

Chapter VIII

Automated Overlay of Infrared and Visual Medical Images

G. Schaefer

Aston University, UK

R. Tait

Nottingham Trent University, UK

K. Howell

Royal Free Hospital, UK

A. Hopgood

De Montfort University, UK

P. Woo

Great Ormond Street Hospital, London, UK

J. Harper

Great Ormond Street Hospital, London, UK

ABSTRACT

Medical infrared imaging captures the temperature distribution of the human skin and is employed in various medical applications. Often it is useful to cross-reference the resulting thermograms with visual images of the patient, either to see which part of the anatomy is affected by a certain disease or to judge the efficacy of the treatment. In this chapter, we show that image registration techniques can be effectively used to generate an overlay of visual and thermal images and provide a useful diagnostic visualisation for the clinician.

INTRODUCTION

Medical infrared imaging captures the natural thermal radiation generated by an object at a temperature above absolute zero. It is non-invasive, non-contact, passive, radiation-free and complementary to anatomical investigations based on x-rays and three-dimensional scanning techniques such as CT and MRI, and often reveals problems when the anatomy is otherwise normal. Often, visual and infrared images of the patient are taken to relate inflamed skin areas to the human anatomy which in turn is useful for medial diagnosis as well as for assessing the efficacy of any treatment. Currently, this process requires great expertise and is subject to the individual clinician's ability to mentally map the two distinctly different images.

Image registration is one of the most important medical image processing techniques and is used to geometrically align or overlay two images taken from different sensors, viewpoints or instances in time. Both images are typically aligned through a combination of scaling, translation and rotation. Registration is often used to monitor growth, verify the effects of treatment and make comparisons of patient data with anatomically normal subjects.

In this chapter, we show how image registration can be effectively used to overlay medical infrared images and visual images of a patient in order to relate areas that are of interest due to their thermal pattern to the human anatomy. After capturing both thermal and visual images, the visual image is pre-processed with a skin detection technique to separate the patient from the background. Using an intensity-based registration algorithm which requires no user interaction, visual and infrared images are then superimposed and the generated overlay presented to the user for visualisation purpose. The generated system is currently in use at Royal Free and Great Ormond Street hospitals to assess patients suffering from morphea.

BACKGROUND

Thermal Infrared Imaging

Advances in camera technologies and reduced equipment costs are among the factors that have led to an increased interest in the application of thermal imaging in the medical fields (Jones, 1998). Thermal medical imaging (or medical infrared imaging) uses cameras with sensitivities in the infrared to provide a picture of the temperature distribution of the human body or parts thereof. It is a non-invasive, non-contact, passive, radiation-free technique that is often being used in combination with anatomical investigations based on x-rays and three-dimensional scanning techniques such as CT and MRI and often reveals problems when the anatomy is otherwise normal. It is well known that the radiance from human skin is an exponential function of the surface temperature which in turn is influenced by the level of blood perfusion in the skin. Thermal imaging is hence well suited to pick up changes in blood perfusion which might occur due to inflammation, angiogenesis or other causes. Asymmetrical temperature distributions as well as the presence of hot and cold spots are known to be strong indicators of an underlying dysfunction (Uematsu, 1985). Computerized image processing and pattern recognition techniques have been used in acquiring and evaluating medical thermal images (Plassmann & Ring, 1997; Wiecek, Zwolenik, Jung, & Zuber, 1999) and proved to be important tools for clinical diagnostics. Thermal imaging has been successfully employed in, among others, detecting breast cancer (Head, Wang, Lipari, & Elliott, 2000; Anbar et al., 2001), diagnosing Raynaud's phenomenon (Merla et al., 2002), and local scleroderma (morphea) (Black et al., 2002).

Image Registration

Image registration is a method used to geometrically align or overlay two images taken from different sensors, viewpoints or instances in time. Both reference (fixed) and sensed (moving) images are typically aligned through a combination of scaling, translation and rotation although a universal registration method is not possible due to the wide variety of noise and geometric deformations caused by the diverse methods of image capture available. Registration is often used to monitor growth, verify the effects of treatment and make comparisons of patient data with anatomically normal subjects (Maintz & Viergever, 1998).

