

<p>Institution: University of Manchester</p>
<p>Unit of Assessment: UoA 11 – Computer Science and Informatics</p>
<p>Title of case study: Active Shape and Appearance Models</p>
<p>1. Summary of the impact</p> <p>Our research on Active Shape Models (ASMs) and Active Appearance Models (AAMs) opened up a radically new approach to automated image interpretation, with applications in industrial inspection, medical image analysis, and face tracking/recognition. We identify:</p> <ul style="list-style-type: none"> • direct economic impacts of at least \$20m pa, with much larger indirect impacts; • healthcare impacts in non-invasive cardiac monitoring, accurate detection of growth deficit, and more efficient drug trials; and • cultural impacts in cinema and computer games.
<p>2. Underpinning research</p> <p>The impacts are based on research that took place in Manchester from 1993-date, with the first major publication in 1995. The key researchers were:</p> <p>Prof Chris Taylor (1993-date)</p> <p>Prof Tim Cootes (1993-date: PG Fellow '93, Advanced Fellow '95, Lecturer '01, SL '02, Prof '05)</p> <p>Dr James Graham (1993-date: Senior Lecturer '92, Reader '11)</p> <p>Dr Carole Twining (1999-date: PDRA '99, RCUK Fellow '05, Lecturer '10)</p> <p>Dr Gareth Edwards (1996-2000: PhD Student '96, PDRA '99)</p> <p>Dr Rhodri Davies (1999-2008: PhD Student '99, PDRA '02-'03, Part-time PDRA '04)</p> <p>The aim of the research was to develop automated methods of image interpretation, with applications in industrial inspection, medical image analysis, and face recognition as important targets. We introduced the idea that interpretation of a particular class of images (eg magnetic resonance images of the knee) could be based on generative models of shape and appearance, learnt from a training set of similar images. The key findings/insights were as follows:</p> <ol style="list-style-type: none"> 1. We showed that statistical models of the shapes and spatial arrangements of structures in a given class of images could be learnt and used to interpret unseen images automatically [1]. The modelling approach required a user to annotate landmarks manually in the training images. 2. We showed that the approach could be extended to model both the shape and photo-realistic <i>appearance</i> of structures in a class of images, and used to interpret unseen images automatically [2]. The interpretation method was more efficient than in [1], but building the models still involved manual input. 3. We showed that shape models could be built automatically, without manual annotation of landmarks [3], and subsequently extended this to full appearance models learnt directly from images [4]. 4. We demonstrated applications of these methods in industrial inspection [5], medical image analysis [6], and face tracking [2].
<p>3. References to the research</p> <p>The research was published in leading journals, including the top journals in the field (IEEE Transactions on Medical Imaging, IEEE Transactions on Pattern Analysis and Machine Intelligence, Medical Image Analysis). Output [1] is the most highly cited paper in the journal in which it was published and the 13th most highly cited paper of all time in computer vision or medical image analysis. Outputs [2] and [3] are in the top 0.6% and 2.2% by citation in the journals in which they were published (all citation data 1/10/2013, ISI Web of Science).</p> <p>Key Publications</p> <p>[1] T.F. Cootes, D. Cooper, C.J. Taylor and J. Graham, "Active Shape Models - Their Training and</p>

Application." Computer Vision and Image Understanding (CVIU), Vol. 61, No. 1, Jan. 1995, pp. 38-59, DOI: 10.1006/cviu.1995.1004. (*Citations: 2554, 1/1587 in CVIU all-time citation list*)

[2] T.F.Cootes, G.J. Edwards and C.J.Taylor. "Active Appearance Models", IEEE Pattern Analysis and Machine Intelligence (PAMI), Vol.23, No.6, pp.681-685, 2001 DOI: 10.1109/34.927467. (*Citations: 1174, 29/4722 in PAMI all-time citation list*)

[3] R.H. Davies and C.Twining and T.F. Cootes and C.J. Taylor, "A Minimum Description Length Approach to Statistical Shape Modelling", IEEE Transactions on Medical Imaging (TMI), Vol.21, pp.525-537, 2002 DOI: 10.1109/TMI.2002.1009388. (*Citations: 204, 71/3246 in TMI all-time citation list*)

Other Relevant Publications

[4] T.F. Cootes, C.J.Twining, V.S.Petrovic, K.O. Babalola and C.J.Taylor "Computing Accurate Correspondences across Groups of Images", IEEE Pattern Analysis and Machine Intelligence, Vol. 32 (11) pp.1994-2005, 2010 DOI: 10.1109/TPAMI.2009.193.

[5] T. Cootes, G. Page, C. Jackson, and C. Taylor, "Statistical grey-level models for object location and identification," Image and Vision Computing, vol. 14, pp. 533-540, 1996. DOI: 10.1016/0262-8856(96)01098-0.

