for breast or military thermal images) do not perform well on other types of images. This paper presents a few strategies for the automated segmentation and registration of anatomical landmarks on thermal images of arms and hands. The segmentation method is based on mathematical morphological operations and simple rule based processing easily available through prior knowledge about the objects of interest.**

I. Introduction

THE registration of infrared images of the human body is a key aspect of computerized processing. It facilitates greatly the subsequent processing and analysis of the images and allows fast, automated analysis of images or image sequences. In order to perform accurate registration of landmarks in infrared images, a clean segmented image highlighting the region of interest is necessary. However, the segmentation task can be quite challenging, perhaps more than for visible images for several reasons. First, the intensities depend on object temperature, object surface properties, surface orientation, wavelength, etc. and are not necessarily uniform across the same object. The environmental conditions at the time of imaging are also of great influence on the resulting images. In addition, the range of infrared images is typically much less than that of visible images, leading to low contrast, poor resolution and less texture information [1].

Fortunately, these undesirable effects may be somewhat attenuated by carefully controlled imaging conditions [2] and when the object of interest has a high emissivity (e.g. human skin in the 8-12μm range). For example, infrared imaging of the body under controlled environmental conditions may sometimes produce higher quality images in terms of contrast, Signal to Noise Ratio (SNR) and resolution than typical infrared imaging used in military targets. However, despite carefully controlled protocols, the resulting images are not always ideal. In addition, key features that are successfully used in the processing of visible images cannot be used reliably when dealing with infrared images, for instance color or texture information [1]. In addition, the range of gray levels corresponding to the temperature of the skin can overlap strongly with that of the background, for instance when dealing with images of the extremities (fingers), making the segmentation task challenging.

Most papers on segmentation of clinical infrared images found in the literature focus on asymmetry analysis of breast thermal images [3-6]. The methods using Hough transform and/or morphological operations perform relatively well for the segmentation of breast and breast tumors. Other methods found in the general literature on infrared image processing include fuzzy clustering [7], local thresholding using a standard deviation map [8]. Ginesu et al. compared adaptive thresholding, rank order statistics and morphological processing (skeletonization) to identify foreign bodies in infrared images of food [9]. They found that their three methods are similar in performance although the rank order statistical method produced marginally better results. However, none of the methods mentioned above produced satisfactory results when implemented on our database of infrared images of hands, arms, heads and necks. As a result, we developed our own segmentation techniques and we present some of the results in this paper along with partial registration of infrared images of hands and upper-arms.

II. Methodology

A. Database of thermal images

The database used in this paper consists of a large number (several thousands) of thermal images of various views of body regions (hands, arms, shoulders, head and neck). The images were taken from volunteers participating in distinct projects undertaken by the authors. Ethical approval was obtained for each project and informed consent was given by each volunteer.

Images have been collected since 2003 using a Thermacam SC500 (FLIR Systems Inc.) and an Indigo Omega thermal cameras.

This work was supported in part by a discovery grant from the National Science and Engineering Research Council (NSERC)

C. Herry is a Ph.D. student in the department of Systems and Computer Engineering at Carleton University, Ottawa, Ontario, Canada (corresponding author: herrycl@ieee.org).

M. Frize. is a Professor in the Dept. of Systems and Computer Engineering at Carleton University (with a joint appointment in the School of Information Technology and Engineering at University of Ottawa), Ontario, Canada.

R.A. Goubran is a Professor and Chair of the Department of Systems and Computer Engineering, Carleton University, Ottawa, ON, Canada.

1-4244-0033-3/06/$20.00 ©2006 IEEE.

B. Preprocessing

Preprocessing of thermal images involves correcting for possible artifacts created by the imaging system, including the removal of any noise. The choice of a noise removal method is linked to the type of imaging system used, its quality and the type of noise created. The noise in modern infrared systems is usually well characterized by a Gaussian distribution, but is not always additive. Infrared imaging systems based on photo-detectors are better modeled by Poisson noise and adapted noise removal procedures must be considered [9-10]. In the application considered hereafter (i.e. infrared images of the human body in closed and controlled environment), the signal to noise ratio was typically high enough so that extra processing was required on most images. Noisy images were denoised using a stationary wavelet transform method that kept fine details in the image.

