An overview of the Polartechnics SolarScan melanoma diagnosis algorithms

Hugues Talbot and Leanne Bischof
CSIRO – Mathematical and Information Sciences
Locked Bag 17, Building E6B, Macquarie University
North Ryde NSW 1670 Australia
{Hugues.Talbot,Leanne.Bischof}@csiro.au

Abstract

The Polartechnics SolarScan is an Australian medical instrument designed to help physicians diagnose melanoma. It has been in development in Sydney since 1994 and has now entered its commercial phase. Because of IP issues many aspects of the image analysis subsystems could not be discussed openly until recently. In this paper we describe some of the algorithm designs that makes this instrument powerful and reliable.

1. Introduction

Melanoma is the most deadly form of skin cancer. About 1000 Australians die each year as a result of melanoma. As melanoma can readily spread through the whole body it must be detected and treated early for best survival chances. Most Australians go to their GPs for skin checkups, however not all GPs can be trained to diagnose early melanoma, and as it is a relatively rare skin condition, most GPs will only see a handful of melanomas in their entire career. To err on the side of caution, a lot of benign skin lesions are also needlessly excised.

The idea of the Solarscan instrument is to allow clinics to have access to a reliable diagnostic aid that will lower error rates. This instrument needs to be relatively cheap, fast, accurate, give reproducible results, be able to follow up patients and above all have a low error rate.

The result of a collaboration between Polartechnics Ltd, The Sydney Melanoma Unit at Royal Prince Alfred Hospital and CSIRO-MIS, the Solarscan instrument is one of a few such instruments that could fit the above description.

In this paper we present an overview of the image analysis algorithms that makes this instrument useful.

2. Instrument design

The Solarscan instrument consists of a video capture device connected to a standard PC running Microsoft Windows NT. The video capture device is a special-purpose designed camera with a 3-CCD 760×560 captor, self-contained flat illumination and surface microscopy (epiluminescence) optics. The front cover of the camera is changeable and can be replaced to take standard non-microscopy pictures, for example to save the location of moles on the body in the patient database for future reference.

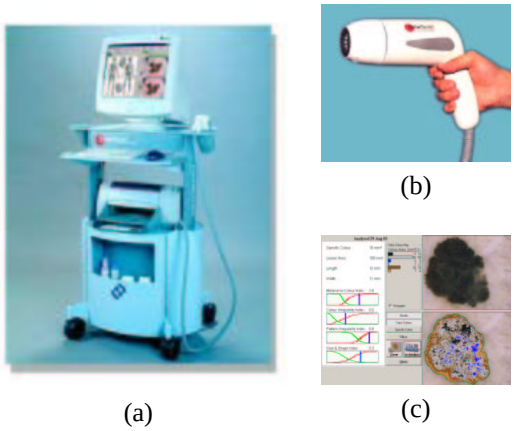

(a) (b) (c)

Figure 1. The SolarScan instrument: (a) global appearance, (b) camera, (c) user interface.

All captured images have 4 colour swatches in each corners for colour calibration.

The diagnosis system has two main components:

- The image analysis pipeline. Its aim is to segment the lesion from the image and to provide feature extraction.

- The diagnosis model. Its aim is to provide diagnosis aid to the physician. It can simply be a probability of a lesion to be a melanoma in the case of a full diagnosis system, or it can provide significant numbers that are in tune with the physician's training (ABCD rules for example [5, 8]).

3. Image analysis pipeline

In this section we describe in more detail the image analysis pipeline leading to feature extraction.

3.1. Image calibration

All images taken with the SolarScan instrument contain four swatches of known reflectance which can be used to calibrate images between session and between instruments. The aim is to be able to refer all pixels to a known fixed colour space such as CIE XYZ.

A standard uniform gray background is also imaged from time to time to provide background illumination correction.

3.2. Artifact removal

Surface microscopy involves a layer of oil between the lesion and the optics. Air bubbles can be trapped in this layer. Additionally, hair is frequently present on images of skin. Because of the danger to the lesion it is not usually shaved and must be detected and masked out.

