Detecting **malignancy** automatically

Professor of Computer Science **Dr Stella Atkins** describes an exciting project that is applying machine learning techniques to aid development of an automated device to diagnose and report on the nature of suspect skin lesions

What inspired you to begin your current project?

I have always been interested in applying computing science to the medical arena. I have explored medical image analysis ever since Magnetic Resonance display technology switched from film-based images to digital monitors.

I became involved in cancer research through a former PhD student, Dr Tim Lee, who works at the British Columbia Cancer Research Centre and the University of British Columbia on biomedical imaging of skin lesions. After he graduated we continued to collaborate, leading to our current project on diagnosis of atypical moles.

Can you outline the main focus of the project?

Diagnosis of malignant melanoma is difficult for family practitioners, who rarely encounter it in daily practice. Melanomas can be confused with benign lesions such as naevi and seborrheic keratoses, which are far more prevalent.

Our goal is to develop a low cost skin lesion screening system to help healthcare professionals to readily determine whether or not a patient should be referred for a skin biopsy or to a dermatologist. We hope that the system will also be used by experts as a screener and 'second reader' in early diagnosis.

Why is early diagnosis of melanoma important?

It is a particularly malignant cancer, often resistant to conventional chemotherapy or radiotherapy treatments. Early diagnosis is critical for it to be completely excised surgically while it is still localised and potentially curable.

What is your current stage of progress with the development of a dermoscopy device?

We have tested three dermoscopy systems: a handheld dermoscope attached to an off-the-shelf digital camera capable of acquiring images in 'raw' (uncorrected) mode; another dermoscope with polarised light and a built-in camera that attaches via USB to a computer; and a dermoscope which attaches to a smartphone camera.

We developed calibration techniques based on imaging a colour chart with known spectral properties and under consistent lighting and have applied appropriate corrections to images acquired with the cameras. We are also developing algorithms for large-scale, automated diagnosis of skin lesion images.

In order to classify images automatically, the lesion must be 'segmented' from the skin. Could you explain how this technique is conducted?

There are two major approaches for automatically segmenting lesions in skin images: using image processing techniques based mainly on colour, and using a supervised machine learning algorithm, which requires a 'training set' of images whose pixels are used as a gold standard, for testing on new, unseen images using established parameters.

With the machine-learning method, for every new image each pixel is assigned a value of the likelihood that it belongs to either 'skin' or 'lesion'. Thresholding the likelihoods for all pixels determines how the pixel is to be labelled. We have obtained a very high per pixel sensitivity and specificity using both approaches.

How do you identify texture features?

From the segmented images, we can automatically calculate image features based on dermoscopic patterns, with reference to gold standard images. This technique enables us to model meaningful texture measures from visual inspection, and correlate these with automated texture measures derived from computer-based techniques.

For whom will the system be initially available?

Initially, the system will be used by dermatologists in our clinic using a specialised dermoscope attached to a smartphone camera. After a small pilot trial, we will tune image processing and classification parameters for maximum accuracy, and then make the system available to other dermatologists in Canada in a multi-site trial. Our aim is that in future work, a beta version may also be made available through a smartphone app planned for release during the 2015 World Congress of Dermatology in Vancouver.

In terms of the project's success to date, are there any achievements of which you are particularly proud?

My PhD student Paul Wighton won the Albert Kligman Young Investigator Scholarship, in 2011. Another PhD student, Maryam Sadeghi, won the gold award out of 2,400 submissions at the World Dermatology Congress in Seoul in 2011, which was a very great honour; she also won the silver medal award in the Grace Hopper Student Research Competition. Maryam led a team to develop a free iPhone app for melanoma prevention, 'UV Canada', which was featured on television across Canada, and which has had over 35,000 downloads and updates – I was very proud when her work was broadcast and received an award at the BC Cancer Research day.