[6] T.G. Williams, A.P.Holmes, J.C. Waterton,R.A. Maciewicz, C.E. Hutchinson, R.J. Moots, A.F.P. Nash, C.J.Taylor. "Anatomically Corresponded Regional Analysis of Cartilage in Asymptomatic and Osteoarthritic Knees by Statistical Shape Modelling of the Bone", IEEE Transactions on Medical Imaging, Vol: 29 Iss: 8, 2010 DOI: 10.1109/TMI.2010.2047653.

4. Details of the impact

Context

Because it established a new generic approach to image interpretation, our research has achieved significant impact reaching across multiple domains. Here we highlight three:

- Printed Circuit Board inspection (with spin-off to Number Plate Recognition)
- Face image analysis
- Medical image analysis

In each case, solutions based on statistical models of shape and appearance offered significant advantages over previous approaches – which generally took the form of expensive to develop and maintain ‘hand-crafted’ solutions to specific problems. By learning shape and appearance from examples, our approach allows solutions to new problems to be developed quickly, whilst dealing reliably with the variability inherent in real data.

Pathways to Impact

Because of the broad potential of the research, we pursued multiple pathways to impact. The key ideas were published and have been adopted widely by the research and product development communities. Some specific developments were patented to provide protection for commercialisation partners. Potential routes to market were explored by Visual Automation Ltd, a knowledge exchange incubation company, wholly owned by the University and embedded in the research group. This resulted in University spin-outs Kestra (PCB inspection), imorphics (medical image analysis), and Genemation (face image analysis), which between them raised around £4 million in external investment. In addition, spin-offs Image Metrics and Optasia were established independently by staff and students from the group. Kestra was subsequently acquired in a trade sale by CyberOptics Corporation, a manufacturer of PCB production equipment, who provided a route to market.

Printed Circuit Board Inspection and Number Plate Recognition

Surface mount technology (SMT) is one of the key enablers of the consumer digital electronics revolution – driving miniaturisation and cost reduction, through automation. Component mounting errors are, however, relatively common, so virtually all production lines incorporate automated optical inspection (AOI) of the finished printed circuit board. Prior to our work, AOI systems used large libraries of handcrafted algorithms to deal with different component types, with parameters that had to be tuned carefully to deal with variation in appearance. Our technology allowed the inspection task to be learnt from good examples, dramatically simplifying set-up and improving true/false positive ratios. The technology is now fundamental to CyberOptics Corporation’s AOI

systems, which generated revenues of \$28.5m in 2011-2012 inclusive – having roughly doubled over the REF period – with around 750 new installations between 2008-2012 [A].

Starting in 2010, 4Sight Imaging, a spin-off from the University-Cyberoptics team, has adapted the AOI technology for use in automatic number plate recognition. Although this is a relatively mature market, the superior performance of their model-based system saw the installed base grow to around 100 systems (~£1m revenue) by February 2013, more than 60 in the year 2012/13 – with applications in police intelligence and parking management [B].

Medical Image Analysis

Sophisticated medical imaging methods have become ever more widely available, resulting in a flood of data and creating a demand for medical image analysis methods to extract information automatically. Image segmentation – identifying and extracting the boundaries (surfaces in 3D) of specific anatomical/pathological structures – is a key underpinning technology. Our research established a new segmentation paradigm, exploiting anatomical knowledge learnt from training images, which has been applied by major medical imaging companies (Siemens, Philips, GE) [E, F] and specialist SMEs – with both economic and healthcare benefits. We provide illustrative examples.

Imorphics: Orthopaedics and Image-Guided Intervention

Spin-out imorphics has a turnover of around £1m pa – based directly on our research – and works with major companies in pharma, orthopaedics and image-guided intervention. In Rheumatoid Arthritis it currently provides the image analysis technology underpinning critical drug trials for four global pharmaceutical companies, including a trial of a potential blockbuster drug. Here the use of automated model-based analysis achieves levels of sensitivity and specificity not previously available, allowing inferences regarding treatment effect to be made from much smaller trials, with significant time and cost savings. In orthopaedics, a global supplier of prostheses has invested around £2m with imorphics over the assessment period to develop technology for designing patient-specific surgical jigs from 2-view radiographs, for a knee replacement system due to enter clinical service in late 2013. In image-guided intervention, global engineering company Renishaw is using imorphics technology routinely in robotic stereotactic neurosurgery – to provide navigation from MRI scans to place deep-brain stimulation electrodes in the sub-thalamic nuclei and, currently in clinical trials, to provide targeted delivery of pharmaceuticals directly into brain. More fundamentally, a leading medical device company is using imorphics technology as the basis for a comprehensive image-guided surgical planning package for a new generation of surgical robots due for high-profile launch in 2014 [C].

