ncaf
Natural Computing Applications Forum

Networks

Dynamics and Dynamic Systems

University of Sheffield, 9-10 September 2008

A major strength of the NCAF meetings is the eclectic mix of academia and industry.

NCAF's autumnal meeting is to be held in Sheffield, hosted by the department of Mechanical Engineering. The setting will be the tranquil Turner Museum of Glass and rooftop garden in the Sir Robert Hadfield building. Hadfield's discovery of manganese steel epitomises Sheffield's engineering excellence. It is fitting that the building should host NCAF, as natural computing adapts and progresses to pave the way for scientific and engineering discoveries in the 21st century.

The theme for the meeting will be dynamics and dynamic systems. Long-time NCAF member Professor Keith Worden (University of Sheffield) will deliver a discussion on the theme of uncertainty in dynamic systems. The dynamics group in Sheffield is currently holding an EPSRC Platform Grant in this area. Dr Neil Lawrence (University of Manchester) will present recent work on Gaussian Process Dynamics. With the fundamentals of Gaussian Processes in a machine learning context having been laid down with rigour over the past decade, Gaussian Processes are finding their way into a variety of specialised areas; Neil's work includes using them for system identification in biological systems.

The Turner Museum of Glass houses one of the UK's most comprehensive collections of 19th and 20th century glass.

Photo: The Turner Museum of Glass

The University of Sheffield has several departments with a strong presence in various aspects of natural computing, including groups in Engineering, Computer Science and Psychology. An interdepartmental forum (the Data Modelling Group) brings together these fields and encourages cross-disciplinary research with a machine learning focus. From this group, Dr Charles Fox will present some work decoding the movement of rat whiskers, and ex-data modeller Dr Michael Dewar (now at the University of Edinburgh) will talk about tracking and decoding the movements and behaviour of animals, while Giovanni Lupica (University of Sheffield) will discuss the identification of sub-cell organisms via supervised learning.

The theme of dynamic systems is one which fits well with current research across the University and beyond. In particular, the Automatic Control and Systems Engineering (ACSE) department, home to the Rolls-Royce University Technical Centre (UTC), has a strong research background in the modelling of dynamic systems. From ACSE (which, they tell me, is the largest control department in Europe) is Dr Rob Harrison, who will discuss the application of the Kalman Filter to geomagnetic storm prediction. Also from the University of Sheffield, Dr Stana Zivanovic will present a model of human and building dynamic interaction, and Dr Maurizio Filippone will discuss the modelling of human behavioural patterns.

A major strength of the NCAF meetings is the eclectic mix of academia and industry, and NCAF has long been a place where the two worlds collide. Maintaining this tradition, the coming meeting welcomes two seasoned speakers from industry, namely Dr Steve Reed (Defence Science and Technology Laboratory, DSTL), and Dr Steve King (Rolls-Royce).

As well as the series of speakers, NCAF provides an opportunity for people in the UK Natural Computing community to come together to share ideas and knowledge, start friendships and collaborations. Cemented in this idea is the regular NCAF social, which for this event is set to involve a trip to Kelham Island museum to experience Sheffield's industrial past, before a visit to one of Kelham Island's fine brewing establishments for a friendly pint and the ever-popular but really-rather-competitive NCAF pub quiz. We look forward to seeing you in Sheffield this Autumn!

James Hensman
University of Sheffield

How computers will keep us healthy

A series of remarkable slides depicting obesity trends in the USA showed the increasing 'weight' of the problem.

NCAF's summer meeting was held at St. Anne's College, Oxford, marking a return to the city after a seven year absence. The theme for the event was Signal Processing and Medical Applications, and a range of speakers presented ideas from both industry and academia.

Following a warm welcome from the local organisers, David Clifton and Mark Ebdon, the meeting opened with Visakan Kadirkamanathan (University of Sheffield), who highlighted the importance of understanding biological metabolic networks for assisting gene analysis and rational drug design. Techniques used to analyse such networks were then presented, using either a method based upon evolutionary algorithms or Markov chain Monte Carlo to solve both stoichiometric and carbon-13 balance equations simultaneously. This would determine the concentration of reactants in a network at steady state. Flux estimation under dynamic conditions was also considered, and a particle-filter solution was proposed. Results using these techniques were shown to work successfully when tested against simulation data.

