

Research

Open Access

Quantitative assessment of pain-related thermal dysfunction through clinical digital infrared thermal imaging

Christophe L Herry*¹ and Monique Frize^{1,2}

Address: ¹Department of Systems and Computer Engineering, Carleton University, 1125 Colonel By Drive, Ottawa, ON K1S 5B6, Canada and ²School of Information Technology and Engineering, University of Ottawa, 161 Louis Pasteur, Ottawa, ON K1N 6N5, Canada

Email: Christophe L Herry* - cherry@sce.carleton.ca; Monique Frize - mfrize@connect.carleton.ca

* Corresponding author

Published: 28 June 2004

Received: 07 April 2004

BioMedical Engineering OnLine 2004, **3**:19 doi:10.1186/1475-925X-3-19

Accepted: 28 June 2004

This article is available from: <http://www.biomedical-engineering-online.com/content/3/1/19>

© 2004 Herry and Frize; licensee BioMed Central Ltd. This is an Open Access article: verbatim copying and redistribution of this article are permitted in all media for any purpose, provided this notice is preserved along with the article's original URL.

Abstract

Background: The skin temperature distribution of a healthy human body exhibits a contralateral symmetry. Some nociceptive and most neuropathic pain pathologies are associated with an alteration of the thermal distribution of the human body. Since the dissipation of heat through the skin occurs for the most part in the form of infrared radiation, infrared thermography is the method of choice to study the physiology of thermoregulation and the thermal dysfunction associated with pain. Assessing thermograms is a complex and subjective task that can be greatly facilitated by computerised techniques.

Methods: This paper presents techniques for automated computerised assessment of thermal images of pain, in order to facilitate the physician's decision making. First, the thermal images are pre-processed to reduce the noise introduced during the initial acquisition and to extract the irrelevant background. Then, potential regions of interest are identified using fixed dermatomal subdivisions of the body, isothermal analysis and segmentation techniques. Finally, we assess the degree of asymmetry between contralateral regions of interest using statistical computations and distance measures between comparable regions.

Results: The wavelet domain-based Poisson noise removal techniques compared favourably against Wiener and other wavelet-based denoising methods, when qualitative criteria were used. It was shown to improve slightly the subsequent analysis. The automated background removal technique based on thresholding and morphological operations was successful for both noisy and denoised images with a correct removal rate of 85% of the images in the database. The automation of the regions of interest (ROIs) delimitation process was achieved successfully for images with a good contralateral symmetry. Isothermal division complemented well the fixed ROIs division based on dermatomes, giving a more accurate map of potentially abnormal regions. The measure of distance between histograms of comparable ROIs allowed us to increase the sensitivity and specificity rate for the classification of 24 images of pain patients when compared to common statistical comparisons.

Conclusions: We developed a complete set of automated techniques for the computerised assessment of thermal images to assess pain-related thermal dysfunction.

Background

The skin temperature distribution of a healthy human body exhibits a contralateral symmetry [1]. Temperature distribution that shows asymmetrical patterns is usually a strong indicator of abnormality [2-4], but the converse is not always true since some pathological conditions may exhibit bilateral thermal dysfunction. In such cases other signs of abnormalities in the temperature distribution need to be found [5,6]. Some nociceptive and most neuropathic pain pathologies are associated with an alteration of the thermal distribution of the human body in the form of hyperthermic or hypothermic regions [5]. Since the dissipation of heat through the skin occurs for the most part in the form of infrared radiation, infrared thermography is the method of choice to study the physiology of thermoregulation and the thermal dysfunction associated with pain. The early literature on medical thermography focused on qualitative interpretation of thermograms; this involved determining abnormal thermal variations of the skin by means of a visual assessment of pseudo coloured or grey-level thermograms with the help of isothermal displays, visual localisation of hot or cold spots, and visual detection of symmetry [7-12].

The task of decrypting thermograms and extracting useful and reliable information was complex, even for highly trained medical thermographers, since it relied upon the subjectivity of the human visual ability to distinguish between variations in intensity levels representing temperature distribution in thermograms. In addition, the use of pseudo-colours for mapping the temperatures of a thermogram was also criticised for its subjectivity due to the psychological effect of certain colours, which may skew the observer's performance [13]. As a result, thermographic research examined general quantification techniques for specific problems in order to reduce the subjectivity of the assessment of thermograms [14]. Many past and recent publications discuss thermal dysfunction associated with pain, however, to our knowledge none so far applied comprehensive computerised techniques to the assessment of thermal images of persons experiencing pain.

Methods

Objectives

The overall goal of this work was to automate as much as possible a computerised assessment of thermal images of pain in order to support clinicians' decision making. Our approach consists of several steps. First, the thermal images are pre-processed to reduce the noise introduced during the initial acquisition of the images and to extract irrelevant background. Then, potential regions of interest are identified in a semi-automated manner, using fixed dermatomal subdivisions of the body; they are also identified in an automated manner based on an isothermal

analysis and segmentation techniques. Finally, we assess the degree of asymmetry between contralateral regions of interest using statistical computations and distance measures between comparable regions.

Data collection

Hundreds of thermal infrared images of pain patients were digitally recorded on magnetic tapes by Monique Frize and her team at the Pain Clinic of the Moncton Hospital, Moncton, New Brunswick, Canada, between 1981 and 1984, using an AGA Thermovision 680 medical infrared camera system and an AGA OSCAR recording and retrieval system. Each image had 128×128 pixels and 256 grey intensity levels. We randomly selected 43 patients from this database, 25 females and 18 males. In June 2002, a second set of thermal images was taken of 15 healthy volunteers, following approval by the Research Ethics Committee at Carleton University. The healthy subjects consisted of 9 males and 6 females, who did not have any known medical history of chronic pain or ongoing acute pain. The same infrared camera system and image format previously used in the 1980s were also used for these new tests. This population served as our control population. The images were transferred to a PC and saved in a Tag-based Image File Format (TIFF) for easy manipulation. The programming work was carried out with MATLAB and the Image Processing toolbox. Examples of thermal images used in this study are provided in Figures 1 to 2. Figure 1 shows original images while Figure 2 shows pseudo-colour isothermal images. Thermographic examination were carried out in a room without windows, maintained at a uniform, constant temperature of 23°C and relative humidity ranging from 50% to 70%. Subjects were asked to follow the patient preparation guidelines published by the American Academy of Thermology [15] prior to the examination. Subjects were asked to disrobe the body areas to be imaged and stand up for 15 minutes to allow equilibration with the ambient temperature. Images were taken of the upper-back (primarily thoracic region), lower-back (primarily lumbar region) and upper- (thighs) and lower-legs (calves).

