ncaf
Natural Computing Applications Forum
Networks

An invitation from intelligent robots

Centre for Hybrid Intelligent Systems, University of Sunderland, 21–22 January 2009.

The Hybrid Intelligent Systems group has developed robots that are capable of natural behaviour in everyday real world situations.

The first NCAF conference of 2009 will be held at the University of Sunderland, and hosted by the Centre for Hybrid Intelligent Systems (HIS) and the NESTCOM project. The event will be held at the Sir Tom Cowie Campus at St.Peter's along the banks of the River Wear with a special evening event at Sunderland's National Glass Centre. On arriving at St.Peter's you will enter into the award winning David Goldman Informatics Centre, which has been described by some students as a 'computing cathedral'.

The focus of this natural computing application forum meeting will be Intelligent Robotics. The Hybrid Intelligent Systems group has an international reputation in this area and has produced award winning intelligent interactive robotic systems. The group has participated and co-ordinated several European Union and several Engineering and Physical Sciences Research Council (EPSRC) projects including EmerNet, Micram, Mirrorbot, Mira, and NESTCOM. The HIS group has developed robots that are capable of natural behaviour in everyday real world situations. This involves investigating many cognitive tasks, for example, navigation, language processing and visual recognition. In a number of projects we use our cognitive robots to study learning and communication in real world environments. The foundation and motivation for much of our research comes from natural systems such as biological systems, neural systems or cognitive systems. We use these foundations in order to build more sophisticated adaptive interaction systems, learning agents, self-organising information systems and robotic engineering systems.

For instance, MIRA is our robot head that we use to research into principles of neural network learning, human-robot interaction, language parsing and understanding, vision and image recognition, auditory localisation and basic emotions. It is capable of head movements, generating emotional features through expressions, limited word generation and recognition and is also able to hear, localise and track sounds. Recently we have implemented new approaches to neural learning of sound localisation or tracking, i.e. turning the head to the sound, limited word recognition and generation within restricted dialogs and emotion generation (via movements of mouth and eyebrows).

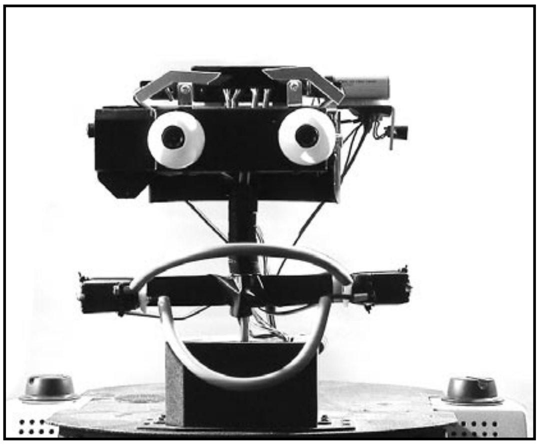
Centre for Hybrid Intelligent Systems, University of Sunderland

The MIRA robot head

The NESTCOM project focuses on bringing together European NEST projects under the initiative 'What it means to be human', to research on multimodal communication. These NEST projects have produced a substantial amount of specific but sometimes isolated knowledge, exploring in particular verbal and visual communication media, which includes such important topics as learning by imitation, studying the neural origins of language, exploring the evolutionary origins of human mind, and researching into verbal and non-verbal communication.

We look forward to welcoming you to the area and introducing those new to Intelligent Robotics to some great talks and current research from students and our departments.

Stefan Wermter
University of Sunderland

INSIDE

BACK PAGE

Sheffield Systems under the spotlight

Much was made at the meeting about the application of pattern recognition and other machine learning techniques to safety critical systems, particularly in the aerospace industry.

The autumn meeting of NCAF was held in Sheffield, hosted by the Department of Mechanical Engineering in the Turner Glass Museum. Attendees enjoyed two days of stimulating discussion about a host of natural computing topics, with a running theme of dynamics and dynamic systems.