Neil Townsend from T-Plus Medical provided the day's industrial focus, enchanting the audience with tales of princesses, knights in shining armour, and the increasing problem of obesity and diabetes. A series of remarkable slides depicting obesity trends in the USA showed the increasing 'weight' of the problem. In the midst of this fairy tale, Neil succinctly explained some of the difficulties in producing industrial products from ideas conceived within a university research group. Using a diabetes monitoring device designed at the University of Oxford as an example, he highlighted issues including the problem of regulations, product robustness, and lack of revenue.

Figure 1

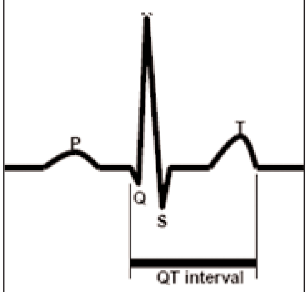

The QT interval is a critical measurement in ECG waveforms. www.mhra.gov.uk

Real-world problems

Following a hearty lunch, Christopher James (University of Southampton) provided a very broad overview of his research in neural engineering, covering everything from brain-computer interfacing to tracking worms. Christopher explained that his research is mainly in the development of advanced Blind Source Separation (BSS) techniques and application of these techniques to real-world problems in biomedicine and neural engineering. BSS techniques such as Independent Component Analysis (ICA) can be used in the analysis of both single-channel and multi-channel recordings of biomedical signals such as electroencephalograms (EEG traces).

Hujun Yin (University of Manchester) discussed the use of topological clustering for decoding population neuronal responses. He showed how a self-organising map can cluster the continuous responses of the discrete spike train recorded in the somatosensory cortex of rats. They were excited by stimulating the vibrissae (the hairs and whiskers on the face) at different frequencies and amplitudes. Hujun demonstrated how the clustering preserved information. A selection of results was presented that showed that clustering can naturally find underlying stimulus-response patterns and preserve information among clusters.

The final talk of the day was given by Lionel Tarassenko (University of Oxford), who updated the audience on progress regarding continuous patient vital sign monitoring. The objective is to use data fusion to detect early deterioration in patient health. Results from recent trials at the University of Pittsburgh Medical Centre using the vital sign system known as Visensia™ (the new name for BioSign™) were presented. One result of particular interest is that critical vital sign events were detected by Visensia far more frequently than by clinicians, who collected observations once in every four hours. Finally, the audience were given a vision of the 'hospital of the future' project, and a wider range of projects were mentioned, including telemedicine systems for primary care health.

After a full day of presentations, it was appropriate to enjoy Oxford's traditional summer weather of light drizzle with a hearty barbeque and a punting trip up to the local public house, the Victoria Arms. The outward journey was made more enjoyable by the quiz compiled by Graham Hesketh. The return journey was made more difficult because the quizmaster had lost his pole, later found floating down the river. The night was rounded off with an exciting musical quiz compiled by David Clifton.

Electrical signals

The second day began with Ji Won Yoon (University of Oxford) speaking about Brain Computer Interfacing (BCI) using information from EEG signals and labelling from EMG data (the electrical signals related to muscle activity) to create models that reduce the time required to train the BCI system. To do this, a state space model was set up and an appropriate Extended Kalman Filter was implemented. The hidden state noise and the observational covariances were marginalised out to improve the results. The technique worked well when tested on a selection of experimental results, i.e. the EEG signal matched the EMG training data.

Iain Strachan (Oxford Biosignals) presented current work in the field of determining cardiac response to drugs. Current FDA regulations require all new drugs to undergo a QT study, aiming to check for significant prolongations of the QT interval (a critical measurement in electrocardiograms or ECG waveforms) that can lead to cardiac arrhythmia (see Figure 1). Such studies are traditionally performed manually, with cardiologists marking up each individual ECG waveforms. It was noted that this process is often inaccurate, so in response a new set of techniques have been created to assist in QT analysis.