The procedure to detect air bubbles involves taking a top-hat filter on the lightness component of the calibrated image, together with strong edges detected with a morphological gradient [11].

Hairs are detected using a more complex procedure involving the algebraic closing consisting of the intersection of a series of closings by linear structuring elements [13]. The artifact detection is illustrated in Fig. 2

Obviously the calibration patches within the field of view are also masked out.

3.3. Lesion segmentation

An accurate lesion segmentation is critical to the whole procedure and has been subject to a significant body of research [15, 9]. The procedure is difficult because of the large variation in skin colour in the population and the equally large variation of lesion appearance.

For reliability an entirely automated procedure based on seeded region growing [1] is first run. If the result is judged unsatisfactory by the operator a semi-automated procedure is tried next, based on colour clustering.

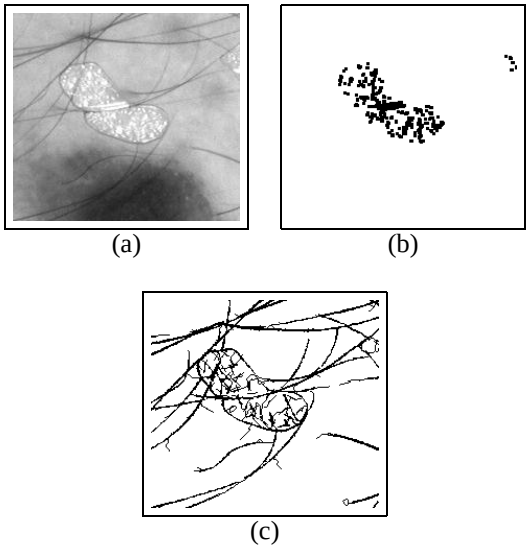

Figure 2. Artifact removal on subset of lesion image: (a) lightness component, (b) air bubbles, (c) lesion hair.

In both approaches, a principal component analysis of a representative sample of calibrated images is peformed. PC1 and PC2 are the first and second principal components, respectively.

3.3.1. SRG-based procedure

The main difficulty in an SRG-based procedure is to find the seeds. In this case the lesion seed is obtained from the PCA analysis. PC1 is roughly equivalent to lightness,the lower 20% of which are used as a seed for the lesion, assumed to be darker than the rest of the image. The brightest 20% are assumed to be skin. A standard SRG algorithm is run on the calibrated data with these seeds.

3.3.2. Colour clustering

This approach is based on clustering the bivariate histogram PC1 vs. PC2 of the PCA transform of the calibrated image. The clustering technique is based on finding peaks in the bivariate histogram and segmenting by watershed with a technique similar to that described in [12]. The colour clusters are then ordered on the basis of increasing lightness, which yields an ordered series of image segmentations, from darkest to lightest. To limit the number of boundaries to a reasonable number, skin and lesion statistics are used from the previous SRG procedure.

(a)

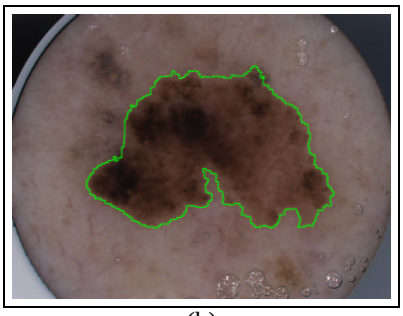

(b)

Figure 3. Example of lesion segmentation.

3.4. Lesion border analysis

Clinical analysis of the border of a lesion has been shown to be significant [7]. A number of researchers have published automated methods of boundary analysis [2, 10]. In this project we systematically reused all available published methods including fractal analysis [4], notch analysis, symmetry analysis and more. One of the most significant variables in subsequent statistical analysis turned out to be the lowly area measurement (i.e: larger lesions were more likely to be melanomas than smaller ones).

Novel methods were also developed. As an example we present our method for determining edge abruptness. Clinical studies have shown that the sharpness, or abruptness of the transition between the lesion and the surrounding skin is diagnostic: melanomas tends to have sharper transitions than benign lesions. To extract the edge profile around the boundary, a distance function (DF) is computed from the edge of the lesion (in both directions, inside and outside). From each point on the boundary a profile is run towards the inside of the lesion and towards the outside, following the upstream [6] of the DF, recording the lesion and skin intensities respectively. This allows the reconstruction of a reliable surrogate for the gradient of the border, as shown on Fig. 4.

