


Research in the UEA School of Computing Sciences





Research in the UEA School of Computing Sciences

Research into Music Processing.

"Features and Classifiers for the Automatic Classification of Musical Audio Signals", In: Proc. 5th Int. Conf. on Music Information Retrieval, Barcelona, 2004.

Today, computing pervades every aspect of modern society, creating numerous research challenges. Reflecting the diversity of these challenges, the School's research encompasses several contemporary themes in computing. Pattern recognition and analysis are applied to computer vision and speech, work that dovetails with graphical modelling of people and the environment. Data mining is applied to business and, increasingly, to medicine, with a focus on the underlying mathematics and algorithms. Biological systems are modelled via 3D-image analysis, and computer languages are designed to simulate cellular processes.

The School provides a vibrant research environment in which to carry out computing research. Sited within the UEA Science Faculty, it houses research labs dedicated to virtual environments, speech, colour and computational biology. Currently, the School has 26 research-active faculty, around 15 research assistants, and over 50 postgraduate students. The School runs an active research seminar programme, attracting speakers from around the world. As well as providing exciting opportunities for PhD projects, the School runs a full postgraduate teaching programme.

The School's research is organised within three Laboratories, whose research activities are described below. There is a great deal of research overlap between these laboratories, and considerable collaboration with other UEA Schools and research institutes across the Norwich Research Park. Through SYS Consulting Ltd, the School also undertakes industrially funded research, with contracts ranging from local authorities to national companies. All-in-all it's an exciting time to be researching in computing at UEA.

The Graphics, Vision and Speech Laboratory

People use audio and visual signals to make sense of the world and to communicate within it. Developing machines with the same (or enhanced) capabilities has been a challenging research theme in computer science for decades. To accomplish this task, computers must process and generate audio and visual signals. Research in the laboratory is concerned with the analysis, processing, recognition and generation of these signals in applications such as colour vision, machine vision, computer graphics, avatars and speech, music and language processing. These technologies have many common theoretical foundations that include signal-processing, machine learning, statistical pattern recognition, time-series estimation, and automata theory.

There are three main areas within the laboratory (see below), each of which is an internationally-leading research group in its own right, and there is collaboration and synergy between them in research in specific application areas. For instance, our research on audio-visual speech synthesis and lip-reading draws on our experience in both automatic speech recognition and computer vision; research in automatic sign language generation integrates research in graphics, avatars and speech synthesis, and research in

colour imaging is closely connected with computer vision. Our laboratories are equipped with fast graphics computers, haptic devices, spectrometers, 3D displays, a well equipped soundproof recording booth and an eight camera video-based motion capture system.

The laboratory is particularly successful in collaboration with industry: the laboratory dominates the University's patent portfolio, the University's spin-out portfolio, and the School's EPSRC grant and consultancy holdings.

Graphics and Virtual Environments

The main focus of our graphics research is in two related areas: crowd simulation and automatic urban modelling. Other areas of interest include haptic feedback for dentistry, protein visualisation and haptic interfaces for people with learning disabilities (a haptic device gives the user a sense of touch or motion by applying force feed-back to the user).

The crowd simulation group (EPSRC funded) is focusing on increasing the realism of models of people's individual behaviour in large groups. This covers individual appearance, movement and interaction with the environment. Typically, crowd rendering engines implement a

Level-of-Detail (LOD) system whereby entities, such as people, are replaced by discrete, successively coarser, approximations as their distance from the camera increases. Currently we are working on adaptive LOD for the control meshes used by real-time approximation techniques. For efficient graphics card implementation, algorithms are being created to transform smoothly between discrete control meshes calculated offline. We are also finding ways to efficiently transfer textures between avatars with different mesh structures.

The automatic modelling of real cities and their population with virtual humans is based on our own software and has evolved into a fruitful collaboration with the School of History (focussing on historic reconstruction) and two successful EPSRC/SUSTRANS grants. In addition, our work on groups of people connects closely with the talking avatar research in the Laboratory.

Below: *Crowd Variety: HSV shift and Skeletal scaling performed on 1000 avatars in real-time.*

"Automatic Generation, Texturing and Population of a Reflective Real-Time Urban Environment", Computers and Graphics Journal, Vol.31, pp.625-636, 2007.