Dynamic systems have long been of interest to the practitioners in the field of neural and natural computing, from the advent of tap delay line neural networks through to Gaussian Process Dynamics. The problems of system identification, prediction and control are interesting ones, and this meeting was devised in order to explore the subject within NCAF's usual friendly environment.

Vibrating structures

The Dynamics Research Group at Sheffield has been active in the area of applying computational techniques to dynamic systems for some years, and it was Professor Keith Worden, former head of the group, who opened the meeting with a discussion on 'Aspects of Uncertainty in Dynamic Systems'. Keith's work involved the classification of damaged and undamaged vibrating structures using Multi-Layer Perceptrons (MLP), and examining the robustness of the trained networks using interval techniques. He also discussed how uncertainty in a system can be modelled (the most popular choice currently being probability theory), how it propagates and how different uncertainties can be fused.

Much was made at the meeting about the application of pattern recognition and other machine learning techniques to safety critical systems, particularly in the aerospace industry. As Steve King of Rolls-Royce explained, this is in part because of a move towards `power by the hour' and similar industry shifts, where the manufacturer takes responsibility for the serviceability of the product through its life. Steve explained the core aspects of Rolls-Royce's Engine Health Management (EHM) system, and how pattern recognition could be used. A key point of Steve's talk, which provoked much discussion amongst the NCAF delegates, was that machine learning techniques must provide information in a manner that is easily interpretable by the engineer.

Geomagnetic storms

After lunch, the focus moved toward the tracking aspect of dynamic systems research, as Rob Harrison (University of Sheffield) discussed Geomagnetic Storm prediction via the Kalman Filter. Rob gave a pedagogical review of the Kalman filter, as well as some interesting insight into the physics of geomagnetic storms, and the disastrous consequences of an undetected storm in terms of power outages and loss of communication systems.

Mike Dewar from Edinburgh University gave an overview of the iBehave informatics project. This involves many aspects of natural computing including image processing, tracking and classification. The project aims to provide a system capable of interpreting videos for use in the drug discovery pipeline. Mike's entertaining videos of mice and flies captivated the delegates with real-time `annotation' of the subject matter.

The day ended with a brief foray into information theory, as Maurizio Filiponne (University of Sheffield) presented some recent research into Information Theoretic Novelty Detection. The parallels between the machine learning field and that of information theory have been explored before, (see e.g. Mackay 2003, Information Theory, Inference and Learning Algorithms), and Maurizio showed that a robust novelty detection algorithm could be devised in this manner.

The social event took place in the evening, starting with a tour of Kelham Island brewery, with the intention of perusing the area's fine selection of pubs and ending in an Indian restaurant. The good people of Kelham Island proved so hospitable, however, that the attending delegates stayed for some hours before deciding that it was time to stop drinking and start eating. After much ensuing merriment, a venerable banquet of curry followed in the Seven Spices Balti, washed down with yet more beer (or an excellent mango latte for those who had already made good use of the brewery's hospitality).

Gaussian Processes

Neil Lawrence started the second day of talks with a presentation entitled 'Gaussian Process Dynamic Systems'. Neil covered some of the basics of the Gaussian Process (GP) methodology, and then described how a covariance function can be formed from differential equations. The result is a normal distribution whose samples appear to describe a dynamic system. Using the usual GP inference methods, one can perform inference about parts of the dynamic process from observations. The presentation raised interesting points about the use of machine learning tools in conjunction with expert knowledge of the system being modelled. It was interesting to see that new methodologies brought to machine learning (like GPs) can be extended in different directions.

The theme of Gaussian Processes continued as Grigorios Skolidis (University of Sheffield) discussed the classification of arrhythmic heart beats. This is a popular application area at NCAF. At the last meeting, Iain Strachan presented some work using Hidden Markov models for a similar task. Greg's work involves using discriminative GP classifiers to distinguish healthy and unhealthy beats.