Pixel perfect diagnosis of melanoma

A project underway at the **Simon Fraser University in Canada** is developing innovative technology to improve diagnosis of melanoma by adapting dermoscopy and photography methods and using artificial intelligence to create handheld devices capable of inspection and rapid analysis of suspect skin lesions

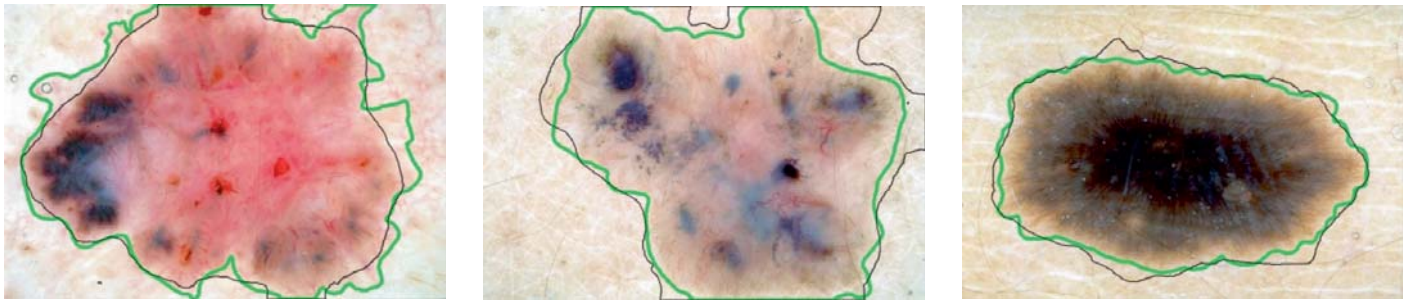

FIGURE 1. Sample segmentation results for our automatic method (denoted in black) compared to ground truth (denoted in green).
From left to right: A complete lesion with variegated colours, a challenging partial lesion with fuzzy borders and variegated colours, and a typical easy to segment lesion.

MALIGNANT MELANOMA IS a virulent form of skin cancer that, if it spreads to the lymph or the other organs of the body, can bring about death within five years of diagnosis. In recent years, more cases of melanoma have emerged and this is thought to be due largely to people's increased exposure to ultraviolet light, which is the main trigger for melanoma, as well as their greater awareness of the incidence and risks of skin cancer.

Melanoma initially appears as small freckles or mole-like marks, usually on the skin but also in the eye, that gradually grow and change shape: clinical diagnosis is usually via visual inspection and characterisation of the freckle or mole according to what are known as the ABCD parameters: Asymmetry, Border, Colour and Diameter. E is sometimes also used: the Evolution of the lesion in terms of changes in shape and size and the feel of the surrounding skin. Diagnosing melanoma is uncertain, but the lesions tend to be asymmetric, to change in size and nature and to display irregular borders and uneven colour. "Experts may also examine a skin lesion under high magnification using a dermoscope, a handheld magnifying device that uses polarised light; and consider further features for classification, especially those involving texture and textural patterns," explains Dr Stella Atkins, Professor of Computing Science at Simon Fraser University, near Vancouver.

Atkins has undertaken many years of research into support for medical science through the application of computing techniques: "All of my work is directed to improving patient care and outcomes using computers, across a wide spectrum – ranging from telepathology systems in Central Asia for improving pathology diagnoses to designing algorithms which enhance low dose computerised tomography images, allowing safer patient scans," she reflects.

Atkins is currently leading a project that is seeking to develop a tool to remove uncertainty from the diagnostic process for melanoma. In addition to improving accuracy of diagnosis, the tool should speed up patient referral for expert attention if malignant melanoma is suspected or diagnosed. Funded by the Canadian Government's Natural Sciences and Engineering Research Council, within their Collaborative Health Research Projects Programme, the project started in 2008 and expects to deliver its results by 2013.