Vision and colour imaging

We undertake research in colour, scale-space methods, face and gesture computation and aspects of medical imaging. An important area is attempting to understand aspects of visual perception and providing algorithms that equip digital cameras with perceptual processing abilities similar to those of humans. This (EPSRC funded) research can be found in Hewlett-Packard cameras and in commercial photo-processing software. Besides being UEA's most prolific generator of patents, we have spun-out several companies of which the most recent, Im-Sense Ltd (www.im-sense.com), is developing image processing pipelines for digital cameras.

Colour is also used for medical imaging, and there has been a successful collaboration with the Norfolk and Norwich University Hospital to investigate how colour and texture processing can be used to understand images automatically from a capsule endoscopy camera (a camera in a pill). Other medical projects include the automatic registration of X-ray images and using scale-spaced compositions of images to aid diagnoses. This latter work complements the research that has developed a unique way to decompose an image, hierarchically, according to the size of image features. This "data-sieve" has been used in many applications, from artistic rendering of photographs, to image de-noising and it is also used in our work on lip-reading. We are at the forefront of research into modelling faces and automatically synthesising talking heads and this research is also funded by EPSRC.

Speech, lip-reading, music processing and avatars

There are several closely-linked areas involving speech, music, language and animated characters that are covered within the Laboratory. For many years, we have done fundamental research into speech and language processing algorithms (e.g. speech recognition in noise, formulaic language modelling, language processing for speech synthesis) and in the development of applications of speech processing (e.g. call-routing, recognition of speech transmitted using VOIP, dysarthric speech). A recent application area of interest is distributed speech recognition (DSR), an emerging technology in




Algorithms developed in the school are embedded in Hewlett Packard cameras.

"Gamut constrained illuminant estimation", International Journal of Computer Vision, Vol.67, pp.93-109, 2006.

which recognition is divided between a client and its server, and we have been active in defining the international DSR standard. We have developed methods of estimating speech information lost in transmission across networks and compensating for this loss in the recognition process.

We were one of the first groups in the world to research audio-visual speech synthesis, and have recently begun an EPSRC-funded project in collaboration with the University of Surrey to research in automatic lip-reading. Our music processing systems developed in conjunction with IMIRSEL at the University of Illinois won the Genre Classification, Artist Identification and Classical Composer

Identification tasks at the 2007 MIREX competition. The algorithms behind these systems have been patented, and we are currently commercialising them with venture capital funds as FindTunes. We have also been developing systems that can animate sign language from a "script" that is based on HamNoSys, a notation system developed by our collaborators at the University of Hamburg for transcribing sign language. Together with speech recognition, this technology is at the core of SiSi (Say It Sign It,) developed by students working with IBM.



Above: Multiple Choice Questions with Sign Language support (ESRC-funded project with Scottish Qualifications Authority).

"Providing Signed Content in the Internet by Synthesized Animation", ACM Transactions on Computer Human Interaction, Vol.14, pp.1-29, 2007.

Left: A figure created using a combination of a mesh and rendering techniques.

"High quality shadows for real-time crowds", In: Proc. EUROGRAPHICS 2006, Vienna, pp.37-42, 2006.

The Mathematical Modelling and Algorithms Laboratory

Knowledge Discovery and Datamining (KDD)

Organisations such as companies, government bodies, and research centres collect vast amounts of data on their customers, providers, employees, processes and the like. Storing and accessing large amounts of data is as difficult as it used to be, but the challenge now is to use this data effectively. Organisations need to recognise patterns and trends within the data, and to present this information to management in a manner that enables it to be properly exploited.

KDD provides a set of techniques to unearth such patterns. The KDD group has developed a methodology for using a range of procedures to deliver patterns of an appropriate standard and quality for management. Members of the group have also developed sophisticated techniques for machine learning including novel approaches to rule induction, time series analysis, clustering and ensemble techniques. These techniques are rooted in artificial intelligence, information theory and statistics. As data becomes more diverse, encompassing diagrams, pictures, charts and sound, so datamining techniques need to become more diverse, and the KDD group is now increasingly working with multimedia.