Fatigue monitoring

Steve Reed of the Defence Science and Technology Laboratory discussed the use of neural networks in fatigue monitoring of military aircraft. Steve described the lengthy process that it has taken in order to persuade safety experts that a neural network approach was a safe one. This included simplifying the neural network in order to make its feed-forward function easily tractable, as well as education of the appropriate people in the field of neural networks, seen for so long and by so many as an untrustable black box. The application involves assessing the level of strain in a part of the aircraft without monitoring it directly, but simply by using the flight data that are routinely recorded by the aircraft throughout the flight. Both in-flight and grounded scenarios provided training data. The result is a small neural network which can give an accurate history of an aircraft's loading regime, which can then be used to assess the fatigue-related airworthiness of the aircraft. The success of Steve's work (the system is currently flying on a number of aircraft) will be an encouraging sign to researchers in the field of applied machine learning.

Following lunch, Professor Gary Green of the York Neuro-imaging Centre discussed non-parametric

Using diffusion to find patterns of neuronal responses

The use of Bayesian methods in imaging neuroscience is now well established, in particular for functional magnetic resonance images (fMRI) and source reconstruction, given magneto or electro-encephalographic (MEG or EEG) data (for example, work by Friston's group at University College London, some of which will be mentioned later). In general, the aim is to find a configuration of parameters representing the brain's response to experimental perturbation that minimizes an energy function given a model. This is equivalent to maximizing the probability of data, given the same model. Given that data are generated by spatially extended brain structures, an explanation of the data requires a spatial model, i.e. one that predicts data at one point in the brain (voxel) in terms of responses within a local neighbourhood. These include Gaussian Process priors, which provide constraints over the space of functions representing neuronal responses through a parameterised covariance function that includes the spatial extent of dependence, in other words its smoothness.

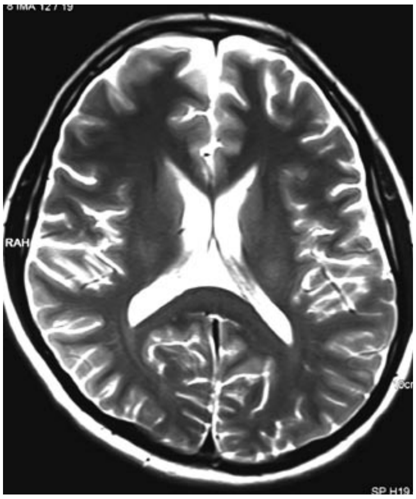

A magnetic resonance image of the brain during ataxia', from www.japi.org/august_2008/c_628.html

Of particular interest is a subset of covariance functions based on diffusion, which describes the effect of random perturbations of an ensemble of particles. The diffusion kernel, i.e. the object that propagates a solution of the diffusion equation from one point in time to another, can be used to represent covariation between the brains response at different voxels, whose neighbourhood depends on the time of diffusion. This provides a general framework for discrete and continuous formulations, where a graph-Laplacian is defined over an arbitrary, i.e. irregular, graph and the Laplace-Beltrami operator defined over, in general, a non-flat continuous space. The Laplace matrix (or operator for continuous domains) takes centre stage as it provides a problem dependent basis set on which vectors (or functions) can be represented. In particular, the diffusion kernel can be used to formulate either a global or local basis set.

The eigensystem of the graph-Laplacian provides a global basis, where eigenmodes of increasing spatial frequency have eigenvalues of increasing magnitude. The diffusion kernel is a function of this eigensystem, which can be decomposed into a weighted sum of components. Each component is the outer product of an eigenmode weighted by the negative exponential of its eigenvalue scaled by the diffusion time. This means that as time increases high frequency modes are down weighted, which leads to the characteristic effect of diffusion, i.e. dispersion. Optimisation of the diffusion time provides a data-driven way to select eigenmodes, when used to model neuronal responses, and has been used to preserve the spatial structure of neuronal responses from high-resolution fMRI data. The alternative is to smooth the data with a fixed Gaussian kernel. However, this leads to blurred images.

