Area Assessment of Psoriatic Lesions based on Variable Thresholding and Subimage Classification

Juha Röning, Robrect Jacques, Jukka Kontinen
Machine Vision and Media Processing Group
Infotech Oulu and Department of Electrical Engineering, University of Oulu, B.O. Box 4500,
FIN-90401 Oulu, Finland
email: jjr@ee.oulu.fi

Abstract

The assessment of psoriatic lesions on human skin by human observers is not objective or consistent, and the estimates typically involve too considerable variation. To facilitate the evaluation of psoriatic patients, an objective assessment of involved skin surface is needed.

In this work, a segmentation scheme was developed to divide the skin surface into non-involved and involved regions. Based on the bi-modality of small subimages, a variable thresholding scheme was generated. To take both local and global impacts into account, an iterative method with variable window size was developed. Multiple passes improve the classification of pixels. A number of post-processing steps are applied. These steps operate on a small 3x3 neighbourhood of the centre pixel. The number of false segments is reduced, defects are removed from the image and smoother transition is obtained.

1 Introduction

In the field of dermatology, digital imaging applications were rare before 1985, but progressive growth has been going on since then. The benefits of digital image processing in the field of dermatology are obvious: the method is objective, quantifiable and interactive, and allows easy permanent storage [2].The number of papers and reports addressing digital imaging in the field of dermatology has increased exponentially [1]. One of the topics is the psoriasis.

Psoriasis is a non-contagious skin disorder of probably genetic origin [3]. The most common type is *plaque psoriasis*, which is manifested as swollen skin lesions covered by a silvery white scales. Although the origin of the disease is not quite clear, it appears that psoriatic lesions are developed because local skin cells behave as if a skin wound were present. Because of this accelerated cell production is found to replace and repair the (non-present) skin wound. Other variations are, however, also possible: pus-like blisters as in *pustular psoriasis*, severe sloughing of the skin, as in *erythrodermic psoriasis*, or smooth inflamed legions, as in *inverse psoriasis*. The severity of skin lesions can also vary from one patient to another.

To facilitate the development of curative treatments and medicines, a *psoriasis area and severity index* (PASI) has been developed [4][5]. It is a way to evaluate the area and severity of psoriatic lesions in a patient. If an objective indicator of disability is available, more research can be done on the progress of the severity of psoriatic lesions under influence of different treatments. The body is divided into four main areas: the head, the upper extremities, the trunk, and the lower extremities. The areas represent, respectively, 10, 20, 30 and 40% of total body surface. The parameters provide an evaluation of the *area* of involved skin (*A*) and the severity of the lesion by three target symptoms, namely *erythema* (*E*), *infiltration* (*I*) and *desquamation* (*D*).

The area of involvement has to be calculated. In analogy with the assessment of burn lesions on skin surface, the "rule of nines" method was used [6]. Studies [6],[7] have shown, however, that this method is not objective. The estimates involve wide inter-observer variation. The lack of objective indicators of disease severity make consistent observations of the evolution of the patients' condition inaccurate or even impossible. There is a need for accurate area calculation for pharmaceutical development and evaluation of disability.

In this paper, we will first give a brief overview of the system by which a doctor can make an objective assessment of the area of involvement in psoriasis, and will then concentrate on the sub-problem of segmentation.

2 Overview of the system

Our goal is to develop an objective method of estimating involved surface area in patients with psoriasis.The current image processing method can be divided into four main phases. The image is pre-processed to remove the information unnecessary for segmentation and highlighting and to clarify the necessary parts. Next, the image is segmented by using colour information to extract unhealthy skin to a maximum extent. Classification of the segments affected by psoriasis is then carried out manually. Finally, the white skin surface area of the whole body is determined, to calculate the percentage area of psoriasis involvement. A schematic overview of the system is given in Figure 1.

The assessment is based on segmentation with a previously developed hierarchical connected components analysis [13, 16, 17]. One important aspect of the system is preprocessing of the image, including extraction of Ohta I2 [18], the normalized green (g) colour feature, SNN (Symmetric Nearest Neighbor) and median filterings. The purpose of these filterings is to enable accurate segmentation without too many false segments.