The first technique is a fully automated measurement of the QT interval that uses a hidden Markov Model (HMM), operating on a time series of observation vectors obtained by wavelet decomposition of the raw signal. However, a semi-automated system is required in order to provide verification of the automated analysis and hence the second technique is a semi-automated measurement that uses Dynamic Time Warping (DTW). A selection of results was presented using clinical data where QT prolongation had been manually observed, and it was clear that the techniques presented demonstrated the effect equally well.

Martin Terry, from the patent attorneys Kilburn and Strode, provided a quick overview of the patenting process, highlighting the difficulties in patents

Evolving Schumacher

It is likely that few of the readers of Networks will need persuading of the power and versatility of techniques such as artificial neural networks or genetic algorithms. What may be of more interest, however, is a novel application of these techniques grounded in the fields of emergent behaviour, artificial life, and computer gaming. Throw in an obscure but expressive programming language by the name of POP-11, and you have the makings of a rather interesting project. The brief was simple: can embodied agents in a simulation learn how to race a car?

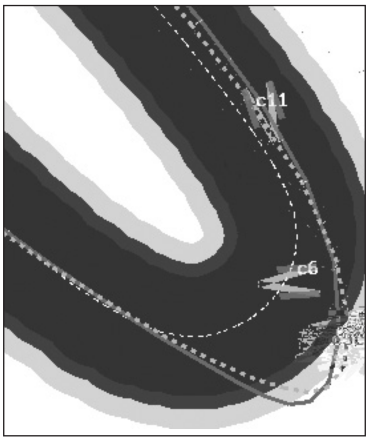

The black-ice slingshot-powerslide manoeuvre

The agents are modelled by neural networks, each in control of a simulated vehicle on a pre-drawn racetrack. Straightforward Newtonian physics was used to model such aspects as acceleration and friction, and a series of waypoints are drawn along the course of the track. The inputs to the neural networks consist of the speeds of the left and right wheels of the car, the car's current heading, a choice of which waypoint to observe, and the bearing to the chosen waypoint; essentially, these are equivalent input sources to those you could derive from real-world GPS co-ordinates on a way-pointed track.

Using random neural network initialisation; running the simulation causes a variety of different behaviours: spinning, looping, driving in reverse, crashing into barriers – and invariably not progressing very far along the track at all. The neural networks obviously require some training. Instead of performing supervised training to learn from a human driver (via backpropagation), the cars are given a set amount of time in which they can do whatever they choose to do. After the set time has elapsed, the grim reaper visits and culls all the agents, with the exception of the lucky two which have strayed by chance further along the track than their peers. The neural networks of these two agents are then combined using a genetic algorithm, and the resulting offspring replace the other cars which were less successful.

Within a few generations it can be seen that the cars are learning a preference for driving along the general direction of the track, even taking account of the bends in the track and adjusting their turning speeds accordingly. By this time the race is on, and the fastest two cars of each epoch (i.e. genetic algorithm iteration) are selected. The agents learn to cut the corners in order to better their race times – but not by too much, or they will crash off the track and get stuck – and before long, intelligent acceleration and deceleration through corners is observed. After approximately 15 epochs the cars are not only optimising their driving styles, but are taking advantage of the slipperiness of the road and actually power-sliding through tight corners.

The learnt behaviours can generalise to tracks other than the track on which the cars are 'born'. By placing the cars onto a new, totally-unseen track, all it takes is one or two further epochs before the new racing lines are being made and optimised. By altering the coefficients of friction on the track to make it appear as black ice, it is possible to confuse the drivers into overshooting all the corners. But soon enough, behaviours emerge such as hugging the high-friction edges of the track next to the barriers and sling-shotting through the bends are observed.

There is only one in-built rule in this system: the fastest cars to complete the track without cheating will survive. All this complex behaviour emerges from a simple combination of a genetic algorithm and some neural networks. Perhaps they are powerful techniques after all.