A local basis can also be formed from columns of the diffusion kernel. For a fixed diffusion time, each column provides a 'bump' function, one for each point in the brain whose spatial extent is determined by the diffusion time. A weighted sum of outer products of these can then be used to construct a spatial covariance matrix as above. When combined with priors over the weights this facilitates the selective removal of covariance components (similar to Automatic Relevance Determination) from regions of the brain where there is no evidence of response. The crucial difference here is that components are localised in space and not global as with eigenmodes. This provides a way to learn sparse representations of neuroimaging data, which has been implemented in an approach to source reconstruction of MEG and EEG data called Multiple Sparse Priors. This can be generalised by using columns of diffusion kernels at different diffusion times to provide a multiscale and overcomplete basis set from which a sparse representation can be selected. Similar ideas have been developed by Coifman and colleagues to construct compact, multiscale representations using ideas from diffusion to generalise wavelets to irregular graphs and non-flat continuous spaces, called diffusion wavelets.

In short, diffusion provides a general framework to inform flexible representations of the sources, or causes, of data we observe that are multiscale, selective and sparse. This application mirrors key concepts in the work of Olshausen and others based on the sparse encoding of retinal images to explain spatio-temporal receptive field properties of cortical simple-cells and the nature of their response.

Lee Harrison
York Neuroimaging Centre, University of York

The article describes work undertaken at the Wellcome Trust Centre for Neuroimaging.

system identification, with a view to locating and characterising activity in the brain. On a further medical note, the meeting ended with a discussion of the application of neural networks to image recognition problems in the medical field, presented by Nick Granville.

'Puzzle Corner' aims to provide light relief (as well as educate) at NCAF meetings, but it was not so for bewildered NCAF newcomer Maria Pavlou. She was coaxed onto the stage in order to assist 'Fenella the Rottweiler' with the meeting's puzzle. The puzzle involved betting on a series of rounders games, with the objective of returning a pre-defined profit with a constraint on the maximum loss. The problem turned out to be fully constrained, which proved a furore amongst NCAF members who had placed hypothetical money on the hypothetical game, and then demanded a hypothetical refund.

NCAF's autumn meeting was an excellent experience due not only to the attending speakers, but also to the thought provoking questions which their presentations stirred in the audience. The Turner Glass Museum hummed with the sound of friendly discourse, as the delegates dissected and consumed the natural computing topics of the moment.

James Hensman
University of Sheffield

PUZZLE CORNER

Number 40

The CEO of Magnetic Humanities thought he was onto a winner with his new American Football mini-game for the Y-Box. Charging £39.95 for what was simply a look-up table payoff matrix was marketing genius. He was asking Lisa to do some last minute play testing.

The 'game' consisted of the Offense secretly selecting one of 4 offensive plays (designated W, X, Y and Z) whilst the Defense secretly selected one of 4 defensive plays (designated A, B, C and D). The payoff matrix determined the number of yards gained by the Offense, with negative numbers indicating a loss. The matrix entries for row W were -5, -1, +3, +3 for the corresponding defensive plays A, B, C and D; for row X they were -1, +5, -1, -5; for row Y they were +5, +5, -5, -5; and for row Z they were +5, -5, -1, +3. Hence, for example, the play combination ZC would result in a one yard loss for the Offense.

The CEO, pointing out that the sum of the matrix entries was zero, claimed the game was fair (i.e. that the expected average gain per play was zero if either or both adopted an optimal strategy). However, Lisa pointed out that this was just coincidence, and that by changing only two numbers in the grid whilst keeping the overall sum at zero, the expected average gain per play became almost half a yard.