The segmentation method used in this work is basically a

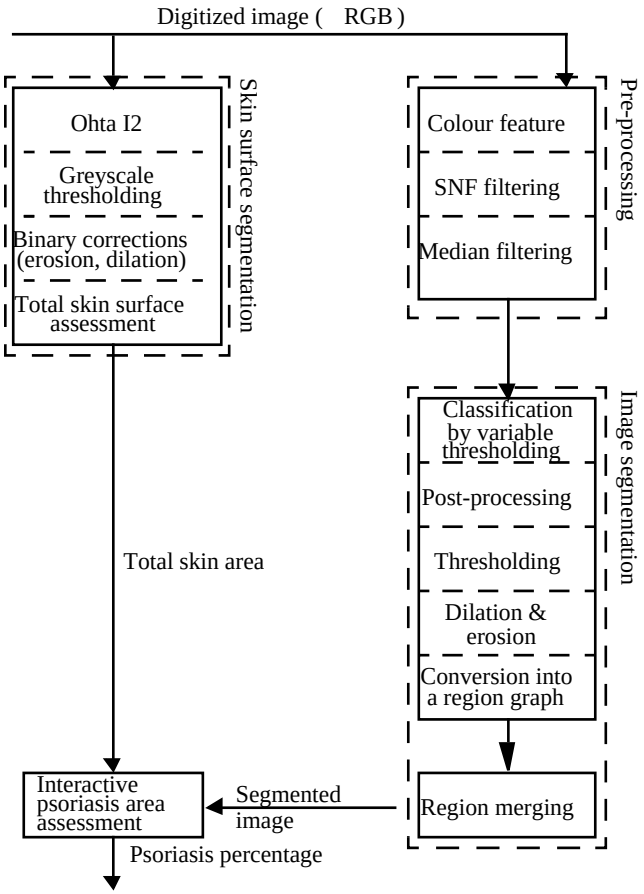

Figure 1 Schematic overview of the system.

variable thresholding method [9-12]. A classification scheme with overlapping classes was developed. Different passes with variable window size are performed to improve classification on a global and local scale. The resulting image is thresholded to obtain the different objects (involved skin and non-involved skin regions).

3 Segmentation

A schematic overview of segmentation is given in Figure 2. Based on the typically bi-modal histogram of skin surface with psoriatic lesions, a new method to create the initial 'region graph' is presented in section 3.1. The image is then divided into several small images (e.g. 50x50 pixels) where the local variable thresholds are determined. Compared to a globally determined threshold, better independence of the irregular lighting on skin surface is gained. The segmentation process is here illustrated with images of one patient. The original colour image, body extraction and enhanced greyscale image are shown in Figure 3.

3.1 Create region graph

The 'create region graph' consists out of four parts: classification of the image by variable thresholding, post-processing of the 'labelled image', thresholding of the labelled image, correcting the labelled image with dilation and erosion operators, and conversion into a region graph format. In the following sections, each part will be discussed.

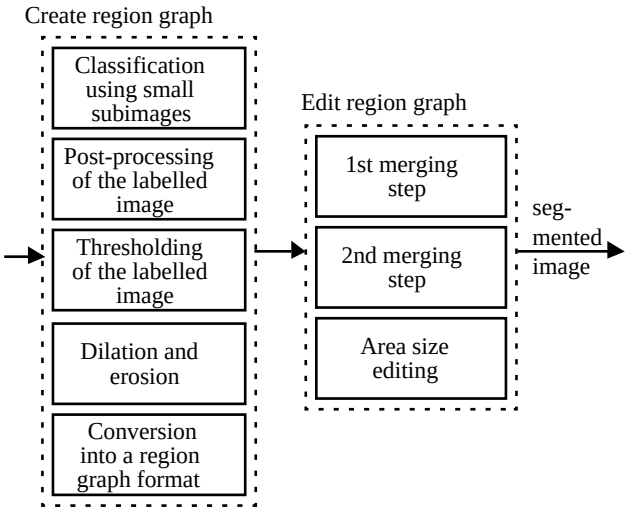

Figure 2 A schematic overview of the new segmentation system.

3.2 Classification using small subimages

Principle

The pixel greyscale level belongs to one of the two Gaussian curves, either that of 'normal skin' or that of 'involved skin'. This assumption makes it possible to classify the pixels into two different classes. The histogram is used to determine a local variable threshold. Pixels with a greyscale level above this threshold belong to one class, while the ones with a greyscale level below the threshold belong to the other class. After thresholding, a binary segmented image is obtained.

In most variable thresholding schemes [9-11] the valley between the two Gaussian curves is chosen as the threshold. This threshold splits the sub-population into two clearly distinctive classes. Some other methods use the curvature of the distributed modal histogram to determine the thresholds [12]. In this work, a different approach was chosen. Altogether three thresholds were used: the mean m of the histogram and the two means m_1 and m_2 of the Gaussian curves.

The two curves overlap. This is inevitable because psoriatic lesions and blend smoothly into normal skin. In the overlapping area, no definite decision can be made about what class a given pixel belongs to. Below the mean m_1 and above the mean m_2 of the two Gaussian curves, each pixel can be classified to a certain class with a fair amount of certainty. The chance of a pixel with a greyscale value larger than m_2 belonging to the other class (curve A) is very small (the bell shape of Gaussian distribution). Pixels with greys-