Methodology

Pre-processing of thermal images

The characteristics of the thermal imaging system used in this study was assumed to be an almost ideal infrared system, where most of the noise types inherent to the design are lower than the photon noise generated by the background and signal itself. It is known that photon noise follows a Poisson distribution and is signal-dependent [16][17]. Common noise removal techniques based on a Gaussian approximation to Poisson noise do not work well due to the signal-dependent, spatially variant nature of photon noise. In order to remove the noise efficiently while retaining the relevant information contained in

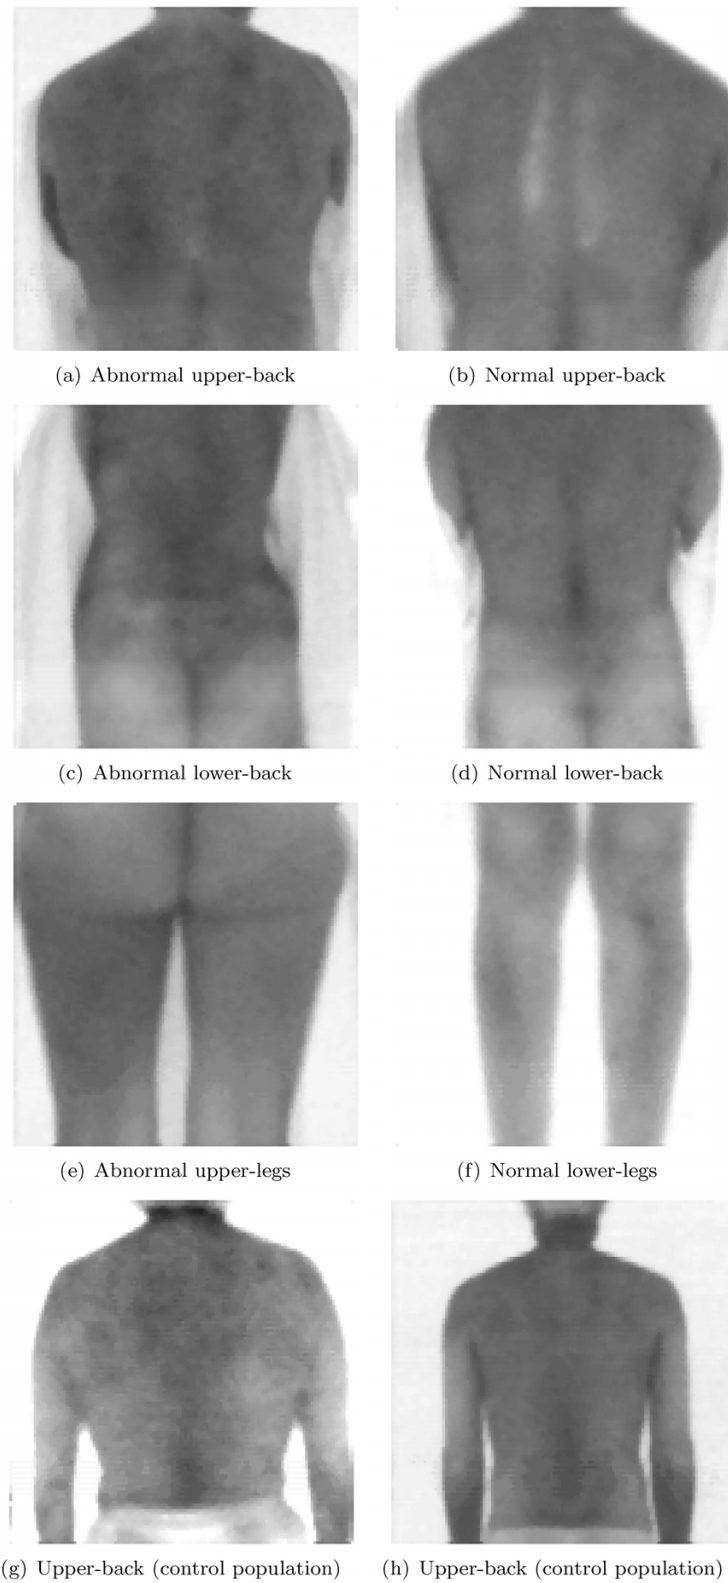


Figure 1
Examples of thermograms considered in this study. Thermograms of the upper-back, lower-back, upper- and lower-legs for pain patients and control subjects.