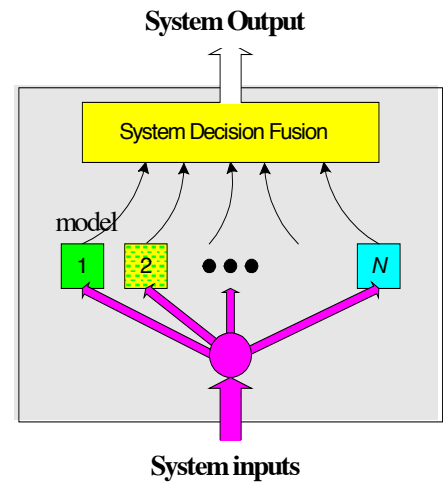
The UEA KDD Roadmap methodology as described in "Building the KDD Roadmap: A Methodology for Knowledge Discovery", In: *Industrial Knowledge Management*, pp.179-196, Springer, 2001.

The KDD group comprises seven faculty, eight research assistants and between 10 and 20 PhD students. As such, it is one of the largest such groups in the UK. It runs an MSc in KDD and over 70 of Aviva's staff have now graduated from that course. As well as working in the financial services sector, practical work at UEA has been undertaken on health related data with major studies in osteoporosis, cardio vascular disease and gastroenterology. Supporting this research, in the last five years funding has been received from BBSRC, EPSRC and The Royal Society as well as numerous companies (including Alston Transport, Derbyshire Police, Lanner Group, Master Foods, MET Office, National Air Traffic Services, Norwich Union, Process Evolution Ltd., Simultec AG (Zurich) and Virgin Money). Graduates in this area are highly sought after as industry and commerce begin to exploit the powers of KDD.

Optimisation

An important and widespread application of optimisation techniques concerns the efficient use of scarce resources to increase productivity. These applications include operational problems such as the distribution of goods, production scheduling and machine sequencing. They also include planning problems such as capital budgeting and facility location, and design problems such as telecommunication and transportation network design. Moreover, optimisation techniques lie at the heart of many KDD tools.

Most of the problems considered in the optimisation group are hard to solve, both theoretically and in practice. Exact methods will generally be slow, to the extent of being unusable. In consequence, heuristics are developed that can also be used to find approximate solutions with performance bounds, and this has been the focus of much of the group's research into discrete scheduling. In addition, better approximate solutions are found using meta-heuristics. These techniques have been widely applied in industry and often find reasonable approximations to optimal solutions.




Construction of hybrid ensemble of neural networks and decision trees.
 "Super Computer Heterogeneous Classifier Meta-Ensembles", *International Journal of Data Warehousing and Mining*, Vol.3, pp.67-82, 2007.

The optimisation group comprises six members of faculty and has three overlapping areas of interest – meta-heuristics, discrete scheduling and mathematical programming. The group has a rich history of collaboration with industry and government in the OR area: in scheduling, simulation and process optimisation (Derbyshire police, Jarrolds, Lanner, LloydsTSB, Unilever), network optimisation (Nortel, France Telecom), branch & cut and game theory (France Telecom), parallel branch & bound (EPSRC/DERA), Lagrangean relaxation (DERA), and meta-heuristics (DERA, Lanner, Nortel).

A random initial network for a genetic algorithm.

"A PROBE-Based Heuristic for Graph Partitioning", *IEEE Transactions on Computers*, Vol.56, pp.1701-1720, 2007.



The Computational Biology Laboratory

Physically located in between the Schools of Computing and Biological Sciences, and housed within the newly established D'Arcy Thompson Centre, the Laboratory promotes an integrative philosophy to computational biology with areas of research spanning the biological hierarchy from genome through to ecosystem.

In addition to carrying out high quality research with national and international partners, the Laboratory aims to develop computational biology across the Norwich Research Park (NRP). NRP represents Europe's largest and most successful single site concentration of research and development in the areas of plants, food, diet and health, the environment,

and information systems. Members of the Laboratory collaborate closely with scientists from NRP institutions such as the School of Biological Sciences (BIO), the School of Environmental Sciences, the John Innes Centre (JIC), the Institute of Food Research and the Sainsbury Laboratory.