Undeterred, the CEO asked Lisa to analyse the game to determine the likelihood that the Offense could achieve a 'first down', that is, gain 10 or more cumulative yards within the first 3 plays. Lisa did and concluded that, using this table, 'first downs' were going to be like rocking horse manure – pretty rare.

Given optimal play by both sides, what was the likelihood of achieving a 'first down'?

The answer will be given at the next NCAF meeting (21–22 January 2009, University of Sunderland).

Fenella the Rottweiler

Natural Computing Applications Forum

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

MEMBERS' NEWS AND VIEWS

Deadline for contributions for the next edition – 1 March 2009.
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 the Sunderland meeting
Preview of the Swansea meeting

Computerised agents to cope with disasters

A system that will use networks of computerised agents to cope with disaster scenarios such as outbreaks of fires was outlined by Professor Nick Jennings on Guy Fawkes night (5 November). Professor Jennings from the University's School of Electronics and Computer Science chose the busiest evening of the Fire and Rescue Service's year to tell a select audience at Park Centre, Farnborough, about ALADDIN, an ambitious £5.5 million five-year research programme aimed at developing computerised agents for use in disaster recovery or terrorist attacks. The project has just reached its halfway point.

ALADDIN stands for Autonomous Learning Agents in Decentralized Data and Information Networks.

Professor Jennings believes that the UK Fire and Rescue Service is a likely potential client for some of the applications already developed by ALADDIN. "We are developing decentralised information systems that can continue to operate effectively when there is a fire or in other extremely difficult circumstances," said Professor Jennings and continued "We use computerised agents which can sense, act and interact in order to achieve individual and collective aims. Central to this endeavour is the effective coordination of the different actors and, to this end, we have developed a rich series of algorithms for inter-agent co-operation and negotiation."

ALADDIN is one of a number of programmes that bring leading academic groups to key industrial challenges through the BAE Systems/Engineering and Physical Sciences Research Council (EPSRC) Strategic Partnership, with the objective of delivering high-quality research and enhancing industrial capabilities. The ALADDIN team comprises internationally leading researchers in complex adaptive systems from the Universities of Southampton, Bristol and Imperial College; in fusion, inference and learning from the University of Oxford and Imperial; and in decentralised architectures from BAE Systems.

One of the challenges facing the researchers is to bring together work from a number of previously distinct fields, such as information fusion, inference, decision-making, and machine learning. This work then needs to be combined with work from multi-agent systems, game theory, mechanism design and mathematical modelling of collective behaviour in order to give a collective view on behaviour.

"As ever more information sources become available – through environmental sensors, intranets, and so on – the problem of obtaining and fusing the right information when making decisions and taking actions is becoming increasingly pressing," said Professor Jennings.

For further information about ALADDIN, visit: www.aladdinproject.org/

Based upon an article from the University of Southampton.

2009 Annual General Meeting

University of Sunderland, 22 January 2009

The Annual General Meeting of NCAF will be held during the Sunderland meeting. This is your opportunity to voice your views to the committee. Even better, you can join the committee and help to run the organisation. Existing committee members can explain what we do, or contact the Chairman, Graham Hesketh.

Nick Granville, Editor

DIARY DATES 2009

15–17 April – EvoBIO 2009, 7th European conference on evolutionary computation, machine learning and data mining in bioinformatics, Tubingen, Finland
http://evostar.na.icar.cnr.it/EvoBIO/EvoBIO.html

20–21 May – NCAF meeting at Swansea University. The theme will be the Grand Challenges in Information Driven Healthcare. For information, email enquiries@ncaf.org.uk or telephone +44 (0)1332 246989

13–16 July – BIOCOMP'09, the 2009 International Conference on Bioinformatics and Computational Biology, Las Vegas, USA.
http://www.world-academy-of-science.org/worldcomp09/ws/conferences/biocomp09

27–29 August – EANN 2009, 11th International Conference on Engineering Applications of Neural Networks, London, UK
http://www.uel.ac.uk/eann2009/