Visiana: Bone Age Measurement

About 5% of children attend a paediatric endocrinology clinic, at some point, for investigation of possible growth disorders. Of these, around 1 in 50 receives expensive treatment with Human Growth Hormone, typically increasing quality-adjusted life years by around 5. Bone age (BA – a measure of developmental age), relative to chronological age, is a key factor in the decision to treat. The standard clinical work-up involves visual assessment of BA from a hand radiograph, but can easily involve errors of 1.2 years (2 SDs). The poor precision, and need for specialist expertise, limit the confidence with which treatment decisions can be made, and inhibit the development of assessment services. Visiana's BoneXpert system, which automates the assessment of BA from hand radiographs, depends critically on the use of AAMs. It produces more reliable estimates of BA, reducing errors to around 0.4 years (2 SDs), without requiring specialist expertise. It was launched in September 2011, and by April 2013 had been adopted by 22 major childrens' hospitals across northern Europe (increasing by around 1 per month), currently conducting around 13,500 diagnoses per annum [D].

Siemens: 3D Cardiac Ultrasound

Advances in medical image acquisition have been dramatic over the past decade. In order to exploit new capabilities effectively, leading medical imaging companies increasingly embed automated analysis facilities in their systems, and have made extensive use of ASM/AAMs. A concrete example is 3D/4D (real-time 3D) cardiac ultrasound. The technology for imaging the 3D structure of the beating heart, in real-time, reached maturity around 2005, but the raw data is exceptionally difficult to interpret. Consequently, embedded image segmentation is *essential* to realising a clinically usable system, allowing structured visualisation (via surface rendering), and

providing quantitative information (eg ejection fraction). Because it provides a non-invasive method of accurately assessing cardiac function, this has become the investigation of choice for cardiac patients – with significant healthcare benefits (that are, however, difficult to quantify). The Siemens 3D/4D cardiac ultrasound product uses image segmentation based directly on ASM/AAMs [F]. The company has a 12% share (\$150m pa) of the rapidly growing global 3D/4D market of \$1.25bn pa – with cardiac imaging the dominant application.

Face Image Analysis

AAMs and ASMs have been used extensively for face tracking and recognition, where the ability to model variability due to person, viewing conditions and expression are critical. Below we provide illustrative examples.

Microsoft Kinect and GE

Microsoft Kinect for Windows (KfW – a variant of the Xbox Kinect 3D sensor) was launched in 2011, targeting companies developing interactive products and services. Current partners include Nissan, Boeing, Pepsi, Siemens, Bloomingdale's, and Mattel, and the device is sold in 39 territories world-wide. Microsoft supply a software development kit (SDK) for the KfW, providing access to high-level user interface functions, including AAM-based face tracking. By February 2013, there had been 500k downloads of the SDK [G]. GE have also used AAMs in face tracking and analysis in contracts for GE businesses and government customers [E].

Film and Computer Game Production

Image Metrics (now Faceware Technology), established by PGR students Gareth Edwards and Kevin Walker, uses ASM/AAM-based technology for facial motion capture in film and video game production. Examples over the assessment period include feature films *The Wolfman* (2010), *The Mummy: Tomb of the Dragon Emperor* (2008), *Meet Dave* (2008), *The Curious Case of Benjamin Button* (2009 – Academy Award for Visual Effects), and computer games *Grand Theft Auto IV* (2009 – \$500 million sales in first week) and *Red Dead Redemption* (2010 – more than 5 million copies in its first two weeks). In 2011 (most recent accounts publicly available) Image Metrics' revenues were \$7 million pa [H]. Similarly, researchers at Carnegie Mellon University (Baker, now Microsoft; Matthews, now Disney), who developed efficient implementations of the AAM, licensed their work to Weta, where they developed the facial motion capture for feature films *Avatar* (2009 – Academy Award for Visual Effects), *The Adventures of Tintin* (2011), and *Rise of the Planet of the Apes* (2011) [I].

5. Sources to corroborate the impact

- [A] CyberOptics Annual Report (10K filing), March 2013, and covering letter.
Sales of AOI systems, dependence on the research.
- [B] Letter from 4Sight Imaging.
Sales of ANPR systems, dependence on the research
- [C] Letter from imorphics
Turnover, contracts with end-user companies, dependence on the research
- [D] Letter from Visiana.
Clinical context, hospital deployments, diagnoses pa, dependence on the research
- [E] Letter from GE Global Research
Routine use of the research in product development (eg face recognition, medical imaging)
- [F] ACUSON SC2000 White Paper and cited publication
Features of 3D/4D cardiac ultrasound product, dependence on the research
- [G] Microsoft online Kinect for Windows announcements
Face tracking SDK launch, dependence on the research, SDK downloads
- [H] Image Metrics Annual Report (10K filing), September 2011, and covering letter.
Motion capture revenues, contribution to specific films/games, dependence on the research
- [I] Web pages for Simon Baker (Microsoft Research) and Iain Matthews (Disney Research).
Licensing of AAM technology to Weta, use in production of specific films