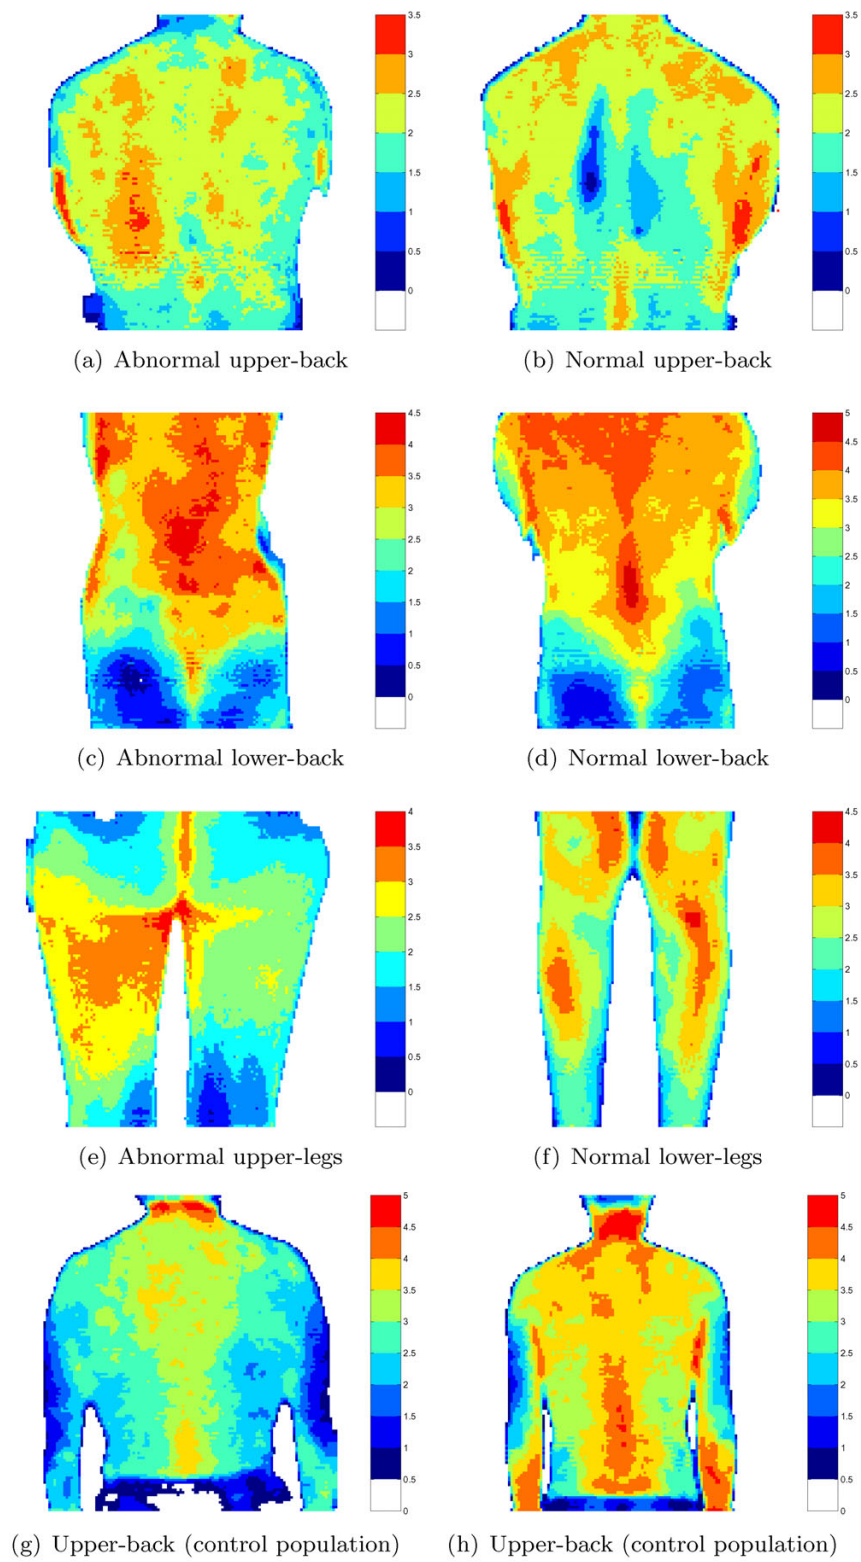


Figure 2

Pseudo-colour isothermal display of thermograms. Pseudo-colour isothermal display of thermograms of the upper-back, lower-back, upper- and lower-legs for pain patients and control subjects. The colorbar on the right gives the range of temperature in °C.

thermal images, we adapted a wavelet-based Poisson noise removal technique that was suggested by Nowak and Baraniuk [18]. They considered the successive images with increasing intensity levels from the photon counting process, until they obtain the final image at the end of the observation period. The authors measured the error between a wavelet-domain filtered version of each image with the next image, until the final image. The sum of the errors is called the PREdictive Sum of Squares (PRESS), which is dependent on the filter (or threshold) used. Minimising the PRESS criterion yields the optimal filter (or threshold). The authors also showed that the sub-images are not required in order to compute the optimal threshold, since the latter reduces to:

$$\mathbf{h}^{(PRESS)} = \left(\frac{\theta^2 - \alpha \widehat{\sigma}^2}{\theta^2} \right)_+ \quad (1)$$

where θ refers to the wavelet transform of the image, $\widehat{\sigma}^2$ is the corresponding matrix of noise variance estimates, α is a regularisation parameter (the threshold is optimal for $\alpha = 1$), and $(\cdot)_+$ sets to zero the negative thresholds.

The noise variance is estimated by projecting the image onto the square of the wavelet and scaling functions, yielding what the authors called a Discrete Wavelet Squared Transform (DWST). The DWST may be implemented efficiently using recursions similar to that of the Fast Wavelet Transform. The choice of the wavelet bases is important to preserve the original features of the thermal images. In particular, they need to be well localised and have a certain degree of smoothness. Possible candidates among orthogonal wavelets are Daubechies wavelets and Battle-Lemarié spline orthogonal wavelets. They have been shown to provide acceptable results in medical image processing.

The next step is to remove the background, which in this work refers to parts of the thermal images unnecessary for a correct assessment. Correct segmentation of images is usually a difficult task when images represent complex objects and scenes. However, thermal images of body areas are usually simple enough that thresholding methods of segmentation perform satisfactorily, with additional image processing steps if necessary. As a result we propose a background extraction method combining an optimal thresholding method adapted from Tsai's work [19] and additional morphological processing steps. The algorithm is summarised as follows:

1. The histogram of the image is computed and smoothed with a Gaussian function to obtain a desired number of modes.

2. The extreme values of the resulting function are then computed as well as their corresponding grey-level.

3. A first threshold is set to the strongest minimum, that is the last minimum remaining after successive smoothing. A second threshold is selected to be the first stable minimum after successive smoothing. The final threshold is then given by a weighted sum of the two thresholds.

4. Isolated pixels wrongly labelled as background pixels are corrected. The approximate boundaries of the image are superposed to the thresholded images and any holes created are filled.

5. The remaining regions of interest are then roughly segmented and the smaller regions are discarded as they likely correspond to background parts. The same procedure is performed on the background regions.

6. Finally the remaining possible holes in the foreground regions are filled based on the size and shape of the regions.

Identification of region of interests (ROIs)

In order to compare the results that were derived from our population of normal patients with the values published in the literature [3,4], we divided the body into multiple squares or rectangles that other authors have used, for example Goodman [3] or Uematsu [4].

The manual division by Uematsu followed the areas of the skin called dermatomes that are innervated by the major peripheral nerves. Being able to link the data from those regions of interest with the corresponding nervous system that feeds them is believed to be more meaningful for the assessment of the thermograms. Goodman used a grid of squares, for the back and front, and rectangles for the arms, legs and extremities. The grid was adapted to the bodily dimensions of each subject and could be moved around to cover any desired areas. The statistics were computed for various numbers of slices or frames of the unit square or rectangle and results from the analysis compared to that of Uematsu's. For comparison purpose, we devised a hybrid of the two divisions for the back, chest and legs, which consists in Uematsu's division with two additional regions in the back.

Our algorithm first looks for some easily identifiable landmarks, namely the shoulders and the hips for the chest and back, the buttock cleft for the back, the area around the knee for the legs. This is achieved through the following steps:

1. The vertical symmetry line separating contralateral parts of the body is first sought by detecting the edges and

averaging the middle distance between the outer boundaries of the body.

2. For the right and left sides of the body, we defined the function corresponding to the distance between the outer boundaries and the symmetry line, thus obtaining two functions representing the left and right profiles of the body.

3. The profile resulting from averaging both left and right profiles is smoothed with a Gaussian function and we look for extreme values of the resulting function.

4. The first significant maximum is found to be representative of the extremity of the shoulder and the level corresponding to 80% of the maximum is used to find the coordinate of the extreme point for choosing the ROIs.

5. The flank level is found through a similar method for the upper back images but looking for the first significant minimum of the pro-file function instead of the maximum. For the lower back images, a search for the minimum distance between left and right side, above the upper point of the buttock cleft, provides acceptable results.

6. The buttock cleft is found through an analysis of the horizontal profile lines across the lower part of the images, when a middle edge line is present. The distance between the lower point and higher points is computed for each profile lines, yielding a function of this distance. The first minimum of this function is chosen to be approximately the upper extremity of the buttock cleft.