While there is a wealth of registration techniques available, most of them can be divided into two categories: landmark-based and intensity-based algorithms (Zitova & Flusser, 2003). Landmark-based methods rely on a selection of several corresponding control points in both images and then seek the best image transform between these landmark pairs. Obviously, registration accuracy is determined by the selection of the landmarks. While manual selection might provide good registration results it is a cumbersome method that relies heavily on the user. On the other hand, automatic control point detection and extraction methods are difficult to derive and typically depend on certain image characteristics. In contrast, intensity-based registration approaches do not require identification of landmarks, rather they utilize the image data directly. Additional masking can be introduced to emphasize special features. The basic intensity approach consists of transform optimization, image re-sampling and feature-matching stages. Feature matching is the most fundamental stage and is achieved through the use of a similarity metric in which a degree of likeness between corresponding images is calculated. These steps are iteratively repeated until a good enough match between the two images is achieved.

In multi-modal image registration applications, the data to be registered stem from two different capturing devices, as opposed to single-modal tasks where images are retrieved using the same sensor type. The measure of alignment between images of differing modality can be based on the assumption that although different in value, regions of similar intensity in the fixed image will correspond to regions of similar intensity in the moving image (Woods, Mazziotta, & Cherry, 1993). Also, for all pixels in corresponding regions, the ratio of their intensities should vary only slightly. As a consequence, alignment is achieved when the average variance of this ratio is minimized. Crucially, this idea can be realized through the construction of a feature space also commonly referred to as a joint probability distribution (Hill, Hawkes, Harrison, & Ru, 1993). The joint probability distribution represents a two-dimensional plot containing combinations of intensities taken from corresponding coordinates in both images. Instead of identifying regions of similar intensity directly within the images, combinations of intensities are analyzed using the joint probability distribution.

During the registration process, a variation in alignment between the images causes changes in appearance of the joint probability distribution. When correctly aligned, corresponding structures in both images overlap causing a clustering of intensities combinations. In contrast, misalignment causes structures in the fixed image to overlap with structures in the moving image which are not their counterpart. This results in the dispersal of intensity combinations within the joint probability distribution. Pluim, Maintz, and Viergever (2003) have demonstrated the effects of registering an image with itself using varying degrees of translation and rotation. Based on the changing regions within the joint probability distribution, measures of dispersion which guide the registration process have been proposed and successfully implemented (Maes, Collignon, Vandermeulen, Marchal, & Suetens, 1997; Viola & Wells, 1997).

Importantly, the concept of mutual information which appears low when intensity combinations are dispersed and high when intensity combinations are clustered, is a well recognized and accepted method of similarity calculation. The main advantage of employing mutual information is that the type of dependency between two variables does not have to be specified and as a result complex mappings can be modeled.

As no assumptions about the nature of the capture device need to be made, image registration algorithm can be generalized and applied to data representing a variety of modalities.

METHODS

Our system is designed to provide an overlay of visual and thermal images to be presented to a clinician. Superimposing both image types allows relating areas of certain thermal patterns to the anatomy of the person as well as monitor efficacy of any treatment. After pre-processing both images to segment the patient from the background, an intensity-based registration algorithm with mutual information similarity metric is employed to geometrically align the two images. This is realized using a series of intelligent agents collaborating on a blackboard architecture to provide an efficient and effective framework for image registration. Superimposed images are presented to the user for visualization. The user can adjust the relative importance of the individual modalities in an interactive manner.