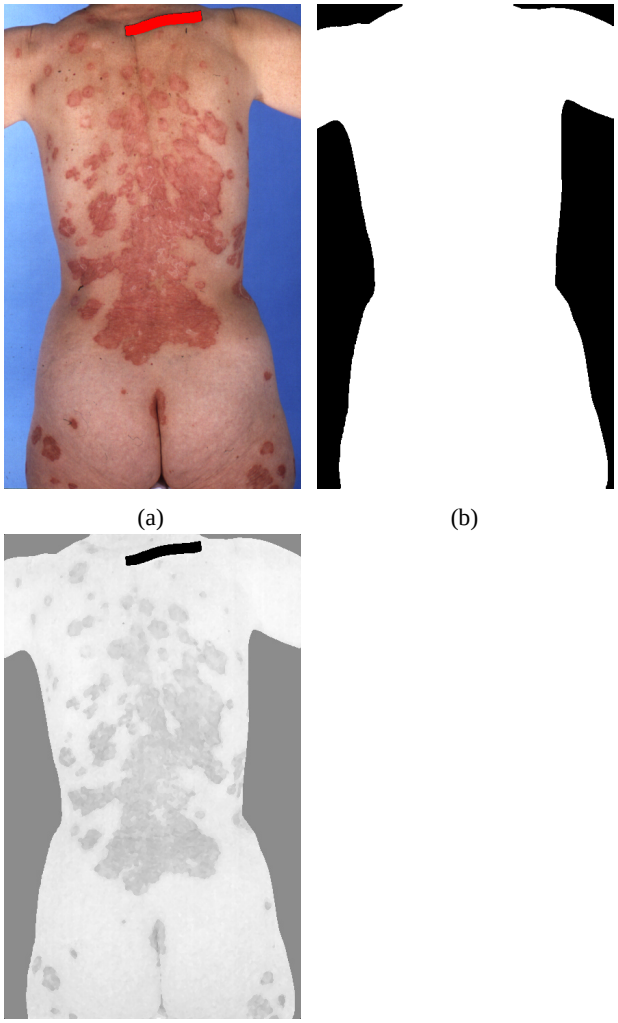

(a) (b)

(c)

Figure 3 (a) Original colour image, (b) body extraction, and (c) pre-processed image.

cale values below m_1 are very likely to belong to curve *A* and pixels with greyscale values above m_2 to curve *B*. Therefore, the two means of the Gaussian curves are used as thresholds. This results in three distinct classes: 'involved skin' (below m_1 and thus belonging to curve *A*), 'normal skin' (above m_2 and thus belonging to curve *B*) and 'unknown' (between m_1 and m_2).

Some information is omitted if the classification is only done into three classes. It is better to specify the 'unknown' class more precisely. Consider the mean *m* of the whole distribution. This mean is between m_1 and m_2. Pixels with a greyscale level lower than *m*, but higher than m_1, are more likely to belong to curve *A* than to curve *B*, and are thus labelled as 'likely to belong to involved skin'. Analogously, pixels with a greyscale level higher than *m*, but lower than m_2, are labelled as 'likely to belong to normal skin'. In this way, a more precise distinction is made between the classes.

In this scheme, the resulting image is not a binary one, but has four classes. Each of the four classes can be given a corresponding label or greyscale level for visual presentation. 'Normal skin' is given the value 20 (white), 'probably normal skin' 15 (light grey), 'probably involved skin' 5 (dark grey) and 'involved skin' 0 (black). A number of classified areas are shown in Figure 4. For the construction of these images, a simplified version of the following algorithm was used.

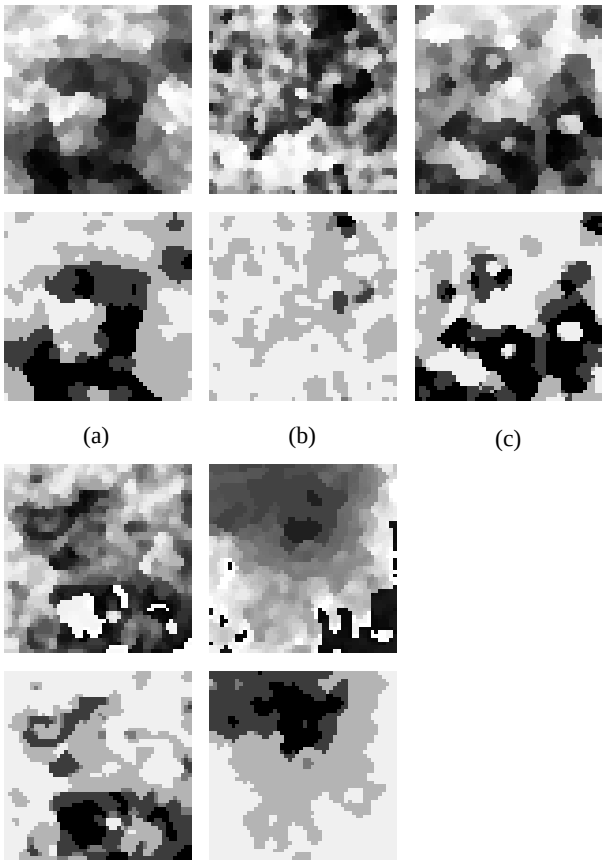

(a) (b) (c)

(d) (e)

Figure 4 A number of small (50x50 pixels) image areas (top row) with their classification (bottom row).

Algorithm

It can be seen from the images that the overall segmentation of the subimages is good. The simple principle could result in a very powerful segmentation algorithm. Some problems, including the implementation itself, should, however, still be solved. The following remarks elaborate the algorithm some more.

The algorithm has to find out if the distribution is Gaussian or not. If a bi-modal histogram is found, the classification can be done by locating the thresholds *m*, m_1 and m_2 and thresholding the image. A Gaussian distribution, however, would indicate that the whole area is normal skin or in-

volved skin. The histogram itself contains no information as to which one of the classes it belongs to. One possible method would be to use the thresholds of the neighbouring subimages. In [11], for instance, the threshold (the valley) of each subimage is calculated. The threshold of the subimage with a Gaussian distribution is calculated by a weighted average of the threshold of the eight neighbouring subimages. After that, the thresholds are smoothed and assigned to the central pixel. For each individual pixel, the threshold is then calculated as an interpolation of the thresholds in the four nearest centre pixels. A continuous variable threshold can be formed this way. The complexity of this scheme is quite high if the three thresholds have to be dealt with.