C. Segmentation

The typical background of clinical infrared images is relatively simple with few objects present. However, the intensities of the objects of interest and that of the background regions often overlap, for example when imaging a person with cold fingers at ambient temperature. Also, the background intensity is often not constant across the image, due to variations in the emissivity of background objects. As a result, thresholding techniques (global, locally adaptive) and clustering methods (for example fuzzy C-means) for example, did not perform very well on our database of infrared images. They often missed important regions of the object of interest (e.g. whole fingers in poorly contrasted images) or returned larger than expected regions where all fine details were lost. In view of these shortcomings, we decided to opt for a method based on morphological processing of the edges detected with a Canny edge detector. The Canny edge detection technique is well known and has the advantage of being able to detect most weak edges[12], which is essential if one wants to capture fine details.

Our segmentation procedure for thermal images of hands works as follows:

1. Detect edges using Canny edge detector with default parameters (automated high and low thresholds based on 30% of edge pixels).
2. Add binary thresholded image using Otsu's optimal global threshold method.
3. Eliminate connected components around the borders and line-like components based on the eccentricity of equivalent best fitting ellipse.
4. Morphological closing of the connected components to consolidate the largest region.
5. Find the smallest convex region enclosing the largest connected component and remove outside components.
6. Remove small components within convex region using automated threshold.
7. Morphological closing using a line as a structuring element, the orientation of which is determined by the direction of the largest connected component in the image. This is followed by a flood-fill operation. The result of this step is to effectively bridge gaps (if any) between main component (e.g. wrist and palm) and extremities (tip of fingers).
8. Final cleaning of spurious objects using morphological opening with a small structural element as well as a threshold on the area of each objects remaining.

The contour from the resulting region is extracted and compared with ground-truth manually segmented images. Similar processing was performed on thermal images of the upper-arm (dorsal and anterior views).

D. Identification of landmarks

Once the main region of interest is found, it is rotated so that the direction is approximately horizontal. This is chosen arbitrarily to facilitate the identification of anatomical landmarks. Fingers can be characterized relatively easily using profile lines across the hand. The junction of the wrist and the thumb is found by analyzing the top and bottom contour of the hand. Using these landmarks, any part of the hand region may be automatically segmented. This is useful when analyzing large databases of thermal images.

Similarly, we processed thermal images of the upper-arm and looked for the junction of the wrist and hand as well as the elbow area. This in turn allowed us to highlight automatically predefined anatomical regions of the upper-arm (for instance a superficial muscle or a vein) and to analyze the evolution of the surface temperature across a sequence of images.

III. RESULTS

All algorithms were implemented in Matlab R14 (The MathWorks Inc.) and tested on 658 thermal images of upper-arms (anterior view) and 1216 thermal images of hands (left and right, dorsal and palmar views).

For each image, a manually segmented contour was used as ground truth segmentation. Contours produced by our method were compared with the ground truth segmentation, using a modified version of the Hausdorff distance between the two contours as an indicator of performance. The estimated distance $d(.,.)$ between two contours was therefore calculated as follows:

$$d(A,B) = \frac{\left(sum\left(A \cap \widetilde{B}\right) + sum\left(\widetilde{A} \cap B\right)\right)}{2 \times \max\left(sum(A), sum(B)\right)} \tag{1}$$

Where A, B are the binary contour images, $\widetilde{A}, \widetilde{B}$ are the dilated versions of A, B respectively (with a 3x3 square structural element), $\cap$ is the logical AND operation and

$sum(.)$ returns the number of non-zero pixels.

The results were classified into three subjective categories (good, acceptable, poor) based on their percentage of correctly identified contour pixels. Good segmentation implied that at least 90% of the experimental contour agrees with the manual segmentation (ground truth). Segmented images were deemed acceptable when at least 65% of the experimental contour agreed with the actual contour. This included most images where it is difficult to separate fingers close to each other. Finally, failed segmentation implied that the contour did not match the actual contour very well. Note that the resulting contour still enclosed the whole hand even if the segmentation of fingers failed. This means that further (finer) segmentation is possible within the region found.