Image Pre-Processing and Segmentation

As the overlay is achieved through application of an intensity-based image registration algorithm, and as the background of thermal images taken in a temperature controlled lab is typically in stark contrast to the patient's body heat, the same

contrast must be achieved for the visual image. That is, the background needs to be separated from the foreground (i.e., the patient). In our approach, we make use the fact that thermal imaging picks up the skin temperature and hence employ a skin detection technique on the visual image (Schaefer, Tait, & Zhu, 2006). We adopt, with some variations, a computationally simple method introduced in (Fleck, Forsyth, & Bregler, 1996) which is based on the fact that the hues of human skin occupy only a small region in color space. The algorithm which operates on the visual (RGB) image proceeds in the following steps:

1. The R, G, and B values at each pixel are transformed into a log-opponent color representation by

$$\begin{pmatrix} I \\ R_g \\ B_y \end{pmatrix} = \begin{pmatrix} L(G) \\ L(R) - L(G) \\ L(B) - \frac{L(G) - L(R)}{2} \end{pmatrix} \quad (1)$$

with

$$L(C) = 105 \log_{10} (C + 1 + n) \quad C = \{R, G, B\} \quad (2)$$

where n represents some random noise (in the range (0;1)) to prevent banding artifacts in dark regions.

2. A measure of texture amplitude T is then derived from the intensity channel by building a difference image of the original image and a median filtered version of it. The resulting texture channel is then again median filtered as are the chromaticity components (at a finer scale compared to the texture channel).
3. Next, hue H and saturation S are calculated as:

$$H = \tan^{-1} \left(\frac{R_g}{B_y} \right)$$

$$S = \sqrt{R_g^2 + B_g^2} \quad (3)$$

4. Pixels that fall within a certain hue-saturation range and do not exceed a texture threshold are identified. In particular, all pixels that fall within

$$\{T < 5 \text{ and } 110 < H < 155 \text{ and } 5 < S < 60\}$$

or

$$\{T < 5 \text{ and } 130 < H < 170 \text{ and } 30 < S < 130\}$$

are marked as skin pixels.

5. Using morphological operations holes are filled and edges smoothed to provide the final output of the skin detector.

Registration

Using the detector described above regions in the visual image are identified that correspond to skin colors and hence to the patient. Non-skin areas are removed by setting their pixel values to 0 (black). In the thermal images, patients are usually well separated from the background in controlled lab conditions so little pre-processing is required. If necessary, an adaptive thresholding algorithm can be applied to improve the separation of patient from non-patient areas.

Once both image types have been prepared, the image registration process is initiated. As we are interested in a fully automatic method that should be applicable on a wide range of different images (different patients, poses, etc.) we adopted an intensity-based approach to registering the two images. In intensity-based techniques pixel information is utilized and the best alignment is derived as that which optimises a pre-defined similarity metric between the registered images. The steps involved are transform optimization, image re-sampling and similarity computation which are applied in an iterative manner until the process has converged. We employ a gradient-descent optimizer, B-spline interpolation for the re-sampling and a mutual information measure as similarity metric.

The similarity computation stage is complex and represents a considerable performance bottleneck when employed in an iterative registration process. Based on a worker/manager model, a distributed blackboard system is employed to spread computational workload between a number of intelligent agents and improve performance of the algorithm (Tait, Schaefer, Hopgood, & Nolle, 2006).

The alignment process begins with partitioning of visual and thermal images into segments; regions of interest and initial transform parameters are also added to the blackboard. On addition of transform parameters, for each sample point in the visual segment a corresponding intensity in the thermal segment is calculated. Importantly, interpolation with a B-spline basis function is used to calculate intensities at non-grid positions. A local joint probability distribution realized as a Parzen histogram is then generated from retrieved intensities by a worker process and placed onto the blackboard. As histograms are progressively generated for all image segments, worker processes become inactive.

Deactivation of all worker processes triggers the manager process to construct a global Parzen histogram from local histograms stored on the blackboard. By estimating the density distribution of the global histogram, an entropy value in the form of a gradient is then calculated. Regular step gradient descent optimization is employed by the manager process to advance transform parameters in the direction of the gradient. In each iteration of the optimization process, the step length through the transform search space is calculated using a bipartition scheme. Once updated, transform parameters are propagated to all worker processes and the procedure is repeated. The algorithm converges when the step length fails to exceed a threshold, or a maximum number of iterations have been reached. On convergence, image segments are accumulated by the manager process and a registered image is assembled.

The framework is flexible in nature and can be employed to either distribute the registration

of single images as outlined above or to distribute the processing of many images in batch mode.