In this work, a different approach is used to handle Gaussian distributed histograms. Let *m* be the mean of the whole distribution. The pixels with a greyscale value lower than *m* are then classified as belonging to 'probably involved skin', while the pixels with a greyscale value higher than *m* are said to belong to 'probably normal skin'. No assignments are made to 'normal skin' or 'involved skin'.

The idea behind this is as follows. If an area really belongs to one category only, e.g. 'normal skin', then the distribution of assigned labels to the pixels is statistical. In other words, neighbouring pixels are given different labels without certain labels clustering together. However, if the area has a (small) psoriasis lesion inside it, the lower labels will cluster together around the psoriasis lesions.

An example of this is shown in Figure 4.b. In the image, the second Gaussian curve (skin lesion) is so small as to be hardly noticeable. Because the algorithm does not detect the second curve, it categorises the subimage as uniform 'normal skin' or 'skin lesion'. Small areas of involved skin are, however, still noticeable in the image because of the clustering of identical labels.

Instead of dividing the image into smaller subimages once, it is also possible to do the classification through a number of passes. Every pass can have a different window size (e.g. 50x50, 100x100 or 500x500). In this way, two things can be combined:

Larger windows give a more global overview of the image, which allows detection of large regions of skin or psoriatic lesions. The details and smaller regions, however, may be overlooked because of local quality differences in the image or irregular lighting.

Smaller windows may detect even very small psoriatic areas. If the windows are too small, however, many areas are classified as 'unknown', and no information about the image area can therefore be obtained.

Four discrete classification levels are not sufficient to allow good classification in the different passes. Pixels may change the class in successive passes. There should be a way retain information of the class to which the pixel belonged in the previous pass.

In theory, it is possible to calculate the possibility of each pixel to belong to one of the two Gaussian curves. Successive passes would then result in a combined probability that can be calculated with the 'rule of Bayes'. A continuous changing label presenting the probability level could then be assigned instead of current four discrete labels. This would lead too far, however, which is why different approach was used here.

To save the earlier classification, a gradual distinction between normal skin and involved skin is needed. In this work, every pixel was assigned a label from 0 (involved skin) through 10 (unknown) to 20 (normal skin). The initial value for every pixel in the image is naturally 10 (unknown).

Every pass changes the label of the pixel as follows:

if it belongs to 'involved skin', 3 is subtracted from the present label.

if it belongs to 'probably involved skin', 1 is subtracted from the present label.

if it belongs to 'probably normal skin', 1 is added to the present label.

if it belongs to 'normal skin', 3 is added to the present label.

Like this, pixels that really belong to 'involved skin' and thus have 3 subtracted from their label in every pass, will ultimately have the label 0 and be classified as 'involved skin' after a certain number of passes. The same is true of 'normal skin' pixels.

In the case of 'probably involved skin' 1 is subtracted. Pixels on the edges of lesions (in the smooth transition between skin and psoriasis) are classified as 'involved skin' after a number of passes. This is agreement with the 'rule of Bayes': if a pixel has a large enough probability to belong to a certain class in every pass, the probability that it really belongs to that class is high.

Pixels belonging to 'normal skin' in one pass and to 'involved skin' in another are classified as 'unknown'. Examples of such pixels are the ones that are classified when the histogram is pure by Gaussian. The statistical distribution of labels in these regions allows pixels to switch between involved and normal skin. The expected value of the label is 10 ('unknown').

The windows are moved through the image. A skin lesion that lies on the edge of a window, is divided into two, one part in each of the windows. It is then possible that, because of different threshold values, the skin lesion is classified differently in the two subimages. For instance, the lesion could be considered 'probably normal skin' in one subimage and 'probably involved skin' in another.

To avoid this, the different windows should overlap. Skin lesions that lie on the edges are centred in the next pass, when the windows are shifted. The position of the windows in the different passes is shown in Figure 5.

The determination of the means of the two Gaussian distributions is not straightforward. It is possible to find an optimal solution (fitting, [15]), but this would require time-

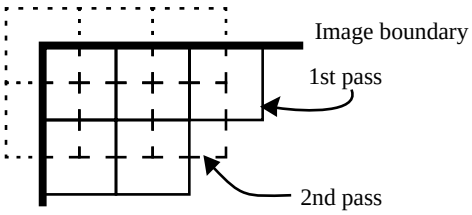

Figure 5 Windows overlap each other in successive passes to avoid wrong classification of objects that lie on the edges.

consuming calculations. The thresholds must be calculated separately for multiple passes. In this work, an approximation is made that is quicker to calculate.

The mean of a Gaussian curve corresponds to the maximum of the distribution. Therefore, the search for the two means is basically a search for the two maxima.

The output of the algorithm is a '*labelled image*'. This structure is an image that contains for every pixel the assigned label instead of the pixel greyscale level. The 'labelled image' obtained by applying the algorithm is shown in Figure 6.a.