7. Finally, the knee is easily detected by looking at the minimal distance between each edge lines for both legs.

Based on the previous landmarks, the positioning of the various boxes delimiting the ROIs defined by both Uematsu and Goodman was determined in a straightforward manner. A second method was implemented to determine more specific regions of interest than dermatomes. It is based on an isothermal analysis of the entire pre-processed image, to determine hotter and colder regions. The step for the computation of each isotherm is chosen according to the range of intensity values covered in the image to provide some adaptivity to the generation of isotherms. In practice, the step is set to one tenth of the total range of intensity values.

Also, whenever the sensitivity setting is available, we looked for isothermal regions at regular intervals of 0.5°C, for comparison with common isothermal analysis found in the literature. The same rule is applied for the hotter regions, starting from the maximum and subtract-

ing successive increments of the step. Typically, only the first two to three steps are necessary to yield significant regions. In addition, smaller regions of less than 10 pixels were discarded by thresholding the areas of the regions returned by the previous isothermal computation. Statistical comparisons with regions smaller than the threshold area would have limited significance in most cases.

Moreover, whenever two regions of very different sizes (one large and one small) are found to be too close to each other, they are merged into one larger region. The positions and pixels in the remaining regions are kept for the statistical analysis.

Finally, in order to avoid artifact regions located around the contour of the body area of interest, the image is eroded along the boundaries using a morphological eroding operator. For all methods, a simple algorithm is used to compute the reflection of the regions of interest with respect to the vertical line of symmetry, allowing us to assess the asymmetry of the resulting intensity/temperature distribution.

Asymmetry analysis

We considered a comprehensive set of statistical features that may be able to capture more completely the nature of the distributions in the regions of interest. Those features include the classical moments (mean, variance, skewness and kurtosis), the extreme values, the area, the heat content defined by Montoro and Anbar as the area times the mean temperature [20] and finally the entropy of the region, which may be approximated as:

$$E = - \sum_{k=1}^L p_k \log_2 (p_k) \quad (2)$$

where p_k is the number of pixels in each grey-level $k \in [1, L]$, where L is the number of grey levels, which we chose to be 256.

The outcome for a particular region of interest usually relies on the computation of the difference of statistics between this region and its contralateral counterpart. Some authors also considered the statistics of the pixel-by-pixel difference between two comparable regions [21]. However, we felt that both these approaches were omitting important parameters of the temperature distribution of ROIs, since they looked at a few statistical features derived from assumptions on the nature of this distribution. As a result, we proposed a new statistical approach using distances between histograms of comparable ROIs in order to assess their degree of similitude as it is extensively used in image registration and pattern recognition. The following distance measures were considered: Man-

hattan or absolute distance, Euclidian distance, maximum distance, chi-squared distance, Jeffrey-divergence distance and Mallows distance [22,23]. The first three distances are simply the first three moments of the classical Minkowski distance based on the metric of the same name. The general formula for the Minkowski distance is:

$$D(H,G) = \left(\sum_{i=1}^L |H(i) - G(i)|^p \right)^{1/p} \quad (3)$$

where H and G denote the histograms and $p = 1$ for the Manhattan distance, $p = 2$ for the Euclidian distance and $p = \infty$ for the maximum distance. The chi-square distance is given by:

$$D_{chisquare}(H,G) = \sum_{i=1}^L \left(\frac{H(i) - \frac{H(i)+G(i)}{2}}{\frac{H(i)+G(i)}{2}} \right)^2 \quad (4)$$

The Jeffrey divergence is defined by:

$$D_{JD}(H,G) = \sum_{i=1}^L H(i) \log \frac{H(i)}{\frac{H(i)+G(i)}{2}} + G(i) \log \frac{G(i)}{\frac{H(i)+G(i)}{2}} \quad (5)$$

The Mallows distance is:

$$D_{Mallows}(H,G) = \left(\frac{1}{L} \sum_{i=1}^L |H(i) - G(i)|^p \right)^{1/p} \quad (6)$$

where $H_{(i)}$ and $G_{(i)}$ are the sorted histogram values and p was chosen to be 2.

In addition, we also calculated the Kolmogorov-Smirnov statistic, which is defined as the maximal distance between two cumulative distributions:

$$D_{KS}(H,G) = \max_i |H_{cum}(i) - G_{cum}(i)| \quad (7)$$

where H_{cum} and G_{cum} are the cumulative histograms corresponding to H and G respectively. The Kolmogorov-Smirnov test derived from this statistic was performed for each pair of symmetric regions of interest, since it showed encouraging results in Vavilov *et al.*'s paper [24]. This test has the advantage of being non-parametric, that is, it does not assume that the population of pixels within each region follows a specific distribution.

Results and discussion

Performance of the denoising method

The Discrete Wavelet Squared Transform (DWST) was implemented in a relatively straightforward manner using the one-dimensional and two-dimensional wavelet transforms available in the Wavelab library from the department of Statistics of Stanford University [25]. In order to reduce the memory requirements of the direct algorithm, we chose to decompose each 128×128 thermal image into four 64×64 sub-images and remove the noise from each of these sub-images. In practice, it was found that artifacts affect the quality of the resulting image at the border of each sub-image.

In order to reduce these artifacts, we increased the number of sub-images to nine partially overlapping 64×64 images and we kept the areas away from the interior borders for each sub-image (see Figure 3). This division into sub-images does not completely eliminate the border effect but reduces it to an acceptable level that does not affect the subsequent analysis. Note that a DWST applied to the whole image would remove these border artifacts. A translation invariant DWST may provide a good alternative, although it would increase the computational complexity. In addition, we assumed that our subdivision kept the signal-dependence and Poisson nature of the noise.

A few orthogonal wavelet bases were investigated on thermal images from our database with various noise levels. Nowak and Baraniuk [18] suggested using a simple Haar basis or a Daubechies wavelet basis for nuclear medicine images. Our tests on thermal images showed that the best results quantitatively (based on the reduction in noise variance) and visually (presence of artifacts and ringing effect for instance) were achieved with a Battle-Lemarié spline wavelet basis of the 3rd order.

Furthermore, we varied the parameter α in equation 1 and found that $\alpha > 1$ yielded better denoising results. In particular, $4.5 < \alpha < 7.5$ consistently produced the best trade off between reduction of noise and conservation of the image structure.

Since the computation of mean square errors between original images and denoised images was not readily available, the performance of the denoising was assessed in terms of percentage of edges and contours remaining after denoising, for a comparable noise level. We found that the resulting images from the Press-optimal wavelet-domain filter retained the contours and edges of the thermal images better than when a Fourier domain Wiener filter or a wavelet-domain filter (applied to the square root of the image) was used.