Although the system by default performs intensity-based registration to provide a fully automated system, in certain cases intensity-based approaches do not produce an accurate overlay. In such instances the user has the possibility to switch to a landmark-based algorithm. Corresponding control points then need to be specified by the user before an accurate overlay is generated.

Once an appropriate transform has been found and image registration performed (typically the visual image is selected as reference and the thermal one as sensed image), a composite image is created. This is simply performed by computing a weighted sum of the respective pixel values of the original visual image and the thresholded thermogram. Equal weights will generate an aver-

age of the two images whereas different weight factors will put more emphasis on one of the two modalities. The actual weightings between the two modalities can be determined by the user and can be controlled interactively.

RESULTS

We have used a set of thermal-visual image pairs to evaluate our proposed method. Two examples are provided in Figures 1 and 2, each of which shows the original visual image, the thermogram, the visual image segmented based on the output of the skin detection step, and the final overlaid image. In both cases, the final image was weighted as 80% visual and 20% thermal. As can be seen, in both cases an accurate overlay of the two image types is achieved.

Figure 1. Example 1 of thermal-visual overlay: Original visual image, thermogram, segmented visual image, composite image (from left to right, top to bottom).

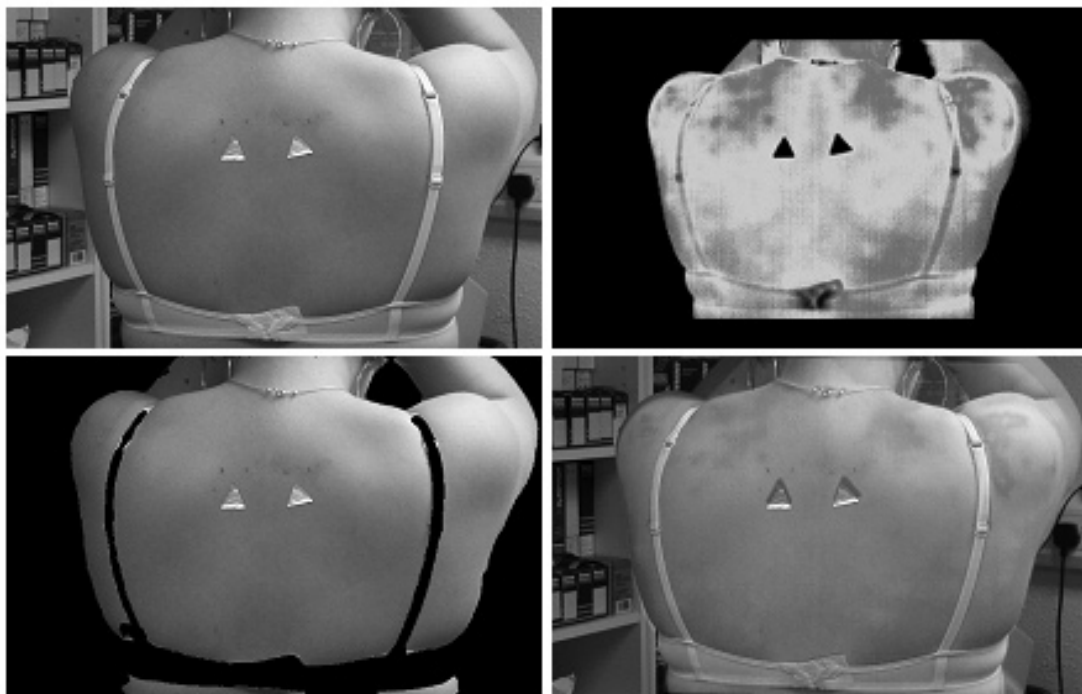


Figure 2. Example 2 of thermal-visual overlay: Original visual image, thermogram, segmented visual image, composite image (from left to right, top to bottom).




Figure 3. Example of different weightings between visual and infrared images in generating the overlay: 75% visual-25% infrared; 50%-50%; 25%-75%; 90%-10% (from left to right, top to bottom).




Figure 4. Example of landmark-based registration. Control points have been set for the visual image and copied over to the infrared image.