3.3 Post-processing of labelled image

The classifying algorithm with thresholding works quite well for a large number of images. An example is given in Figure 6.a. The overall classification is good compared with Figure 3.

The result is not, however, optimal yet. Some pixels adjacent to a skin lesion are labelled differently, although they clearly belong to the lesion. This is because the lesions blend into normal skin, and a gradual evolution of the label can be expected if one proceeds from the centre of the lesion to normal skin. Other small pixel groups labelled 'involved skin' lie in normal skin areas, but clearly belong to the normal skin. This may be due to the small window size. Local deviations are treated as skin lesions and classified accordingly.

These deficiencies can be eliminated or at least reduced by applying some post-processing. In the previous derivation of the algorithm, no attention was given to the neighbourhood in which the pixel occurs. Smaller windows focus on a small neighbourhood of the pixel, but still take 10x10=100 pixels into account. The immediate or adjacent pixels play no greater role in the decision than the pixels located further away.

It is logical for a pixel to belong to a skin lesion if the adjacent pixels in a 3x3 neighbourhood belong to a skin lesion, too. The same is true of pixels in a normal skin neighbourhood. The following post-processing step is thus proposed: if a pixel is in a 'mainly normal skin' neighbourhood, the label of the pixel is increased, and if it is in a 'mainly involved skin' neighbourhood the label is decreased.

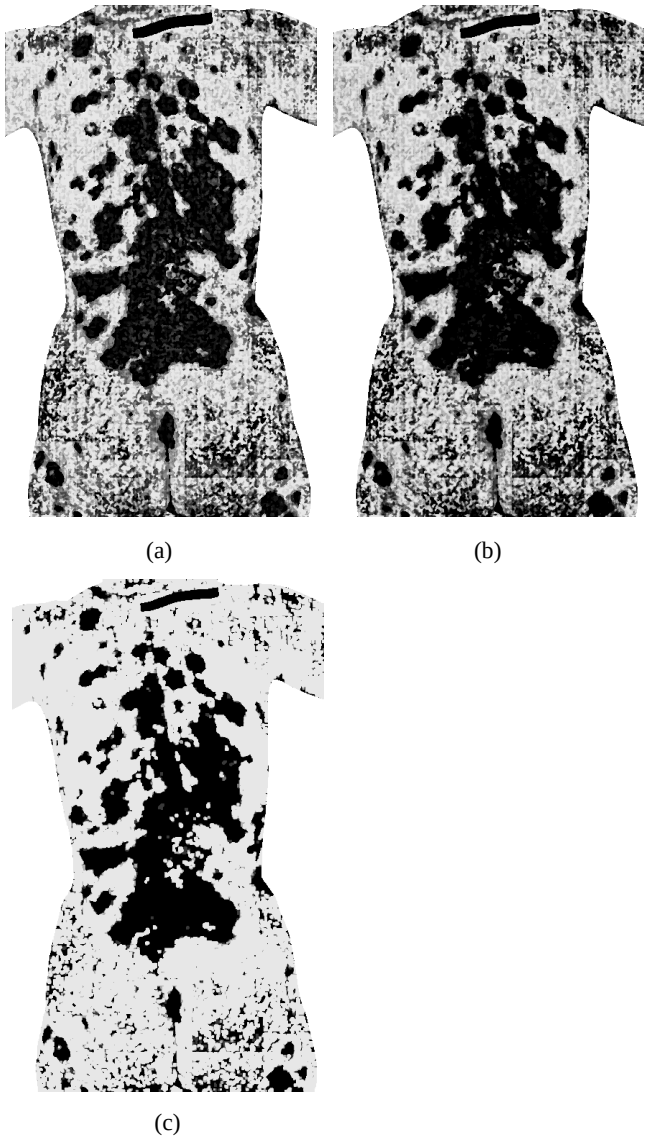

Figure 6 'Labelled image' of the example image (a) after classification, (b) after post-processing, and (c) after thresholding.

The classification is now slightly better. The isolated pixels are mostly upgraded to the right label. Isolated pixels are ones that lie in a uniform area, but are labelled differently. The border pixels around skin lesions are also labelled closer to the area they belong to. An example is given in Figure 6.b.

3.3.1 Thresholding of the labelled image

In the algorithm described above, the labels range from 0 to 20. The idea, however, was to segment the image into two distinct classes: involved and not involved skin. The labels should be thresholded so that labels below a certain value are classified as 'involved skin', while labels above another value are classified as 'normal skin'.

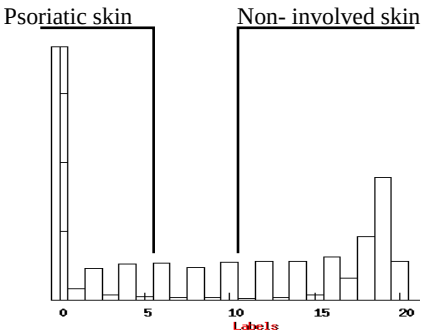

Figure 7 Typical histogram of the labels after classification.