The Laboratory provides an interdisciplinary environment for research and education, including an MSc programme. It runs its own research seminar series and has also been involved in running a number of international workshops. Main themes and recent highlights in research include:

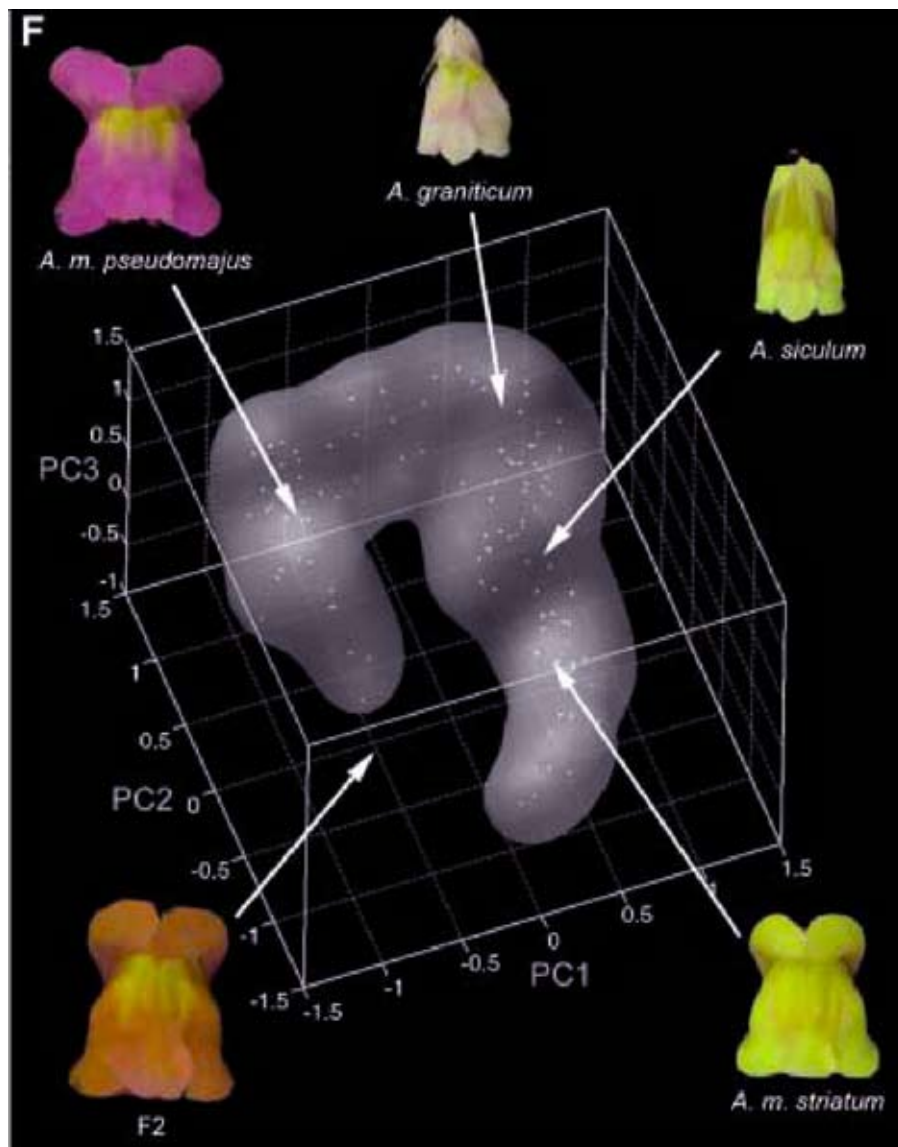
● **Analysis and modelling of plant growth:** Together with scientists at JIC, we have used active shape/appearance models (based on principal component analysis) to conceptualise the process of evolution. The method analyses the shape and colour of flowers that represent all possible genotypes to produce a three-dimensional 'cloud' that characterises the possible range in *Antirrhinum* species (see Figure). The resulting insights evoked newspaper commentary (Daily Telegraph) and a sharp correspondence in *Science*. We are the first to use optical projection tomography (OPT) to image the 3D structure of growing plants which is being used for inform models (BBSRC-funded), and we wrote the reference software that is now distributed with the OPT machines.

● **Protein structure, dynamics and function:** We are using both bioinformatics and simulation based approaches to help understand the relationship between a protein's structure and its dynamical behaviour in relation to function. The resulting DynDom program produced with a BBSRC-funded e-science project receives over 700 hits/month. It collates worldwide data on protein structures into a mineable database itself available on the web. (<http://fizz.cmp.uea.ac.uk/dyndom/>).