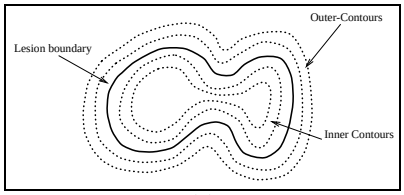

(a)

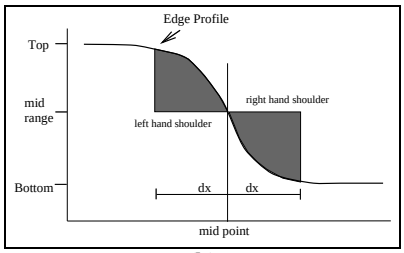

(b)

Figure 4. Edge abruptness measure.

3.5. Colour segmentation

Colour analysis is also significant for the melanoma diagnosis. Melanomas tend to be more colourful than benign lesions. Both absolute colour (i.e: calibrated colours) and relative colour (variation in colours within a given image) measures were derived.

3.5.1. Absolute colours

An RGB cube of calibrated colours derived from a significant subset of images was classified into various clusters. Several methods were used for this, including a complete watershed-based segmentation of the cube, which resulted in a fixed 3D look-up table. Applying this LUT to any calibrated colour skin image results in a pixel-wise classification. From the statistics of the classified regions various measures were derived, such as relative areas of colours, presence or absence of particular colours, etc.

3.5.2. Relative colours

For this an automated clustering of the first two PCA components of individual lesions was performed, resulting again in pixel classification. Because these regions do not correspond to known colours, the resulting regions were used as input for a categorical symmetry analysis (symmetry analysis between regions) based on simple Euclidean distances between clusters. Various other inputs were used for categorical symmetry analysis, such as 1-D histogram segmentation (in all 3 channels) and the absolute colour classification.

Figure 5 shows a sample absolute colour segmentation.

(a)

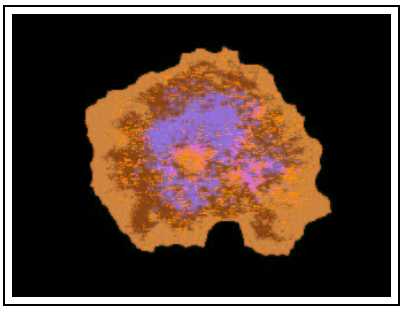

(b)

Figure 5. Lesion absolute colour segmentation. (a) original image, (b) colour segmentation.

3.6. Texture analysis

Some malignant lesions are not colourful, but can still be distinguished from benign lesions by texture analysis. More precisely melanomas tend to have less regular textures and less symmetry than benign lesions. However defining regularity and symmetry in skin lesions is challenging [14]. To capture the notion many different features were extracted. We present a few examples:

3.6.1. Simple measures

Relatively obvious measures were obtained, such as the total intensity variance in lightness or in each of the colour channels. Some extension of symmetry measurements applied to the binary mask of the lesion were also extended to grey-level and colour, such as the flip and rotation measures. These consist of finding the axes of symmetry of the lesion (by moment-based methods), and then measuring the amount of correspondence between parts of the lesions on each side of the axes. This can be done by flipping the lesion about each axis and measuring the average absolute pixelwise difference over the overlap.

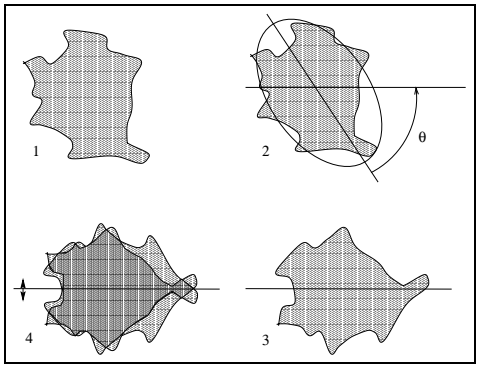

Figure 6. The flip symmetry measures.