Mark Rowan
BAE Systems
(Project undertaken while at Birmingham University)

involving computer software. It was remarked that patenting applications differ between countries, and of particular note was the fact that computer programs cannot be patented in the UK, but may be patented in the US. The presentation generated many interesting talking points, raising questions about copyright and 'copyleft' (see http://www.copyrightservice.co.uk/copyright/p20_copyleft for more information), and continued well into the rest of the day.

The morning concluded with the solution to Puzzle Corner, where the puzzle master once again selected unsuspecting volunteers to re-enact the puzzle, and then showed how to solve the problem starting with the simple case of two boxes.

After lunch, Neil Thacker (University of Manchester) presented on techniques for multi-spectral segmentation of MR images. An approach for analysis of partial volumes measurement in MR images, where a measured signal at a point is a sum of two signals was described. The technique is a quantitative application of Bayes theorem in which prior terms are estimated to be consistent with the sample under analysis. This technique was applied to the brain tissues volume and a selection of results comparing the new methodology with previously published techniques was presented.

The final talk of the forum was from Stephen Payne (University of Oxford) who is looking into the using information theory for decoding the cardiovascular system. He started by highlighting how cardiovascular disease is one of the leading causes of ill-health in the western world and is becoming a serious problem in the developing world. A quick overview on how the body regulates itself was presented. This showed the challenges involved in interpreting cardiovascular system behaviour. Stephen showed how concepts of coupling and information transfer can be used between measurable parameters (such as blood pressure and heart rate) and then used to understand the cardiovascular system. A selection of results presented showed these techniques provide discrimination between healthy and unhealthy patients.

Many thanks are due to the local organisers, David Clifton and Mark Ebdon, to St. Anne's college, and to the speakers for creating another successful NCAF event.

David Wong and Busi Vilakazi
University of Oxford

PUZZLE CORNER

Number 39

The pinnacle of the sporting season was the upcoming 'Riotous Rounders Championship', a best of 7 game series between the 'Beaston Dread Locks' and the 'Old York Vets'. A lifelong Old Yorkonian himself, the 'Lord High Roller' was not averse to a modest flutter on the outcome. He knew that you could get even money on the 'Vets', so he gave Lisa £100 to bet on them to win, expecting to get back £200 if all went well. Of course, he expected nothing back if the 'Vets' were not ultimately crowned champions.

Lisa went to place the bet and was stunned to hear that she could not place a single bet on the overall series outcome. Instead, she could only bet on the outcome of each individual game (at even money). Of course, she was at liberty to change the amount of her bet after each game, but she had no other funds to draw upon beyond the £100 stake from the 'Lord High Roller'. Nevertheless, Lisa reasoned that there was a betting strategy which could guarantee to win the necessary £100 in the event of an overall 'Vets' victory.

What was Lisa's opening bet?

The answer will be given at the next NCAF meeting (9–10 September 2008, University of Sheffield).

Fenella the Rottweiler

COMMITTEE NOTES

Please contact NCAF
through Graham Hesketh,
Chairman – NCAF
PO Box 5944
Derby DE24 8ZD U.K.
Tel: +44 (0) 1332 246989
Fax: +44 (0) 1332 247129
e-mail: enquiries@ncaf.org.uk
http://www.ncaf.org.uk

Can model selection fail? Yes, it can!

Every scientific discovery starts with an observation. Nowadays that is normally some kind of measurement process. The inference process follows the observation. The goal of the inference is to find a plausible explanation for the data. For mathematically minded people, this means finding a model to fit the data. The first model that is considered tends to be – and should be – the 'simplest' one that comes to mind. Simple models are (near) linear and have a small number of parameters. Even with very simple models, the modelling task leads to a better understanding of the data. Modelling also provides insight into the current model's abilities and shortcomings, i.e. the conditions when the data cannot be explained. To remedy the situation the model is refined, meaning a process of remodelling the observations using more sophisticated techniques.