● **Phylogenetics:** Working closely together with groups in New Zealand, Korea, and China, in EPSRC-funded work we develop theory and algorithms for constructing evolutionary trees and networks. These are used to understand how organisms such as plants and viruses evolved. In 2007, several members of the laboratory took part in (and one acted as principal organiser for) a four-month program entitled "Phylogenetics" at the Isaac Newton Institute for Mathematical Sciences, Cambridge, UK.

Active shape/appearance model analysis of the shape and colour of flowers representing the possible genotypes in Antirrhinum.

"Evolutionary paths underlying flower color variation in Antirrhinum", Science, Vol.313, pp.963-966, 2006.




● **Computational biology of RNA:**

Together with BIO, we are data-mining for “short RNAs” (small genes that do not code for proteins) using in house developed algorithms. The rapid expansion of this area is driven by the surge of cheap, massively parallel sequencing, and in BBSRC-funded work we have recently used this data to uncover some of the molecular pathways involved in crop-plant development such as fruit ripening.

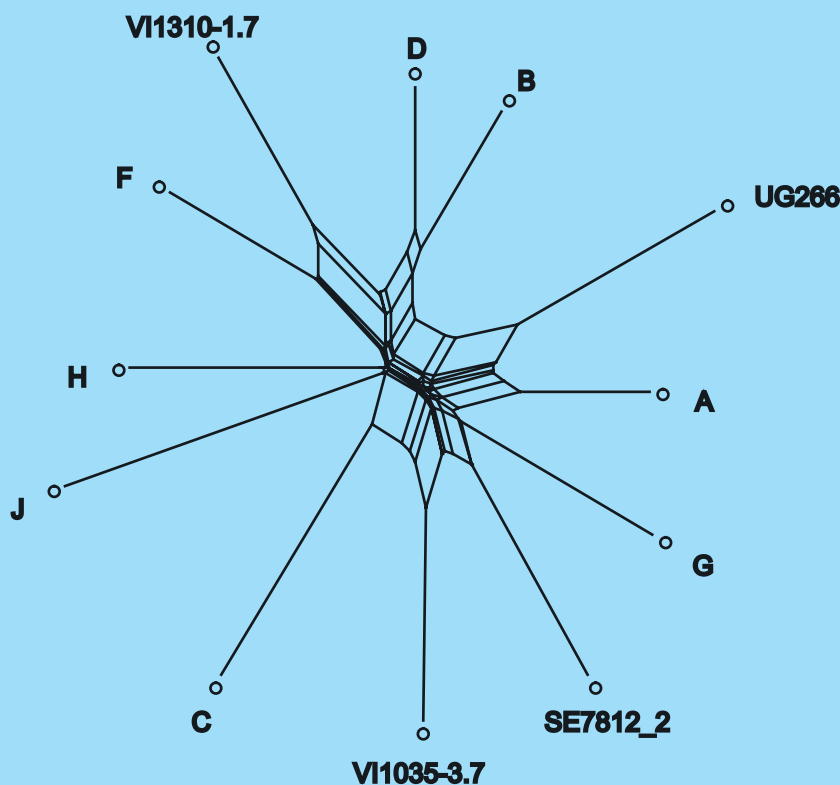
● **Computational systems biology:**

To model the regulatory networks that control the inner workings of the cell, we are developing the transsys language, a language that enables network description in a compact, object-oriented format, and that facilitates computational simulation of gene expression dynamics. Recent BBSRC funding supports the use of transsys to model the processes involved in plant wounding together with BIO.



Logo for the BBSRC-funded DynDom computer program.

“A comprehensive and non-redundant database of protein domain movements” *Bioinformatics*, Vol.21, pp.2832-2838, 2005.



A phylogenetic network representing the evolutionary relationships between AIDS viruses.

“Computing Phylogenetic Diversity for Split Systems”, *IEEE/ACM Computational Biology and Bioinformatics*, Vol.5, pp.235-244, 2008.

School of Computing Sciences

University of East Anglia

Norwich NR4 7TJ UK

www.uea.ac.uk/cm/home/schools/sci/computing

Tel: +44 (0)1603 592847 Fax: +44 (0)1603 593345

For more details contact prof. Vincent Moulton (vincent.moulton@cmp.uea.ac.uk)