The generated images are currently being used in the assessment of morphea (localized scleroderma) patients. In Figure 2, the warmer area of the chest overlay indicates the distribution of a morphea lesion.

In Figure 3, we show how the weights between the two image modalities can be set so as to put more or less emphasis on one of the original images.

Finally, Figure 4 gives an example of the application where the user is performing a landmark- rather than an intensity-based registration. Control points are placed in one of the modalities which can then be copied to the other image. Control points can be adjusted as a complete set through rotation, translation and scaling operations or individually. If control points are placed correctly, landmark-based registration provides high accuracy overlay images.

CONCLUSION

In this chapter, we have shown how image registration can be employed to perform multi-modal medical image visualization and overlay. In particular we have demonstrated its application to overlay thermal and visual images for medical diagnosis. Following a pre-processing step based on skin detection to perform background segmentation, intensity-based multi-modal registration is performed through a set of intelligent agents communicating via a blackboard structure.

Future Research Directions

The system is currently in use at Royal Free and Great Ormond Street hospitals where it is used in the diagnosis and treatment of morphea (local scleroderma) patients. The methods described can

however be used in any scenario where overlay of infrared and visual images is deemed useful. Furthermore, it is sufficiently generic to be used for overlaying medical images of other modalities.

It should also be noted that, although outside the scope of this chapter, the framework can also be employed for registering 3D volume datasets in an equally efficient and effective way (Tait, Schaefer, Hopgood, Zhu, 2006).

While currently the application is restricted to aligning static images, future versions will incorporate superimposed dynamic sequences. In addition, we are currently working at providing overlays of images that have been taken at different times, typically many months apart, of the same patient. These should provide a visualization of how a disease is developing and whether the current treatment shows any effects, and should hence represent a valuable tool for medical diagnosis.

REFERENCES

- Anbar, N., Milescu, L., Naumov, A., Brown, C., Button, T., Carly, C., & AlDulaimi, K. (2001). Detection of cancerous breasts by dynamic area telethermometry. *IEEE Engineering in Medicine and Biology Magazine*, 20(5), 80-91.
- Black, C., Murray, K., Howell, K., Harper, J., Atherton, D., Woo, P., et al. (2002). Juvenile-onset localized scleroderma activity detection by infrared thermography. *Rheumatology*, 41(10), 1178-1182.
- Fleck, M., Forsyth, D., & Bregler, C. (1996). Finding naked people. *4th European Conference on Computer Vision*, 2, 593-602.
- Head, J., Wang, F., Lipari, C., & Elliott, R. (2000). The important role of infrared imaging in breast cancer. *IEEE Engineering in Medicine and Biology Magazine*, 19, 52-57.
- Hill, D., Hawkes, D., Harrison, N., & Ru, C. (1993). A strategy for automated multi-modal image registration incorporating anatomical knowledge and imager characteristics. *Information Processing in Medical Imaging*, 687, 182-196.
- Jones, B. (1998). A reappraisal of infrared thermal image analysis for medicine. *IEEE Trans. Medical Imaging*, 17(6), 1019-1027.
- Maes, F., Collignon, A., Vandermeulen, D., Marchal, G., & Suetens, P. (1997). Multi-modality image registration by maximization of mutual information. *IEEE Transactions on Medical Imaging*, 16, 187-198.
- Maintz, J., & Viergever, A. (1998). A survey of medical image registration. *Medical Image Analysis*, 2, 1-36.
- Merla, A., Di Donato, L., Di Luzio, S., Farina, G., Pisarri, S., Proietti, M., et al. (2002). Infrared functional imaging applied to Raynaud's phenomenon. *IEEE Engineering in Medicine and Biology Magazine*, 21(6), 73-79.
- Plassmann, P., & Ring, E. (1997). An open system for the acquisition and evaluation of medical thermological images. *European Journal of Thermology*, 7, 216-220.
- Pluim, J., Maintz, J., & Viergever, M. (2003). Mutual information-based registration of medical images: A survey. *IEEE Transactions on Medical Imaging*, 22, 986-1004.
- Schaefer, G., Tait, R., & Zhu, S. (2006). Overlay of thermal and visual medical images using skin detection and image registration. *28th IEEE Int. Conference Engineering in Medicine and Biology*, (pp. 965-967).
- Tait, R., Schaefer, G., Hopgood, A., & Nolle, L. (2006). Automated visual inspection using a distributed blackboard architecture. *Int. Journal of Simulation: Systems, Science & Technology*, 7(3), 12-20.
- Tait, R., Schaefer, G., Hopgood, A., & Zhu, S. (2006). Efficient 3-D medical image registration using a distributed blackboard system. *28th IEEE*