Many images have label histograms of the kind shown in Figure 7. Most pixels are classified as psoriatic (0) or non-involved skin (20), and only a few pixels in between are not certain yet. These pixels are divided into three groups, using the thresholds 6 and 10. The pixels between 6 and 10 are classified into 5 unknown classes. Some methods have been tentatively applied to find a threshold that would classify all the pixels correctly in one of the two classes. Both fixed thresholds and thresholds calculated on the basis of the label histogram have been applied. For most images, however, these methods were not successful. Shadows are mostly also classified as psoriatic involved skin. The threshold would then merge both areas as one object. This should be avoided. Therefore, some uncertainty remains in the classification. At the 'edit region graph' step, this uncertainty is removed by merging regions based on global pixel greyscale information. A 'Labelled image' of the example image after thresholding is shown in Figure 6.c.

3.3.2 Dilation of the thresholded labelled image

Another post-processing step is then applied to the labelled image. The previous step improved the classification, but some psoriatic regions are connected with very thin lines of low classification. A 'region graph' made connects regions that are located far from each other and are in no way connected. This problem is present in Figure 8.a. In this image, the connection of different regions is due to the shadows of the ridges on the back of the hand. The shadows between the fingers or strange objects in the image, such as underpants, will also cause this problem. Because of the assumption of bi-modality, subimages classify these pixels as 'involved skin', because it has clearly a lower greyscale value than the surrounding skin in every pass and with every window size. Figure 8.b shows the initial 'region graph' constructed out of this 'labelled image'. The two distinct psoriatic lesions are connected to each other by a thin line. At the subsequent steps, this connection cannot be undone. The result is a wrong region (Figure 8.c).

To remove the thin lines, a dilation and erosion step is applied. Dilation expands the higher label (i.e. non-involved skin surface) and removes very thin structures, e.g. the small shadow line. Erosion then expands the bigger classified psoriatic regions back to their original size. The overall classification remains the same, and only the thin structures, such as shadows and really small psoriatic lesions, are removed. Figure 8.d shows the same image. The two regions are no longer connected to each other. The two psoriatic lesions are individually assessable.

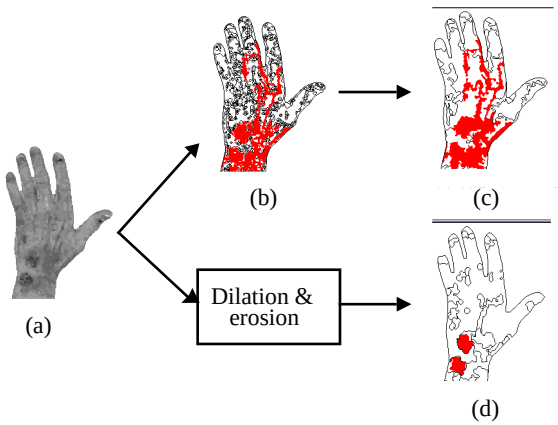

Figure 8 (a) Enhanced image, (b) a thin line connects different regions, (c) in the end result this could not be undone. (d) After a dilation and erosion step, the result is much better.

3.3.3 Conversion into the 'region graph' format

Next, the classified image is converted into a 'region graph' for better post-processing. The classification information is preserved in the region graph format. The initial regions are made of the labelled image, so that the pixels belonging to one region have the same classification label. The information about pixels is also restored. It is possible to access the greyscale level of each individual pixel in the region, but it is also possible to calculate such parameters as the mean intensity of a region, the variation of greyscale in a region, and the average border contrast with an adjacent region. This information is averaged out over a large number of pixels. The influence of local deviations within a region is reduced.

3.4 Region merging

In this work, the initial region graph is obtained by variable thresholding. Unfortunately, it also gives a lot of false segments, e.g. shadows in areas without psoriatic lesions. Region merging is applied to remove these false segments. A merging step requires two parameters: the merging method and the merging threshold.

3.4.1 Methods

Methods are the cost functions on which the two regions are evaluated. If this cost function is smaller than a specific

threshold, the two adjacent regions are merged into one bigger region.

Merge weaker

This method is the original method also used in the HCC/SNF procedure [13].

> In 'merge weaker', two adjacent regions are merged if the average boundary contrast is lower than a certain threshold.

The average boundary contrast is defined as the sum of the difference in greyscale between pixels that lie on both sides of the common boundary divided by the number of pixels on that boundary. The idea behind this method is that if the contrast between two boundaries is small, the two regions are likely to belong to the same object. After all, two regions belonging to the same object blend smoothly in one another.

On body surfaces, this may cause a problem because shadow regions and normal skin regions also blend smoothly into one another. The boundary of skin lesions is also mostly not well-defined and blends into normal skin. If this method is used with a threshold that is too high, it is inevitable that wrong regions may merge.

Merge similar

Pixels that belong to the same classification (normal skin or involved skin) have greyscale values that are close to each other (the colour feature is derived based on this assumption, see section 3.3). It is equally likely that regions belonging to the same object or classification have similar average greyscale values.

> In 'merge similar', two adjacent regions are merged if the difference in average intensity is lower than a certain threshold.

The average intensity of a region is defined as the sum of all the greyscale values of the pixels belonging to the region divided by the number of pixels in that region.

The comparison between two adjacent regions does not take place on a local scale (the boundary between the two regions), but applies to the whole region. Local deviations on region boundaries have less influence.