Fig. 1, 2, and 3 show a few representative examples of good, acceptable and poor segmentation respectively.

We found that the overall success rate of the segmentation technique is quite satisfactory with our database. 78.8 % of segmented images were either acceptable or good (12.2% and 66.6% respectively). Most segmented images in the acceptable category only had minor deviations from the ideal contour (see for example Fig. 2) that could easily be fixed by simple post-processing steps (using a priori information about the shape of the hand for example). It is important to note, however, that a database comprising images captured in extremely poor conditions would significantly lower the number of successful segmentations.

Fig. 4 shows an example of delimited regions on the upper-arm corresponding to superficial muscles (from left to right, top to bottom: extensor carpi radialis longus and brevis, extensor digitorum, extensor carpi ulnaris and flexor carpi ulnaris). The regions were obtained automatically based on the detection of landmarks at the wrist and elbow (not shown).

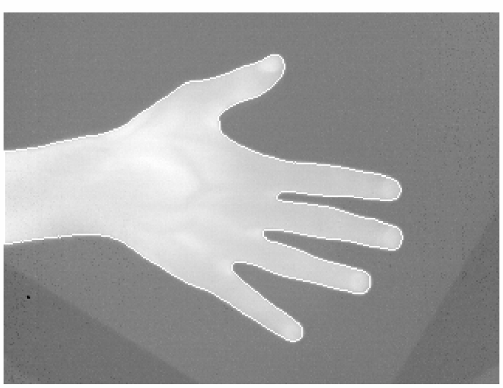

Fig. 1. Example of a good segmentation of a dorsal view of the right hand. 99% of the pixels on the calculated contour match the true contour.

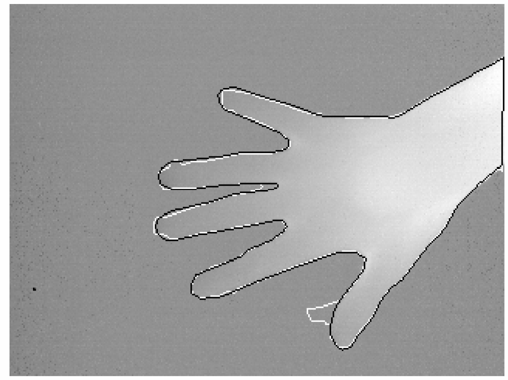

Fig. 2. Example of a segmented hand. The white and black contours are the experimental and true contours respectively. Note the poor contrast between fingertips and background. 85% of the pixels on the calculated contour match the true contour.

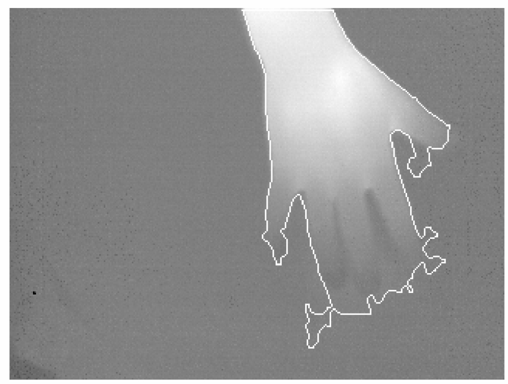

Fig. 3. Example of a failed segmentation. Although no critical part of the hand was cropped, the contour encloses extra regions that are not relevant for the analysis. Only 55% of the pixels on the calculated contour match the true contour.

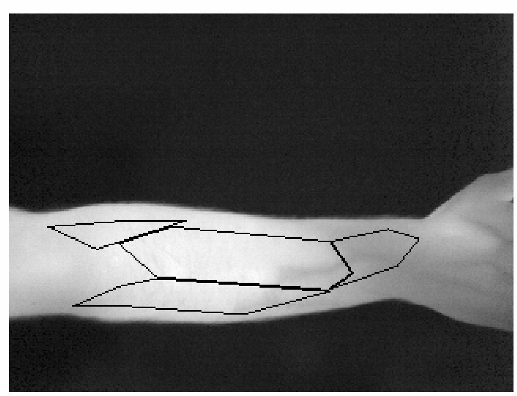

Fig. 4. Example of segmented regions of the upper-arm. The regions correspond to some of the superficial muscles of the upper arm.

