Networks

Human-Computer Interaction at Cambridge

3 – 4 September 2003

Many interaction modalities with computers will be discussed, ranging from user interfaces for digital photography, hands free interfaces to Brain Computer interfacing, computer vision, and computer imaging.

Human-Computer Interaction (HCI) will be the theme of this year's autumn meeting, which will be held in Cambridge on 3 and 4 September. This is the first joint meeting of NCAF and Microsoft Research, which is also generously sponsoring the event. We have defined HCI in the widest sense, covering not only technical aspects of humans interacting with machines, but also philosophical, design and psychological aspects of HCI.

Many interaction modalities with computers will be discussed, ranging from user interfaces for digital photography (K. Wood, Microsoft Research), hands free interfaces (D. MacKay, University of Cambridge), to Brain Computer interfacing (P. Sykacek, University of Oxford), computer vision (A. Criminisi, Microsoft Research), and computer imaging (P. Perez, Microsoft Research).

Each day will commence with a tutorial. On the first day, Prof. R. Spence (University of Cambridge), will cover the principles of Human Computer Interfacing. Prof. Chris. Bishop (Microsoft Research) will be giving 'An Introduction to Mixture Models, EM, and all that' to begin the second day of the meeting. Both are well known in their respective fields and are excellent speakers, promising an exciting start on both days.

Other presentations will include reports on progress of the VIBES project (Variational Inference Engine for Bayesian Networks) by J. Winn (Microsoft Research), job-shop scheduling using software agents by P. Swann, (Rolls-Royce) and co-operative behaviour of complex systems by Ian Nabney (University of Aston). In short, the autumn programme promises to contain many exciting talks on a large variety of HCI tools by well-known experts in their fields. If this is not enough we have many presentations on non-HCI related issues which will reach the parts that other talks could not reach.

Dr. Iead Rezek
University of Oxford

Responding to the survey

In order to return to some of the core elements of NCAF's original remit and in response to returns from the recent *Networks* readers' questionnaire, it is proposed that the Winter 2003/4 NCAF meeting will be themed on *Fundamentals of Natural Computing*. It is hoped that during the first day there will be a tutorial on the fundamental techniques used in machine learning and artificial intelligence applications. To consolidate this, a workshop on one or more of these techniques is planned for the second day. This is an ideal opportunity for people new to Natural Computing techniques to be introduced to the area. It may be particularly attractive to students planning to utilise computational methods in their research.

The provisional date of the meeting is 7–8 January 2004 to fit in with the academic calendar and the industry post-Christmas lull. The basic details of the meeting, including the majority of speakers, will (honest!) be available for members to peruse at the September meeting and the finalised agenda will be available by the end of that month…watch this space!

'Bert Bullen
BAE Systems

Natural Computing is Safe in Our Hands

…the highlight of Graham's day, as he transformed himself into the mystery man of natural computing, the Great Enchilada, and entertained us with his latest trick of digital prestidigitation and managed (after an inordinate amount of faffing, assistance from his accomplice, Guacamole, and complete suspension of disbelief by the audience) to predict the value of a hidden card from a pack.

4 June 2003. Outside lane of the M5. Smile on my face. All the hassles of the world left behind me. Yes it's time for another couple of days in the invigorating world of the latest NCAF meeting. I have to confess I always look forward to these events, with their mix of technology, ideas and people, all set against a relaxed, friendly background. (To see just how relaxed, see later under 'Social Event…').

The June event was held at Exeter University's Crossmead Conference Centre, and Richard Everson of Exeter University had put together an appetising mix of papers and events. As usual with NCAF meetings, the event had a mix of themed and general papers. The theme for June's meeting was 'Safety Critical Systems', a domain offering rich pickings for the natural computing world. The use of neural networks in safety critical systems has often been considered an anathema, and I was interested to see how these two often diametrically opposed and incompatible, worlds could be reconciled.

The meeting opened with a talk on 'Bayesian Reverse Jump Markov Chain Monte Carlo' (RJMCMC) Methods for Safety Critical Systems' by Wojtek Krzanowski (University of Exeter). Wojtek outlined work aimed at developing classifiers that could also deliver measures of the confidence of their classifications, as required by safety critical and safety related applications. Building on datasets from aircraft tracking and trauma treatment domains, Wojtek and his team had developed techniques based on RJMCMC to improve a KNN classifier, and early results were encouraging. Further work was under way, addressing enhancements to distance measures and feature detection.