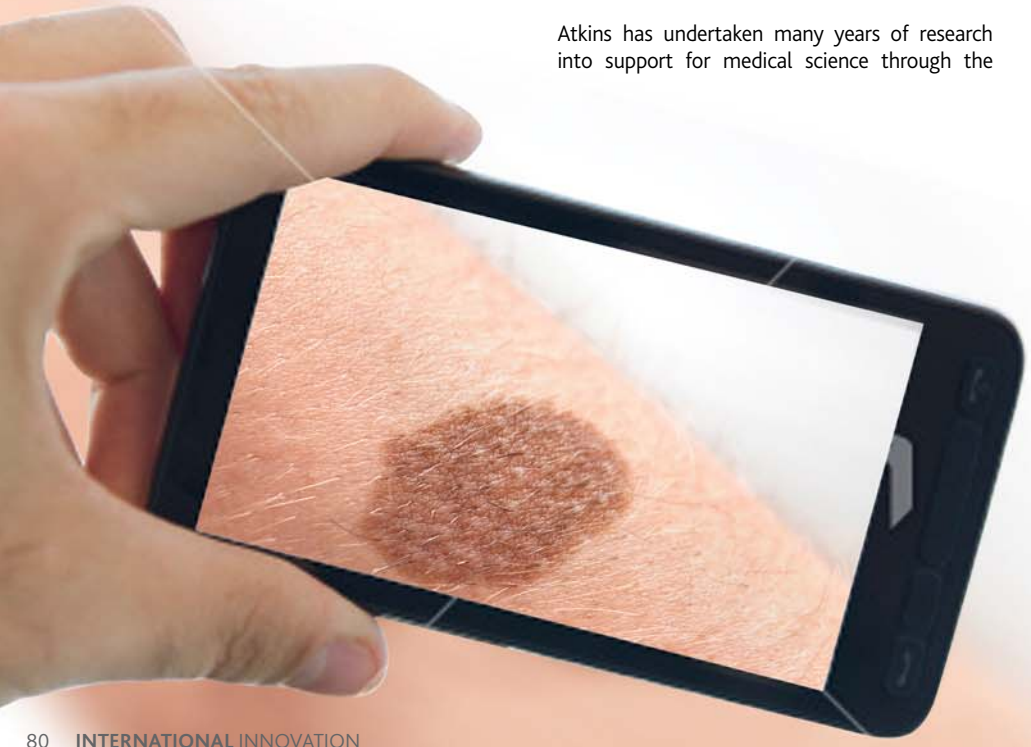

IMPROVING DIAGNOSIS OF MELANOMA

The project scope includes both melanocytic (moles) and non-melanocytic lesions, both of which can be cancerous. Variations and, indeed, similarities between types of skin lesions mean that any attempt to classify them automatically requires comprehensive and clear information about their characteristics as well as the ability to determine the likelihood of their nature by reference to impeccable sources. The project therefore employs machine learning, a variation of artificial intelligence, which is designed to derive determinations from complex empirical inputs by reference to 'experience' of examples of the types of inputs, the characteristics of which are amassed into a body of knowledge. In the case of Atkins' project, the input data is a previously 'unseen' image of a lesion, and the machine learning algorithm will classify the image according to the previously amassed body of knowledge.

IMPROVING ATYPICAL DIAGNOSIS OF MOLES

OBJECTIVES

• To develop and evaluate new low-cost dermoscope systems which will provide the necessary high-resolution images for an automated diagnosis system for moles

• To develop machine learning and image processing techniques to identify major features used for classification of images

• To develop methods to automatically, and reliably, process a large number of skin lesion images

KEY COLLABORATORS

Dr Harvey Lui; **Dr David McLean**, Vancouver Skin Care Centre
Technical collaborators at SFU School of Computing Science and the BC Cancer Agency:

Dr Tim Lee

Paul Wighton, PhD student

Maryam Sadeghi, PhD student

FUNDING

Canadian Institutes of Health Research

CONTACT

Dr Stella Atkins
Professor, School of Computing Science, Simon Fraser University (SFU)

Adjunct Professor, Department of Skin Science and Dermatology, University of Britich Columbia (UBC)

Technology and Science Complex (TASC-I) Building, Room 9237
Simon Fraser University
Burnaby, British Columbia
Canada
V5A 1S6

T +1 778 782 4288
F +1 778 782 3045
E stella@cs.sfu.ca

www.cs.sfu.ca/~stella

M STELLA ATKINS received a BSc degree in Chemistry from Nottingham University, and a PhD degree in Computer Science from the University of British Columbia. Her research interests include medical image analysis for skin images, MR image analysis, and the use of eyegaze trackers, particularly for surgery training.