Merge weaker and similar

One obvious generalisation is to merge the regions that have sufficient boundary contrast as well as similar greyscale intensities. The two thresholds can be set independently of each other.

> In 'merge weaker and similar', two adjacent regions are merged if the average boundary contrast is lower than a certain threshold AND the difference in average intensity is lower than another threshold.

3.4.2 Calculation of thresholds

In general, the calculation of thresholds is no longer done on the basis of individual pixels, but rather on the basis of the different regions and the chosen method. A histogram is calculated and the mean, median or maximum of this histogram is chosen as the threshold.

Calculation of histograms

Two histograms are calculated. The histogram of average boundary contrast between adjacent regions is used to calculate of 'merge weaker' thresholds. The histogram of average intensity difference between adjacent regions is used to calculate of 'merge similar' thresholds. These two histograms are calculated by looping through all the regions and their adjacent regions, calculating the specific cost-function and incrementing the entry in the histogram.

Calculation of thresholds

Of these two histograms, three values are calculated: the mean, the value corresponding to the maximum, and the median.

These three values are taken because they are significant in view of probability. For example, when the mean value of the histogram is used it is likely for two regions to belong together if the difference in the cost function used (difference in average intensity or average boundary contrast) is smaller than the 'normal' difference between regions. The same reasoning can be done for the maximum (most regions are so close to each other) and the median (50% of the regions are so close to each other).

3.4.3 Description of merging steps

In this work, three merging steps are applied. The initial regions are well defined in most cases. They are, however, still surrounded by thin and small regions. An example of this is shown in Figure 9.

Two 'merge similar and weaker' steps are performed on the 'region graph'. The thresholds are determined in a different way: in the first merging step, the median of the similar histogram and the mean of the weaker histogram are taken as the two thresholds. In the second merging step, the maximum and the median are taken, respectively.

After the conversion of the labelled image into a region graph many small regions are present round the psoriatic lesions. In the first merging step a relatively higher threshold is taken, which means that most of these regions merge with one of their adjacent regions. The second step is not needed for all the images, but it gives an improvement for some. Lower thresholds must be chosen (relatively). The regions are already well defined. Too high thresholds would merge the psoriatic regions with normal skin. After these two merging steps, there are still some very small false segments in the region graph.These regions are sometimes on borders of psoriatic regions, and sometimes they are small false segments of skin surface. The regions should be removed to make assessment of the psoriasis area easier. Regions smaller than 50 pixels in area are removed.

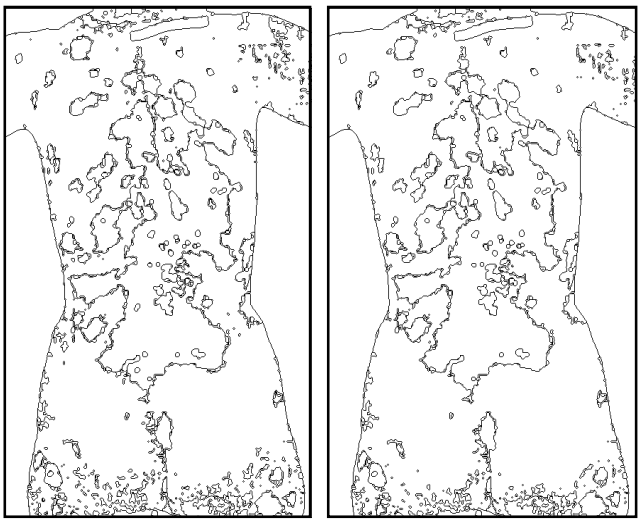

(a) (b)

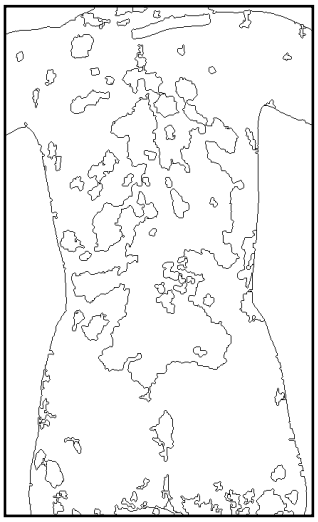

(c)

Figure 9 The gradual improvement of the region graph after (a, b) merging of similar and weaker regions and (c) smaller regions.

The final psoriasis area assessment is done with the interactive tool described in more detail previously [14], and it gives the result shown in Figure 10. The reliability and accuracy of the assessment has been tested with reference analysis by colouring manually the involved areas in the images with the aid of projected slides [15]. The segmentation algorithm works well for a large number of images without adjustment of the system parameters. Reference patients were assessed with the system, and the measured percentages compared with accurate reference values. The error was in all cases smaller than 4%. Compared to the interval of classification in the PASI formula, this error is small. Misclassification occurs, however, in borderline cases. This caused misclassifiaction in less than 10% of the assessed patients.