The next paper, by Jonathan Fieldsend, (University of Exeter), addressed the problem of 'Tuning Legacy Critical Systems with Multi-objective Optimisation'. Jonathan's talk described work underway to tune an existing application, Short Term Conflict Alert (STCA), which is in use by many national air traffic services to predict when two aircraft appear to be likely to get too close. Jonathan's work uses multi-objective evolution strategy to explore the boundary of locally optimal solutions, generating an achievable Receiver Operating Characteristic (ROC) curve, and so allowing an informed selection of operating point, trading off likelihood of true positives against false positives to best suit the application. Initial work on trial databases had proved encouraging, and the team are starting to address STCA optimisation.

After coffee Paulo Lisboa (Liverpool John Moores University) presented a paper titled 'Confidence in Neural Networks: Methodological Issues Arising from a Review of Safety-Related Applications'. Paul's talk looked beyond the technical issues of developing safety critical systems, and addressed all the related issues of making successful applications, including good practice throughout the software development cycle.

Safety Cycle

After a splendid lunch, and plenty of chances to chat, make new contacts and renew old friendships, we resumed with a talk from Zeshan Kurd from the

High Integrity Software Engineering department at the University of York. Zeshan's talk, entitled 'Developing Artificial Neural Networks for Safety Critical Applications', included a discussion of the importance of addressing the safety lifecycle in safety critical system development. He described a hybrid ANN design approach that had been developed at York, which enabled some of the advantages of neural networks to be delivered within the strict constraints of safety critical development.

Zeshan was followed by Alastair Faulkner (CSE Limited & University of Warwick), who raised the issue of data integrity in safety related systems. Alastair pointed out that when we are developing and managing systems, it is too easy to concentrate on the integrity of the design, verification and validation of the software, and forget that many applications are critically dependent on the data they use, and quality, development, maintenance and assurance of the data is often outside the control of the developer or user. He recommended that the data used in and by safety related systems should be subjected to similar hazard and risk analyses that we apply to the system design and that attention should be paid to the structure and validation of data to ensure that potential errors are uncovered and resolved. In addition to an excellent talk, Alastair also managed the previously unknown achievement of finishing with a quote from Hiram K Hackenbacker, one of the best-known technical brains of the last century! (From *Thunderbirds*, for those who had forgotten.)

After tea, two NCAF stalwarts Mandy Bradley and Chris Satchwell (Technical Forecasts Limited), presented a general paper on the subject of 'Forecasting Values of Commercial and Residential Property Using Non-Linear Mathematical and Statistical Techniques'. As well as describing the approach that Technical Forecasts have been taking to forecasting, employing an array of multiple models and data sources, Mandy and Chris presented an important lesson on dealing with errors and uncertainty in source data.

A short break allowed me to slip in a quick advert for the work I am undertaking concerning the non-technical aspects of developing AI based systems. (For more details see www.ai-research. org.uk).

Next – the highlight of Graham's day, as he transformed himself into the mystery man of natural computing, the Great Enchilada, and entertained us with his latest trick of digital prestidigitation and managed (after an inordinate amount of faffing, assistance from his accomplice, Guacamole, and complete suspension of disbelief by the audience) to predict the value of a hidden card from a pack. Luckily, it ended before rigor mortise took hold.

That completed the formal part of the day – all in all a great success – not too many equations and a good focus on real practical issues. So it was off to the social event – canoeing down the river to one of Exeter's best pubs, the Double Locks. But… the weather was against us, and instead of a sun kissed evening enjoyed against the sound of the quiet

splash of paddles in the calm water, we were faced with the prospect of rain and Richard had made the sad but sensible decision of cancelling the canoes. Undaunted, we set out for the pub by taxi, only to find that the weather was chilly but dry all evening! No problem! We had a great time, quaffing the Double Locks' excellent beer and wolfing down large portions of their home cooked food. Once we were satiated we walked back into town to the bowling alley, and started the second event of the evening! Here we discovered that Mark 'Bowling Balls' Cheeseman was a bit of an expert (casually tossing in that he used to bowl for England or something like that). Mark wiped the floor with us in the first game, but we were pleased to see that his skills miraculously evaporated in the return match. A short walk took us back to the conference halls of residence and to well earned sleep (except for a diligent few who undertook the advanced card techniques school.)