Int. Conference Engineering in Medicine and Biology, (pp. 3045-3048).

Uematsu, S. (1985). Symmetry of skin temperature comparing one side of the body to the other. *Thermology*, 1, 4-7.

Viola, P., & Wells, W. (1997). Alignment by maximization of mutual information, *International Journal of Computer Vision*, 24, 137-154.

Wiecek, B., Zwolenik, S., Jung, A., & Zuber, J. (1999). Advanced thermal, visual and radiological image processing for clinical diagnostics. *21st IEEE Int. Conference on Engineering in Medicine and Biology*, (p. 1108).

Woods, R., Mazziotta, J., & Cherry, S. (1993). MRI-PET registration with an automated algorithm. *Journal of Computer Assisted Tomography*, 17, 536-546.

Zitova, B., & Flusser, J. (2003). Image registration methods: A survey. *Image and Vision Computing*, 21, 977-1000.

Additional Reading

Allen, R., Ansell, B., Clark, R., Goff, M., Waller, R., & Williamson, S. (1987). Localized scleroderma: Treatment response measured by infrared thermography. *Thermology*, 2, 550-553.

Bankman, I. (2000). *Handbook of medical imaging*. Academic Press.

Birdi, N., Shore, A., Rush, P., Laxer, R., Silverman, E., & Krafchik, B. (1992). Childhood linear scleroderma: A possible role of thermography for evaluation. *Journal of Rheumatology*, 19, 968-973.

Black, C. (1999). Scleroderma in children. *Advances in Experimental Medicine and Biology*, 455, 35-48.

Brown, L. (1992). A survey of image registration techniques. *ACM Computing Surveys*, 325-376.

Diakides, N., & Bronzino, J. (Eds.). (2007). *Medical infrared imaging*. CRC Press.

Hill, D., Hawkes, D., Harrison, N., & Ru, C. (1993). A strategy for automated multi-modal image registration incorporating anatomical knowledge and imager characteristics. *Information Processing in Medical Imaging*, 687, 182-196.

Jeongtae, K., & Fessler, J. (2004). Intensity-based image registration using robust correlation coefficients. *IEEE Transactions on Medical Imaging*, 23, 1430-1444.

Kakumanu, P., Makrogiannis, S., & Bourbakis, N. A survey of skin-color modeling and detection methods. *Pattern Recognition*, 3, 1106-1122.

Martini, G., Murray, K., Howell, K., Harper, J., Atherton, D., Woo, P., et al. (2002). Juvenile-onset localized scleroderma activity detection by infrared thermography. *Rheumatology*, 41, 1178-1182.

Mattes, D., Haynor, D., Vesselle, H., Lewellen, T., & Eubank, W. (2001). Non-rigid multi-modality image registration. *Medical Imaging 2001: Image Processing*, (pp. 1609-1620).

Nolle, L., Wong, K., & Hopgood, A. (2001). DARBS: A distributed blackboard system. *Research and Development in Intelligent Systems*, 18, 161-70.

Pratt, W. (1974). Correlation techniques of image registration. *IEEE Transactions on Aerospace and Electronic Systems*, 10, 353-358.

Ring, E., & Ammer, K. (2000). The technique of infrared imaging in medicine. *Thermology International*, 10, 7-14.

Roche, A., Malandain, G., Ayache, N., & Prima, S. (1999). Towards better comprehension of similarity measures used in medical image registration. *Lecture Notes in Computer Science*, 1679, 555-566.