Thus one enters a cycle of model modification and revision. Each model is typically more complex than the previous one, i.e. the number of parameters increases and/or the model functional form changes. When to stop the cycle of model refinement is the topic of *model selection*. The guiding principle of model selection is *Occam's Razor* which states that, assuming similar prediction accuracy can be achieved by a set of models, then the simplest one of the set should be selected. This philosophy merely reflects common sense. For example, there is no need to factor in many variables if consumer spending can be predicted from personal income alone. The modelling task is performed hidden from view and inside the human onboard computer (the brain).

The philosophy is so general and simple that one rarely considers entertaining the option of it failing. But model selection does fail. I want to illustrate this on two representative cases.

Model selection is often misunderstood as a process of selecting the *true* model. For example, within the domain harmonic decomposition of autoregressive modelling the selection procedure involves choosing the *true* number of harmonics or coefficients that make up the signal. However, there is no *true* model or model order. The Wold Decomposition Theorem states that a finite-order moving-average model has an autoregressive model counter-part, but one with an infinite number of coefficients. Critical common sense is the best protection against falling into this trap. Prior knowledge should indicate that a particular model order value is simply inexplicable. For example, smearing due to cerebral spinal fluid and skull cannot produce autoregressive models of orders greater than 50 for a 1 second long window of a typical EEG recording.

Apart from producing unreasonable results, model selection may simply *fail to converge* to any result. This can happen during model selection for data with a strong fractal nature, such as financial time series. Larger model orders describe the data at smaller scales with almost equal accuracy. To test this, one can produce a Sierpinski Gasket-type of mixture model in which each mixture component is split again repeatedly into smaller components. Every choice for the number of mixture components captures a different scale of the division process and a plot of penalised or marginal model likelihood does not produce the expected 'knee' (before the training data size is too small, relative to model order, to allow accurate parameter estimates).

Model selection may fail to give a result or fail to give an accurate result, and it does so frequently. Although the theory of model selection is sound, the assumptions it rests upon are buried too deep in the literature for most of us to remember.

Besides, why choose the 'best' model if, in the same stroke, one acknowledges that there is no 'true' model – only the most practical one. The upshot is that, instead of model selection, model averaging should be the paradigm of choice (excuse the philosophical pun!). Model choice is immaterial as the only true criterion is model *performance*.

Iead Rezek
Imperial College, London

Southampton professor elected to lead global computing organisation

A Southampton professor has been elected president of the world's largest educational and scientific computing society – the first time in 60 years a person outside North America has held the position.

Professor Wendy Hall was announced as the new President of the Association for Computing Machinery (ACM).

Professor Hall, who is Professor of Computer Science at the University of Southampton, said her goal as president is to help ACM reach its full potential by expanding international initiatives and increasing gender diversity in all aspects of computing.

The ACM is the world's largest educational and scientific computing society with nearly 90,000 members around the world, representing business and the professions, education, and research and development.

As a past president of the British Computer Society from 2003-04 and a researcher with many international connections, Professor Hall expressed her commitment to guiding ACM towards more initiatives in India and China.

For further information see the ACM website http//:www.acm.org

MEMBERS' NEWS AND VIEWS

Deadline for contributions for the next edition – 1 September 2008.
Please send to Managing Editor – Nick Granville, e-mail:nicholasgranville@btinternet.com

Edited and Produced by:
Forum Public Relations
Westgate House, Old Ivy Lane
West End, Southampton
Hampshire SO30 3RX
Tel/Fax: 023 8013 9805
e-mail: info@forum-pr.co.uk

NEXT EDITION
Review of Sheffield meeting
Preview of the Winter meeting

DIARY DATES 2008

25–28 November – ICONIP 2008, 15th International Conference on Neural Information Processing, Auckland, New Zealand.
http://www.aut.ac.nz/iconip08/

9–11 December – AI-2008, 28th SGAI International Conference on Artificial Intelligence, Cambridge, UK.
http://www.bcs-sgai.org/ai2008/

15–19 December – ICDM 2008, IEEE International Conference on Data Mining, Pisa, Italy.
http://www.icdm08.isti.cnr.it/

Mid-January – NCAF meeting (theme and venue to be confirmed). For information, email enquiries@ncaf.org.uk or telephone +44 (0)1332 246989