3.6.2. More complex measures

Standard unsupervised texture classification methods (Wavelet, Gabor, etc) were too slow and did not appear to be suitable to the particular kind of biological texture found in skin lesions. However we did persist with some FFT-based texture measures with mixed results.

A number of ad-hoc texture measures were also developed. These measures were designed to match patterns often found in lesions, often labelled as “network”. The associated measures correspond to a coverage percentage of a lesion. The idea being that if a single particular texture covers a large area portion of a lesion, it is probably regular.

3.6.3. Specific features

Some malignant lesions can only be diagnosed as such by looking at some specific visual features that need to be looked for. A list of such significant features can be found in [7].

An example of those are the border spots, which are small dark spots on the boundary of a lesion corresponding to melanocytes in malignant stage. Such features can be extracted by black top-hat restricted to the boundary of the lesion, as seen on Fig. 8

4. Feature extraction and selection

Many more features were extracted and measures derived than can be presented here. As an overview, the features not already present in the literature (the vast majority) were designed in collaboration with an expert specialist clinician. The idea of the design phase was not necessarily to come up with features that by themselves were correlated to malignancy, but which reasonably matched something that the clinician could see *and* was in some way related to the highly visual human diagnosis procedure.

The feature selection was made by letting the data itself drive the selection. The set of all measures was fed to a

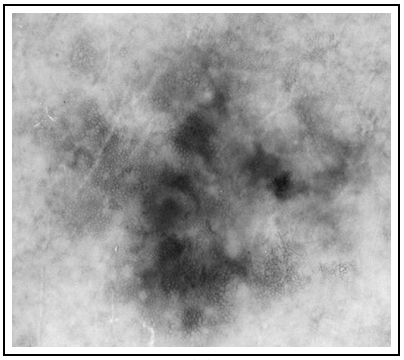

(a)

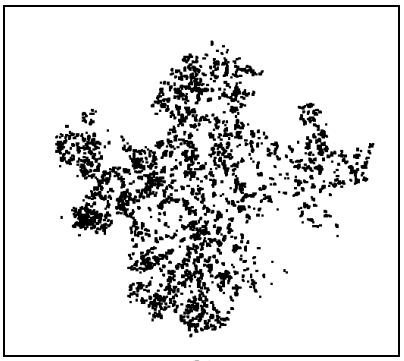

(b)

Figure 7. Example of network extraction: (a) lesion with network, (b) network coverage.

series of classification methods, from logistic regression to regression trees. After culling a number of models were built using cross-correlation.

The database of images used for building these models grew during the length of the project. It started relatively small with about 30 melanomas and a few hundred atypical benign lesions but now contains several hundred melanomas and several thousand both typical and atypical benign lesions.

At some point in the research the number of extracted features grew to over 600, but was later reduced to about 80. Fewer than a dozen are typically needed for a model.

4.1. Diagnostic features

Not much detail can be given here, but it has become apparent that relatively simple measures are often more valuable than complex one. Simple measurements such as lesion area often appear at the root of the model. These measures have the advantage of being robust and highly repeatable whereas more complex one can give signficantly different answers under relatively slight changes in image capture

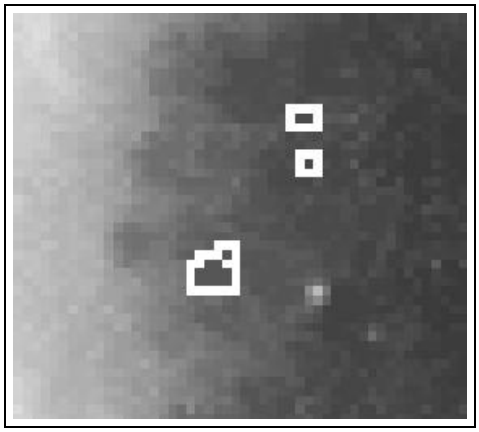

Figure 8. Example of border spot segmentation.

conditions.