IV. Discussion

The method presented performs relatively well on most of our images. Problems occurred when two fingers were too close to each other (a gap of a few pixels). The resulted segmented image would typically show the two fingers together. This is not necessarily problematic for the landmark identification as the fingers can still be separated by guessing where the middle line falls. Another difficulty was to find the contour of a cold finger, when the whole finger (and sometimes all the fingers) was very faint, even when performing manual segmentation. Our algorithm was usually able to overcome "cold gaps" (when the gap is not too large) in fingers with the morphological closing and a line as a structural element. Better performance may be achieved with a Hough transform, for instance by looking for lines in the region where the fingers are. It is then possible to fill any potential gaps in the contour of the fingers. The computational requirement is increased drastically however. Also, as pointed in the previous section, the quality of the database influences greatly the results of the segmentation. For example, we applied our technique on a few images taken in extremely poor conditions where it is almost impossible to distinguish the fingers from the background. The experimental segmentation was not satisfactory but was able to distinguish parts of the contour that the naked eye could not see at first glance.

Finally, let us point out that our implementation has not yet been optimized for real-time processing. Each image may take up to 2 seconds to process on a windows based Pentium IV 1.4Ghz processor. However, we believe that an optimized version of our method would reduce that time by a factor of at least ten.

V. Conclusion

We presented a segmentation technique for infrared images of the human body. Our method achieved good results on a database of thermal images of hands and upper-arms. A more refined and multi-stage segmentation could allow better performance when the contrast is very poor between parts of the region of interest and the background. Future work will investigate other body regions and focus on feature-driven segmentation in order to automatically classify clinical thermal images.

References

[1] Y. Fang, K. Yamada, Y. Ninomiya, B.K.P. Horn and I. Masaki, "A shape independent method for pedestrian detection with far-infrared images," *IEEE Trans. Vehicular Technology*, vol. 35(6), Nov. 2004.

[2] R. Pochaczevsky. Technical guidelines, edition 2. *Thermology*, 2(2), 1986.

[3] E.Y.K. Ng , Y. Chen and L.N. Ung, "Computerized breast thermography: Study of image segmentation and temperature cyclic variations," *J. Med. Eng. Tech.*, vol. 25(1), pp. 12-16, 2001.

[4] N. Scales, C. Herry and M. Frize, "Automated image segmentation for breast analysis using infrared images," *Proc. 26th Annual Conf. IEEE EMBS*, San Francisco, CA, USA, Sept. 2004.

[5] J. Qi and J. Head, "Asymmetry analysis using automatic segmentation and classification for breast cancer detection in thermograms," *IEEE EMBS Conf. Turkey*, 2001.

[6] X. Tang and H. Ding, "Asymmetry analysis of breast thermograms with morphological image segmentation," *Proc. 27th Annual Conf. IEEE EMBS*, Shanghai, China, Sept. 1-4, 2005.

[7] S.G. Sun and S.W. Park, "Automatic target recognition using boundary partitioning and invariant features in forward-looking infrared image," *Opt. Eng.*, vol. 42(2), pp. 524-533, Feb. 2003.

[8] S. Greenberg, S.R. Rotman, H. Guterman, S. Zilberman and A. Gens, "Region-of-interest-based algorithm for automatic target detection in infrared images," *Opt. Eng.*, vol. 44(7), July 2005.

[9] G. Ginesu, D.D. Giusto, V. Märgner and P. Meinlschmidt, "Detection of foreign bodies in food by thermal image processing," *IEEE Trans. Industrial Electronics*, vol. 51(2), April 2004.

[10] C. Herry and M. Frize, "Quantitative assessment of pain through clinical digital infrared thermal imaging", *Biomed. Eng. Online*, vol. 3(19), 2004.

[11] R.D. Nowak and R.G. Baraniuk, "Wavelet-Domain Filtering for Photon Imaging Systems," *IEEE Trans. Image Proc.*, vol. 8(5), pp. 666–678, May 1999.

[12] R.C. Gonzalez and R.E. Woods, *Digital Image Processing*, 2nd edition, Pearson Education, 2002.