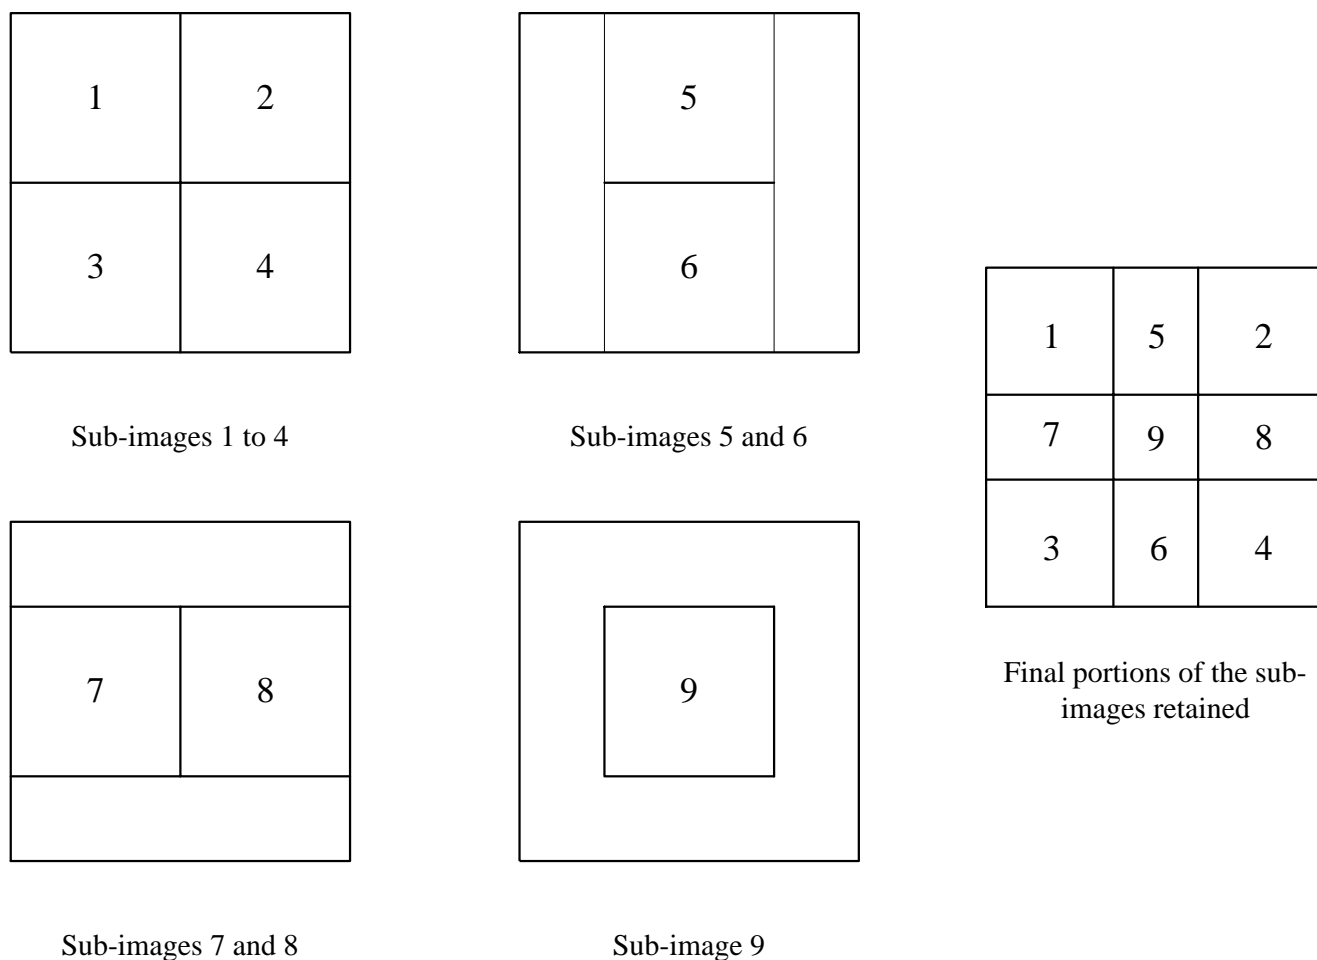


Figure 3
Subdivision of images for the DWST. Division of the original image into 9 sub-images. The numbers indicate the sub-images. The final square on the right shows the final portion of each sub-image that is retained to compute the final, resulting image.

In order to determine the percentage of edges remaining after denoising we summed the total number of pixels returned by a Canny edge detector for increasing threshold level, up to a level where no edge can be detected. We assumed that the Canny edge detector was not affected significantly by the noise and therefore we normalised the number of pixels to that of the noisy image.

Background extraction efficiency

The ratio of the number of images whose background was properly extracted, to the total number of images was slightly higher for the denoised images, as shown in table 1. A correct background extraction, according to our criteria, implies that only the region targeted when the image was taken is present. In our study, for instance, the upper-

back, the lower back or the legs are the targeted regions. The quality of the background extraction was assessed visually and poorly segmented thermal images were discarded manually. The results from table 1 confirm our initial assumption that an efficient denoising could improve the subsequent processing of thermal images. The overall percentage of images whose background was correctly extracted is 88% for the denoised images and 81% for the noisy images. However, it appears that the thresholding part of our background extraction is less affected by the noise than the corrective morphological processing. Therefore, a more complex thresholding strategy may yield similar results for noisy and denoised thermal images.

Table 1: Rate of incorrect background extraction. Rate of incorrect background extraction for images of the lower back, upper-back and legs. The first figure indicates the number of images whose background was incorrectly extracted. The second figure is the total number of images.

	With Denoising	Without Denoising
Lower Back images	9/72	22/72
Upper Back images	12/85	15/85
Legs	22/205	30/205

Automated selection of ROIs

The automated selection of ROIs using a compromise between Uematsu's technique and the more complete technique by Goodman was fairly successful. The determination of the limits for the ROIs based on the shape of the body and on its intensity distribution enabled us to approximately place the boxes corresponding to each ROI. A simple iterative procedure was performed on the boxes to make sure they do not contain any background. However, the difference in body shapes and the lack of symmetry of body areas in some thermal images made it difficult to adapt the ROIs to cover the best area.

Furthermore, we assumed that thermal images would be symmetrical with respect to a vertical line of symmetry; this was verified in most thermal images. However, other thermal images showed areas of the body bent toward one side of the border of the image, or rotated. As a result, even if the limits for the ROIs were correctly identified, the final ROIs did not cover the wanted areas. It may therefore be necessary to consider alternative symmetry curves and define a proper reflection method to get corresponding pixels from contralateral regions.

In addition, we found that division into rectangular boxes based on the dimension of the body may fail to include an abnormality if the latter happens to be in the middle of two boxes. This outlines the necessity for a complementary ROI delimitation method such as the isothermal analysis that works well for asymmetric images and targets potential abnormalities. It produced similar results in most cases: targeted ROIs over possible abnormalities in the thermal pattern distribution in the form of hot or cold spots.