However we have found it impossible to build an accurate model without some of the more advanced measures, especially those that look for specific features in lesions.

4.2. Accuracy

The SolarScan instrument was originally designed to be a diagnostic instrument, however it is now being marketed simply as a diagnostic aid which gives several indications to the operator rather than operate as a black box giving a single probablity or yes/no answer. As such it is not possible to give a concise accuracy figure for this instrument.

However early in the development of the instrument, with only some of the features now being collected and a much smaller image database, this instrument obtained a cross-validated sensitivity (ability to diagnose a melanoma correctly) of 92% and specificity (ability to diagnose a non-melanoma) of 62% [3], which is comparable to an inexperienced skin specialist and better than most GPs. It is presumed that the numbers for the present instrument are higher.

5. Conclusion

The Polartechnics SolarScan melanoma diagnosis assistance system is an innovative Australian instrument which has followed the difficult road from research to commercialisation thanks to a unique collaboration between research organizations and a private company. Image analysis is a key component of this complex system. In this paper we were able to describe parts of the design of the image analysis subsystem.

The proof of the validity of this research will be if the instrument becomes successful as a commercial product, keeping in mind that the final objective of such instruments is really to save lives.

References

[1] R. Adams and L. Bischof. Seeded region growing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(6):641–647, 1994.

[2] M. Binder, H. Kittler, A. Seeber, A. Steiner, H. Pehamberger, and K. Wolff. Epiluminescence microscopy-based classification of pigmented skin lesions using computerized image analysis and and artificial neural network. *Melanoma Research*, 8:261–266, 1998.

[3] L. Bischof, H. Talbot, E. Breen, D. Lovell, D. Chan, G. Stone, S. Menzies, A. Gutenev, and R. Caffin. An automated melanoma diagnosis system. Ballarat, Vic., Australia, 1998. Submitted.

[4] E. Claridge, P. Hall, M. Keefe, and J. Allen. Shape analysis for classification of malignant melanoma. *Journal of Biomedical Engineering*, 14:229–234, May 1992.

[5] NIH Consensus Conference. Diagnosis and treatment of early melanoma. *Journal of the American Medical Association*, 268(10):1314–1319, 1992.

[6] G. Matheron. Example of topological properties of skeletons. In J. Serra, editor, *Image Analysis and Mathematical Morphology*, volume 2, Theoretical Advances, pages 217–238. Academic Press, London, 1988.

[7] S. W. Menzies, K. A. Crotty, C. Ingvar, and W. H. McCarthy. *An Atlas of Surface Microscopy of Pigmented Skin Lesions*. McGraw-Hill, Roseville, Australia, 1996. ISBN 0 07 470206 8.

[8] F. Nachbar, W. Stolz, T. Merkle, and et al. The ABCD rule of dermatoscopy. *J Am Acad Dermatol*, 30:551–559, 1994.

[9] T. Schindewolf, W. Stolz, R. Albert, W. Abmayr, and H. Harms. Comparison of classification rates for conventional and dematoscopic images of malignant and benign melanocytic lesions using computerized colour image analysis. *European Journal of Dermatology*, 3:299–303, 1993.

[10] S. Seidenari, G. Pellacani, and A. Giannetti. Digital videomicroscopy and image analysis with automatic classification for detection of thin melanomas. *Melanoma Research*, 9:163–171, 1999.

[11] J. Serra. *Image Analysis and Mathematical Morphology - Volume II : Theoretical Advances*. Academic Press, London, 1988.

[12] P. Soille. *Morphological Image Analysis, principles and applications*. Springer, 1999.

[13] P. Soille and H. Talbot. Directional morphological filtering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(11):1313–1329, 2000.

[14] W. Stoecker, C.-S. Chiang, and R. Moss. Texture in skin images: comparison of three methods to determine smoothness. *Comp. Med. Imag. Graph.*, 16(3):179–190, 1992.

[15] L. Xu, M. Jackowski, A. Goshtasby, D. Roseman, S. Bines, C. Yu, A. Dhawan, and A. Huntley. Segmentation of skin cancer images. *Image and Vision Computing*, 17:65–74, 1999.