Active Learning Algorithms

Day 2 dawned all too early, but we recovered over a slow and excellent breakfast, before starting the serious business of the day. The first paper was presented by Colin Campbell (Bristol University). Colin presented his work on active learning algorithms, in which a learning machine actively selects unlabelled datapoints from the data set and requests their labels. He showed how this can improve speed of learning in many tasks, and outlined work on extending the principle to regression, using Gaussian Processes.

After a break for coffee, Mike Davies (Queen Mary, University of London), presented an overview of ongoing work in the area of Audio Source Separation and ICA. He outlined methods for blind source separation, where two audio sources recorded by an array of sensors can be separated without any knowledge of the source characteristics. The technique employed a frequency domain framework and Mike explained how to get over the problem of selecting the right ordering of sub-solutions (the 'permutation' solution). A practical

demonstration and a highlighting of the problems of the sensitivity of the system to source/sensor position rounded off an excellent talk.

The last speaker of the morning, Ali Hessami (Atkins Global), treated us to a fascinating and thought provoking presentation entitled 'Ancient Wisdom, Modern Solution'. Drawing on the ideas of Zarathustra, Heraclitus and Tao, Ali described Weighted Factor Analysis (WeFA), a new approach to problem identification, exploration and resolution, based on a guided hierarchical decomposition of a problem, and allowing all aspects to be considered. WeFA has been used in anger on real world problems, and continues to be developed into a practical tool.

After lunch, the final two speakers were both from Oxford University. Natasha Chia covered the topic of 'Neuromorphic Engineering', whose design principles and architecture are biologically inspired (a concept introduced by Carver Mead). The ultimate goal is the realisation of future thinking machines, and the potent blend of neuroscientists, electronic engineers and mathematicians are hoping to deliver this through more closely modelling the spiking behaviour of real neurons. The meeting culminated in Rizwan Choudrey's 'Beyond Independent Component Analysis: Hierarchical Models' talk, in which he not only covered some serious technical material with admirable clarity, but he also managed to (almost) seamlessly integrate a collection of randomly chosen words (Zephyr, Anaconda and Afghanistan) into his talk, for which he deservedly received a mini-Mexican wave from the front row!

All in all, an excellent meeting. Thought provoking and entertaining, and as ever an excellent opportunity to discuss natural computing and similar issues with friends and colleagues in a relaxed environment. Well done to everyone who contributed to the event.

See you in Cambridge!

Richard Ellis
Stratum and BCS SGAI

PUZZLE CORNER

Number 24

Being pushed out of a helicopter hovering six feet above a minefield in the middle of a pitch-black night was not Lisa's ideal way of ending the evening. Nevertheless, that's where she found herself. A sandy strip a mile wide and many miles long, Lisa knew it was the no-mans-land between the elephant grass of the Hellarwi pigmy tribe and the verdant pastures of the bull-worshipping Nagumba.

Out of a warped sense of fair play, the helicopter crew had given her a compass, a pathetically weak flashlight and a yardstick. Using these tools Lisa was capable of accurately following any desired path. Unfortunately she had no idea how the minefield was geographically oriented, so choosing the best direction to get her out of danger quickly was not going to be straightforward (e.g. walking due North until you cross the edge of the strip is not going to work if you're unlucky enough to find the minefield runs North-South). On top of this, the torch could only illuminate up to six feet away so she was going to have to get very close to the edge to spot it.

"Hey, Henri, what's the French for BA-BA-BOOM?" said the pilot as the chopper dusted off, leaving Lisa to contemplate her options. Choosing to minimise the total number of yards she would have to walk in the minefield in the worst-case scenario, she set off on her optimal route.

What is Lisa's walking strategy and how many yards at most would she have to cover in order to guarantee leaving the minefield?

The answers will be given at the next NCAF meeting (3–4 September 2003, Fitzwilliam College, Cambridge).

Fenella the Rottweiler

Researching complexity in clusters

As part of the Novel Computing Initiative for 2003/04, EPSRC recently awarded funding to 16 research clusters in the general area of complexity. One of these research clusters, 'Inference and co-operative behaviour in large interacting systems', is led by the Neural Computing Research Group, Aston University. Members of the research cluster include researchers from Birmingham University, Btexact, King's College London, Oxford University, Reading University / JCMM (Met Office), Royal Holloway, Tokyo Institute of Technology and the University of Florida.