ROIs covering both contralateral parts of the body and ROIs along the line of symmetry were flagged for later analysis, ensuring that the assessment does not fail to detect bilateral abnormalities. A strong gradient of temperature between a flagged region and its surrounding area indicates an abnormality according to our procedure.

We also compared the ROIs produced by the isothermal analysis of fully pre-processed images and of the noisy images with the extracted background. We found that noisy images tend to return a greater number of smaller ROIs that would otherwise be merged into larger ROIs in the case of denoised images. However, smaller regions may not necessarily reflect a true elevation or decrease of temperature but instead indicate spurious, noisy pixels. This justifies discarding regions smaller than a threshold, which was chosen to be 20 pixels for our study. In addition, statistical comparisons of very small samples may not be significant, which further justifies our decision.

Some performance aspects of thermographic analysis

The performance of our quantitative analysis of thermal images in terms of complexity and CPU time was greatly affected by the denoising stage and by our choice of the PRESS-Optimal filter. For instance, the DWST theoretically required $O(21^3 \times 128^2)$ computations per resolution level, but our implementation using smaller sub-images increased the complexity. The average CPU time to process a single image was about 4 minutes, using a Pentium III at 1.2 GHz with 512 MB of memory, MATLAB release 12.1 and Wavelab 802 [25]. In comparison, the rest of the analysis (background extraction, identification of ROIs, asymmetry analysis) required typically less than 3 seconds. Although it is difficult to assess the overall complexity of the thermographic analysis, it is certainly heavily weighted by the denoising step. The other denoising techniques considered require less CPU time than the PRESS-Optimal method, but do not perform as well. If we assume that the CPU time is a decisive factor in the analysis of large databases of thermal images, it would be desirable to improve our implementation of the PRESS-Optimal method or investigate other denoising methods that perform equally well.

Thermal asymmetry of the control population

From our control population of normal healthy subjects, we selected images of the back and legs and performed the pre-processing steps and the selection of regions of interest. In order to validate the normality of the population as far as pain is concerned, we first divided each image using

Table 2: Normal Values for the back. Mean Temperature Difference plus or minus standard deviation for our control population of healthy volunteers.

ROIs	Number of Cases	Mean Temperature Difference ± Standard Deviation
Thoracic (medial)	12	0.09 ± 0.09
Shoulder (posterior)	12	0.12 ± 0.10
Thoracic (lateral)	12	0.09 ± 0.09
Thigh (anterior)	13	0.15 ± 0.11
Thigh (posterior)	14	0.11 ± 0.09
Knee (anterior)	13	0.12 ± 0.13
Knee (posterior)	14	0.15 ± 0.12
Leg (anterior)	13	0.15 ± 0.12
Leg (posterior)	14	0.10 ± 0.06

the predetermined grid from Uematsu et al.'s and Goodman et al.'s paper and compared the statistics derived from these ROIs with the published normal values of thermal asymmetry. Our results are presented in table 2.

All the mean temperature difference values are well within the limits for normal value specified by Ue-matsu and Goodman. Our values are in fact significantly lower, by the order of 0.1 °C, which may partly be explained by the relatively small number of cases in our control population. This certainly had an effect on the standard deviation of the mean temperature differences. Another reason for these low values may be the size of the regions of interest, which for practical reasons was smaller than that of the previous authors.

We pursued our assessment by computing the various distances mentioned previously on all our ROIs and for noisy images and denoised images (without background). Table 3 and 4 summarise the results and give the mean, standard deviation, and upper 99% confidence interval of the mean for the distances between the temperature distribution of ROIs of the back and legs. We did not differentiate between specific regions of the back or legs since we were interested in deriving one threshold for all regions. The threshold would thus be applicable to the ROIs selected by the isothermal analysis, although the approach consisting of defining a threshold for each region is also valid. For each distance, we chose the upper 99% confidence interval as the threshold for determining whether or not a region is abnormal.

Results from noisy and denoised images did not differ significantly and confidence intervals were overlapping. This also confirms that the denoising does not affect the statistical properties of thermal images and thus does not induce a loss of critical information, while we showed that it facilitated the previous analysis.

Using the threshold values from table 3 we assessed the efficiency of each distance measure to discriminate between normal and abnormal regions of interest. For our normal population, this was accomplished by looking at the distances generated from a complete isothermal analysis of thermal images of the back and legs. The distance with the least number of misclassifications may be seen as the more discriminative with respect to normal regions, that is the distance with the best specificity (Number of true negatives with respect to the number of false positives) or the least number of regions identified as positive or abnormal.

Table 5 shows the specificity values for each distance measure. As expected, the specificity values are quite high, since thresholds were applied to the same population from which they were extracted, although on different regions of interests. Table 4 may be viewed as validation results for the thresholds selected, showing their efficiency at identifying normal ROIs.

Finally, we noted that contrary to Vavilov et al. who found that the Kolmogorov-Smirnov statistics allowed the accurate classification of normal regions from abnormal and potentially pathological regions, our experiment suggested that other distance measures might be more suitable to this task. All the other distance measures have similar specificity values. In order to validate our approach and draw some conclusions on the performance of our novel statistical tool for the assessment of pain in thermal images, we first analysed a set of 72 thermal images of the upper back with the fixed division into rectangular ROIs. Then the usual statistics were computed as well as distance between distributions of comparable ROIs.

Out of 222 identified regions of interest (and their symmetric parts when available), 33 regions were classified as

Table 3: Distance results for the back. Mean, standard deviation and threshold of distance measures for the back, for our control population. Mah., Eucl., Max., Chi, JD, Mal, KS stand for: Mahalanobis, Euclidian, Maximum, Chi-square, Jeffrey-Divergence, Mallows, Kolmogorov-Smirnov distances respectively. CI stands for confidence interval.

Distance results: noisy image							
	Mahn. dist.	Eucl. dist.	Max. dist.	Chi sq. dist.	JD dist.	Mallows dist.	KS dist.
Mean	0.492	0.083	0.032	0.110	0.084	1.092	0.171
Std. dev.	0.152	0.025	0.013	0.057	0.048	0.571	0.103
Threshold (99%CI)	0.796	0.133	0.057	0.223	0.180	2.235	0.377
Distance results: denoised image							
	Mahn. dist.	Eucl. dist.	Max. dist.	Chi sq. dist.	JD dist.	Mallows dist.	KS dist.
Mean	0.565	0.105	0.042	0.142	0.108	1.427	0.205
Std. dev.	0.180	0.039	0.020	0.070	0.060	0.798	0.111
Threshold (99%CI)	0.924	0.184	0.082	0.282	0.229	3.022	0.427

Table 4: Distance results for the legs. Mean, standard deviation and threshold of distance measures for the legs, for our control population. Mah., Eucl., Max., Chi, JD, Mal, KS stand for: Mahalanobis, Euclidian, Maximum, Chi-square, Jeffrey-Divergence, Mallows, Kolmogorov-Smirnov distances respectively. CI stands for confidence interval.