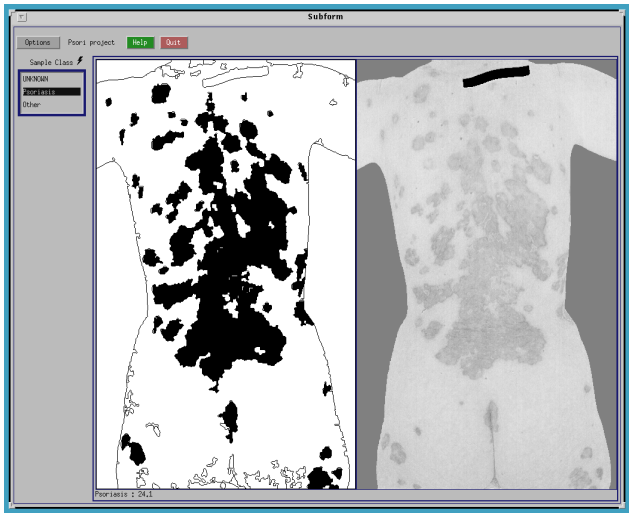

Figure 10 Assessment of the psoriasis area in the example patient.

4 Conclusions

A segmentation algorithm was derived, which is based on the typical bi-modality of skin with lesions. A variable thresholding scheme was used to segment the skin surface. Multiple passes with different window sizes iteratively improved the classification. The defects in the classification were removed with different post-processing steps operating in 3x3 neighbourhoods. Dilation and erosion steps were added, to remove thin structures from the image that caused mismerger of different objects on the skin surface. False segments were removed by applying region merging techniques that use global information of the region, such as average boundary contrast and average intensity.

The system is not accurate in all cases. Some segmentation faults remain in the system. Small psoriatic lesions are removed from the image in the same way as thin structures. The dilation and erosion could be specified differently to minimize this problem.

5 Acknowledgments

The authors gratefully acknowledge the expert assistance of Aarne Oikarinen and Leena Savolainen in the Department of Dermatology and Venereology at Oulu University Hospital. The ERASMUS exchange program is acknowledged for support.

6 References

[1] What dermatologists should know about digital imaging , Douglas A. P erednia, J ournal of the American Academy of Dermatology 25:89-108 (1991).

[2] Editorial: Digital imaging in dermatology , William V . Stoec ker and Randy H. Moss , Computerized medical imaging and graphics , vol. 16, no 3, 145:150 (1992).

[3] http://www .psoriasis .org/whatis .html, National Psoriasis F oundation, 6600 SW 92nd A venue , Suite 300, Portland, OR 97223, USA, tel. (503) 244-7404, fax (503) 245-0626, email 76135.2746@compuserve .com.

[4] Severe psoriasis - oral therapy with a new retinoid, T. F redriksson and U . P ettersson, Dermatologica 157:238-244 (1978).

[5] On the limitations of the psoriasis are and severity index (P ASI), P .C.M van de K erkhof , British J ournal of Dermatology 126:205 (1991).

[6] Measurement of involved surface area in patients with psoriasis , B. Ramsa y and C . M. La wrence , British J ournal of Dermatology 124:565-570 (1991).

[7] Assessment of area of involvement in skin disease: a study using sc hematic figure outlines , S. Tiling-Grosse and J . Rees , British J ournal of Dermatology 128:69-74 (1993).

[8] Color quantification through image processing in dermatology , M. Herbin, A. V enot, J . Devaux and C . Piette , IEEE Transactions on medical imaging 9, 262-269 (1990).

[9] Image Segmentation T echniques , Robert M. Haralick and Linda G . Shapiro , Computer vision, graphics and image processing 29, 100-132 (1985).

[10] Digital image processing , William K. Pratt, W iley-Interscience (1978).

[11] Some experiments on variable thresholding , Yasuo Nakaga wa and Azriel Rosenfield, P attern recognition, vol. 11, 191:204 (1979).

[12] An amplitude segmentation method based on the distribution function of an image , S. Boukharouba, J.M. Rebordao and P .L. W endel, Computer vision, graphics and image processing 29, 47-59 (1985).

[13] Westman T ., Harwood D ., Laitinen T . & Pietikäinen M. (1990) Color Segmentation by Hierarc hical Connected Components Analysis W ith Image Enhancement by Symmetric Neighborhood F ilters . In: Proceedings of The 10th International Conference on Computer V ision and P attern Recognition Systems and Applications , J une 16-21,Atlantic City , New J ersey, pp . 796-802, Los Alamitos .

[14] J ukka K ontinen, Leena Sa volainen, Aarne Oikarinen and J uha Röning: Computer Image Analysing System for Psoriasis Area Assessment, IASTED International Conference Signal and Image Processing (SIP-97), December 4-6, 1997, New Orleans , Louisiana.

[15] Sa volainen, L., K ontinen, J ., Röning , J , Oikarinen, A.: Application of Mac hine V ision to Assess Involved Surface in P atients with Psoriasis , British J ournal of Dermatology , 1997; 137: 395-400.

[16] D. Harwood, M. Subbarao , H. Hakalahti and L. Da v-is, A New Class of Edge-preserving Smoothing F ilters , Pattern Recognition Letter s 6, 1987, 155-162.

[17] Pietikäinen M. Harwood, Segmentation of color images using edge-preserving filters . In: Advances in Image Processing and P attern Recognition (Capellini V. Marconi R., eds). North Holland: Elsevier . 1986:94-9.

[18] Y. Ohta, T . Kanade and T . Sakai, Color Information for Region Segmentation, Computer Graphics and Image Processing 13, 1980, 222-241.