Background

Understanding the behaviour of large complex non-linear systems is of great importance to the analysis and handling of information in modern society. These systems are characterised by a synergetic emergent behaviour of a different nature to that of a similarly large but uncoupled or linearly interacting system. Due to their complexity, especially the type and number of interactions between components, deterministic approaches to analysing such systems are doomed; all feasible approaches rely on a probabilistic definition of the problem using techniques from statistics and statistical physics.

Very large complex systems are characterised by multiple time and resolution scales. Microscopic interacting objects self-organise to construct mesoscopic weakly interacting agents, combined with order parameters which lead to a rich diversity of macroscopically observable phenomena, such as coherence, phase transitions (drastic changes in macroscopic behaviour at specific critical values of the model's parameters), and structural chaos. These systems are driven by non-linearity and, very often, stochasticity; the latter comes from the natural characteristics of the problem, inherent uncertainties and statistical behaviour, as well as from the practical difficulties of modelling and extracting information at the microscopic level. By extending techniques from statistical physics, synergetics, many-body theory, information theory and pattern analysis we wish to address issues in the behaviour of complex systems at each of these resolution scales.

This research cluster brings together expertise from a range of disciplines in order to develop a road map for research in this area and generate a research programme to address the following issues:

- devising feasible inference methods for very large connected systems;
- analysing typical system behaviour in specific instances of large complex systems under given disorder (e.g., topological constraints) and a dynamically changing environment;
- using the insight gained from the analysis to support local decision making as a method for controlling and directing global network properties

Our long term aim is to move from simulation and modelling, which are typically infeasible in large complex systems, to the extraction of higher level variables that allow us to monitor, predict and design such systems. We suggest the exploitation of similarities between many-agent systems such as nervous and immune systems, environmental models, future Internet, ultra-device chips and communication networks and develop principled information processing approaches to describe the resultant behaviour, so that we can design and forecast better for the future.

Research Challenges

One of the main difficulties in analysing large interacting systems is how to extract useful top-level information from the deluge of microscopic data, which tends to be highly fluctuating and very sensitive to changes in local conditions. Using simulations and/or algorithmic solutions of such systems is infeasible as there are no practical deterministic methods to search solution spaces that grow exponentially with the system size.

We anticipate that our research programme will focus on macroscopic and mesoscopic methods as they provide a top-level understanding of the system's behaviour and can identify the emergence of global characteristics. Practical inference methods will also be devised based on the insight gained from the macroscopic analysis; these are essential for obtaining specific solutions in complex systems, in the absence of feasible algorithmic approaches.

The main objective of the cluster is to stimulate interdisciplinary research in the general area of inference and co-operative behaviour in large complex systems. This will be achieved through interaction between theoreticians, researchers and end users, leading to an integrated and coherent research effort. The challenges facing the research cluster, beyond integration, come at several domains: (i) setting up generic theoretical frameworks for the various problems; (ii) carrying out a full analysis of the generic models; (iii) exploiting the insight gained from the analysis for devising feasible inference methods; (iv) employing the inference methods for realistic large scale applications.

Our strength is in the generic nature of our analytic approach, which is not limited to one domain, but can be employed to study realistic large scale interacting systems of many types. Therefore we propose to develop our research programme across three broad domains, which will then motivate the specific projects. These broad domains are: microscopic inference in large connected systems (weather, environment and data mining), distributed communications and computing networks and multi-agent system dynamics (from ultra devices to the immune system).

Planned Activities

The planned activities aim to join theoreticians, application developers and end users to forge a consensus on the best directions for further research in order to overcome the challenges in this domain. In the six-month course of the programme we shall run activities to gather information and ideas, analyse and evaluate them in open forum, and then disseminate the results back to the research and user communities.

The main activities of the cluster will be two workshops: a kick-off meeting held on the 17 July and a second meeting, open for all, towards the middle of October.

We see outreach and community-building as very important components of this project, and intend using both active and passive activities to promote wider involvement. We would like to get in touch with end users to find out their needs and interests to help direct our research towards useful long-term goals. We would like to invite interested NCAF members to participate in our activities and to let us know of their interests.

David Saad Aston University

MEMBERS' NEWS AND VIEWS

Deadline for contributions for the next edition – 6 October 2003. Please send to: Managing Editor – Inspector Rick Adderley, e-mail: Insp1908@aol.com or r.adderley@west-midlands.police.uk