Distance results: noisy image							
	Mahn. dist.	Eucl. dist.	Max. dist.	Chi sq. dist.	JD dist.	Mallows dist.	KS dist.
Mean	0.663	0.099	0.037	0.195	0.101	0.509	0.181
Std. dev.	0.150	0.022	0.013	0.075	0.030	0.247	0.091
Threshold (99%CI)	0.964	0.142	0.063	0.345	0.160	1.004	0.364
Distance results: denoised image							
	Mahn. dist.	Eucl. dist.	Max. dist.	Chi sq. dist.	JD dist.	Mallows dist.	KS dist.
Mean	0.694	0.106	0.038	0.206	0.109	0.528	0.192
Std. dev.	0.166	0.025	0.014	0.083	0.035	0.279	0.100
Threshold (99%CI)	1.027	0.155	0.066	0.373	0.178	1.086	0.392

abnormal using the Manhattan distance, 38 for the Euclidian distance, 27 for the Maximum distance, 36 for the Chi-square distance, 30 for the Jeffrey-Divergence distance, 39 for the Mallows distance and 37 for the Kolmogorov-Smirnov distance. The number of abnormal regions produced by the comparison of mean, variance, skewness, kurtosis, maximum and minimum was only 6. The abnormal regions were not necessarily consistent across the distance measures. The abnormal regions produced by the Euclidian distance, the Chi-squared distance, the Kolmogorov-Smirnov distance and to a lesser extent

those produced by the Manhattan and Jeffrey-divergence distances were generally equivalent. The abnormal regions produced by the three other methods were relatively different. It is interesting to note that the comparison of the mean, variance, skewness, kurtosis, maximum and minimum returned a very low number of abnormal regions as opposed to the distance measures. The threshold for the mean was based on the values of table 2. The other parameters were used to determine whether or not the comparison of the mean was meaningful.

Table 5: Specificity values for the control population. Specificity values of statistical distance measures considered, from a complete isothermal analysis of thermograms of the back and legs for our control population.

Specificity for thermograms of the back							
	Mahn. dist.	Eucl. dist.	Max. dist.	Chi sq. dist.	JD dist.	Mallows dist.	KS dist.
Specificity	0.95	0.95	0.93	0.95	0.94	0.95	0.86
Specificity for thermograms of the legs							
	Mahn. dist.	Eucl. dist.	Max. dist.	Chi sq. dist.	JD dist.	Mallows dist.	KS dist.
Specificity	0.97	0.99	0.98	0.97	0.98	0.98	0.95

Table 6: Sensitivity and specificity values for actual pain patients. Sensitivity and specificity values for thermograms of actual pain patients. Mean ΔT refers to the comparison based on the mean, variance, skewness, kurtosis, maximum and minimum.

	Mahn. dist.	Eucl. dist.	Max. dist.	Chi sq. dist.	JD dist.	Mallows dist.	KS dist.	Mean Temp. diff.
Sensitivity	0.67	0.78	0.33	0.72	0.67	0.55	0.71	0.5
Specificity	0.67	0.83	0.83	0.67	0.67	0.5	0.43	0.67

Furthermore, we noted that the performance of the Mallows distance seemed to be dependent on the size of the sample, where the sample was less than a hundred pixels typically. This is especially problematic when used in conjunction with the isothermal analysis, which may return several small regions. The Jeffrey-Divergence suffered from a similar drawback and as a result we need to consider one of the other distances for very small samples. The Minkowski distances, the chi-square distance and the Kolmogorov-Smirnov distance performed equally well when tested on small samples provided by the isothermal analysis. The five distances correctly identified small abnormal regions for a series of 5 images with apparent abnormalities, while the Jeffrey-Divergence and Mallows distances missed 4 out of 5.

In addition, we compared the outcome from our analysis with the outcome of an assessment by a neurosurgeon at the Moncton Hospital and experienced thermographer, of 24 thermal images of the back and legs. The images were visually assessed by the thermographer since no records or patient history was available for these data. Figure 2 shows an isothermal display at 0.5°C intervals for some of the images submitted to the thermographer. The sensitivity values (Number of true positives with respect to the number of false negatives) and the specificity values (Number of true negatives with respect to the number of

false positives) for the 24 images are summarised in table 6.

These results confirm that the sensitivity of the Euclidian, Chi-square and Kolmogorov-Smirnov distance seem to be significantly higher than that of other distances. Also, the specificity of the Kolmogorov-Smirnov distance is lowest, again confirming the results from the analysis of the control population. Based on the assessment of these 24 images, the best distance appeared to be the Euclidian distance with a fairly high sensitivity and specificity. The Manhattan distance, the Mallows distance and the basic statistics method had relatively low sensitivity. This suggests these methods would dismiss as normal a number of patients with abnormal patterns. Both Manhattan distance and Mallows distance did not perform as well as the results of table 5 suggests.

Conclusions

This paper investigated the computerised assessment of pain through digital infrared thermal imaging. In particular, we looked at the overall digital processing of medical thermograms, from the preprocessing requirements to the decision-support system inputs. In an attempt to increase the objectivity and to facilitate the assessment of pain by medical specialists, we automated as much as possible the assessment of thermal images.

The automation process included the pre-processing and identification of the regions of interest, a statistical analysis and a set of features that can be used by a physician to make a diagnosis.

The pre-processing of thermal images had been overlooked in the medical thermography literature, which focused on the statistical and asymmetry analysis. In order to facilitate the identification of the regions of interest, we proposed applying denoising techniques to thermal images. We modelled the noise by a Poisson distribution with an additive Gaussian white noise component easily removed by classical denoising techniques. Then, we identified the need for a noise removal method that takes the Poisson nature of the noise (and of the signal) into account. This was achieved through a Poisson noise removal technique in the Wavelet domain called PRESS-optimal that was available and used for nuclear medicine images [18]. We showed that the resulting denoised image was more easily processed in the subsequent analysis, without any significant loss of statistical information.

The next problem was to remove the background in order to focus on the areas of interest, again facilitating the following identification of the regions of interest for the statistical comparison. We proposed a simple yet efficient histogram thresholding technique with additional morphological and logical steps, which performed well on both our databases of thermal images.

The identification of the regions of interest consisted of an automated division into a rectangular grid following Uematsu et al. and Goodman et al.'s papers [3,4] for comparison with our control population of normal healthy subjects. This was done by identifying specific reference points on the body, through an analysis of contours and profiles across the image. In addition, an isothermal analysis was performed to determine specific and localised regions, especially the hotter and colder ones that may be indicative of abnormalities.