The rules that the project team are applying are expressed via algorithms based on quantitative data obtained from computer 'training data' about lesion colours, dimensions and textures: "For the training set we developed a tablet PC-based programme for dermatologists to draw the borders around lesions in skin images. We used the dermatologists' outlines on 120 images as our gold standard when validating our methods," explains Atkins, who goes on to point out the painstaking nature of compiling basic gold standard data for computer training: "We initially focused on measuring the dermoscopic structure 'pigment network' (PN) – which is important for diagnosis, and obtained 94 per cent accuracy in identifying lesions with PN Present". With the project well underway, the team are now detailing the dermoscopic structures called 'streaks' and 'pseudopods' of the dark, radially-orientated lines that can be found in melanoma.

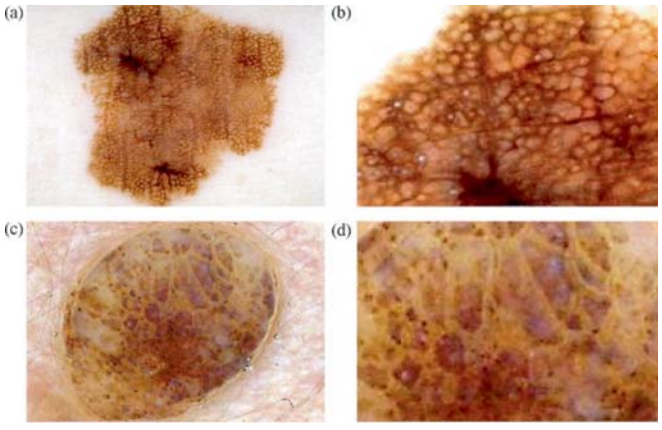

FIGURE 2. a) Present: a lesion containing a pigment network. **(b)** Enlarged pigment network. **(c)** Absent: an image of a lesion without pigment network. **(d)** Enlarged Absent image.

SUPERVISED MACHINE LEARNING

Once the base training data has been acquired, the computer will separate the lesion area from the new input image by mapping the lesion, pixel by pixel, in line with the labels recorded by the dermatologists. During this process, the computer will ascertain whether, according to probability, the part of the image in question is skin or lesion, and if lesion, what can be implied from its position relative to the other lesion pixels – inferring a relationship between the elements of the image of the lesion and determining the probability of its conformance to the criteria for lesion (rather than skin). The machine-learning approach that the project is taking is 'supervised', ie. it has a gold standard training set, thereby maximising the quality of the computer's algorithmic determinations. Supervised machine learning methods are also used to identify lesions having specific texture characteristics such as PNs and streaks, through identification of relevant features associated with these textures. The project hardware platform comprises a variety of PC processors, Apple and other smartphones, hand-held dermoscopy devices and cameras. So far, the project has found that images of lesions taken in uncorrected mode provide better inputs for machine learning, and that other picture formats require additional support: "We are continuing to develop new correction methods based on basic camera parameters observed from imaging standard colour charts," affirms Atkins.

A MULTIDISCIPLINARY PROJECT

The project is truly multidisciplinary, bringing together world-class skills in computing, cancer care and dermatology. Among its key participants are Drs Harvey Lui and David McLean from the Vancouver Skin Care Centre, and Dr Tim Lee from the British Columbia Cancer Agency.

Practical considerations arising from the different perspectives require balance and compromises to be achieved, as Atkins sums up: "We computer scientists want super high-quality skin images obtained using oil on the skin to prevent white scale from appearing in the images. However, the dermatologists point out that in a clinical setting, using oil is very messy; it takes much longer to deal with each patient". Adjustments that have been made to accommodate such requirements include adapting the computer algorithms to take white scale in images into account and using gel or alcohol as a means to obtaining better defined images of lesions.

PROJECT OUTCOMES

The overall objective for Atkins' group is to develop a low cost point of care device for healthcare professionals, and it is intended to release a prototype smartphone version at the World Congress of Dermatology in Vancouver in 2015.

The project has already attracted significant interest from the dermatology world, with interim deliverables winning a number of awards, and has also given rise to a number of published papers on the techniques involved.

In the future Atkins intends to release the compendium of images and diagnostic information that the project is compiling, as a teaching and education aid for others: "We are creating a new atlas of over 300 dermoscopy images, in conjunction with Raman spectroscopy data and near infrared images, together with the pathology diagnosis. We already have over 170 dermoscopy and near infrared images, of which over 80 also have Raman spectroscopy data," she elaborates. Encouragingly, the progress made through this innovative marriage of computer science and dermoscopy looks set to continue.