Finally, we proposed a new statistical tool for the classification of regions of interest as normal or abnormal, based on the computation of several distances between histograms of the ROIs. The results of the quantitative assessment of thermal images from our control population and from pain patients showed that the method based on the Euclidian distance seemed to outperform the other statistical methods considered. Overall, we covered many of the most important steps in the computerised assessment of thermal images and removed the user input, thus providing a highly automated assessment of pain in thermal images.

Authors contributions

CH conceived the algorithms and performed the data analysis and participated in the study design and coordination. MF participated in the study design and coordination and provided the database of images and expertise in infrared imaging. All authors read and approved the final manuscript.

Acknowledgements

The authors thank the National Science and Engineering Research Council of Canada (NSERC) for funding this study. The authors also thank The Terry Fox Foundation in Fredericton, NB for funding the 1984 study and Dr. G. Quartey, at Moncton Hospital, NB, for his expertise in pain assessment, patient base, and clinical results.

References

- Houdas Y, Ring EFJ, Eds: *Human Body Temperature: its Measurement and Regulation* New York, NY: Plenum Press; 1982.
- Aubry-Frize M, Quartey GRC, Evans H, LaPalme D: **The Thermographic Detection of Pain.** In *Proceedings of the 3rd Canadian Clinical Engineering Conference: September 1981; Saskatoon, Canada 1981*:82-83.
- Goodman PH, Murphy MG, Siltanen GL, Kelley MP, Rucker L: **Normal Temperature Asymmetry of the Back and Extremities by Computer-Assisted Infrared Imaging.** *Thermology* 1986, **1**:195-202.
- Uematsu S, Edwin DH, Jankel WR, Kozikowski J, Trattner M: **Quantification of Thermal Asymmetry, Part I: Normal Values and Reproducibility.** *Journal of Neurosurgery* 1988, **69**:552-555.
- Hooshmand H, Hashmi M, Phillips EM: **Infrared Thermal Imaging as a Tool in Pain Management – An 11 Year Study. Part I of II.** *Thermology International* 2001, **11**(2):53-65.
- Hooshmand H, Hashmi M, Phillips EM: **Infrared Thermal Imaging as a Tool in Pain Management – An 11 Year Study. Part II: Clinical Applications.** *Thermology International* 2001, **11**(2):117-129.
- Ellis WV, Morris JM, Swartz AA: **Screening Thermography of Chronic Back Pain Patients with Negative Neuromusculoskeletal Findings.** *Thermology* 1989, **3**(2):.
- Goldberg GS: **Infrared Imaging and Magnetic Resonance Imaging Correlated in 35 Cases.** *Thermology* 1986, **1**(4):.
- Hubbard JE, Hoyt C: **Pain Evaluation in 805 Studies by Infrared Imaging.** *Thermology* 1986, **1**(3):161-166.
- Kim YS, Cho YE: **Correlation of Pain Severity with Thermography.** In *Proceedings of the 17th Annual International Conference of the IEEE Engineering in Medicine and Biology Society and 21st Canadian Medical and Biological Engineering Conference 1995*, **2**:1699-1700.
- LeRoy PL: **Update in Interventional Digitalized Infra-Red thermal Imaging in Pain Management.** In *Proceedings of the 17th Annual International Conference of the IEEE Engineering in Medicine and Biology Society and 21st Canadian Medical and Biological Engineering Conference 1995*, **2**:1707-1708.
- Thomas D, Siahamsi G, Marion M, Boyle C: **Computerised Infrared Thermography and Isotropic Bone Scanning in Tennis Elbow.** *Ann Rheum Dis* 1992, **51**:103-107.
- Anbar M: **Objective Assessment of Clinical Computerized Thermal Images.** In *Medical Imaging V: Image Processing, SPIE-The International Society for Optical Engineers 1991*, **1445**:479-484.
- Collins AJ, Ring EFJ, Cosh JA, Bacon PA: **Quantitation of Thermography in Arthritis using Multiisothermal Analysis.** *Ann Rheum Dis* 1974, **33**:113-115.
- Pochaczewsky R: **Technical Guidelines, Edition 2.** *Thermology* 1986, **2**(2):108-112.
- Wolfe WL, Ed: *Handbook of military infrared technology* Washington DC: Office of Naval Research, Department of Navy; 1965.
- Keyes RJ, Ed: *Optical and infrared detectors. No. 19 in Topics in applied physics* Springer-Verlag; 1977.
- Nowak RD, Baraniuk RG: **Wavelet-Domain Filtering for Photon Imaging Systems.** *IEEE Transactions on Image Processing* 1999, **8**(5):666-678.

19. Tsai DM: **A Fast Thresholding Selection Procedure for Multimodal and Unimodal Histograms.** *Pattern Recognition Letters* 1995, **16(6)**:653-666.
20. Montoro J, Anbar M: **New Modes of Data Handling in Computerized Thermography.** In *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society* 1988:845-847.
21. Mabuchi K, Chinzei T, Fujimasa I, Haeno S, Abe Y, Yonezawa T: **An Image-Processing Program for the Evaluation of Asymmetrical Thermal Distributions.** In *Proceedings of the 19th Annual International Conference of the IEEE Engineering in Medicine and Biology Society; Chicago, IL* 1997:725-728.
22. Puzicha J, Buhmann JM, Rubner Y, Tomasi C: **Empirical Evaluation of Dissimilarity Measures for Color and Texture.** In *Proceedings the IEEE International Conference on Computer Vision (ICCV1999)* 1999:1165-1173.
23. Levina E, Bickel P: **The Earth Mover's Distance is the Mallows Distance: Some Insights from Statistics.** In *Proceedings of the Eight International Conference on Computer Vision (ICCV2001)* 2001, **2**:251-256.
24. Vavilov VP, Vavilova EV, Popov DN: **Statistical Analysis of the Human Body Temperature Asymmetry as the Basis for Detecting Pathologies by Means of IR Thermography.** In *Thermosense XXIII Volume 3061*. Edited by: Rozlosnik AE, Dinwiddie RB. SPIE-The International Society for Optical Engineers; 2001:482-491.
25. **Wavelab 802: Matlab wavelet toolbox from the Department of Statistics, Stanford University** [<http://www-stat.stanford.edu/~wavelab/>]

Publish with **BioMed Central** and every scientist can read your work free of charge

"BioMed Central will be the most significant development for disseminating the results of biomedical research in our lifetime."

Sir Paul Nurse, Cancer Research UK

Your research papers will be:

- available free of charge to the entire biomedical community
- peer reviewed and published immediately upon acceptance
- cited in PubMed and archived on PubMed Central
- yours — you keep the copyright

Submit your manuscript here:
http://www.biomedcentral.com/info/publishing_adv.asp

