

Fundamentals of Natural Computing

Back₆ at Bath 21–22 January 2004

The overall aim is to encourage those new to the field to understand the practical benefits that natural computation offers and to consolidate the fundamental knowledge of existing members.

One of the key ideas behind NCAF has always been to inform and introduce new people to the idea of computational intelligence. In fact to quote from the NCAF mission statement, the forum is:

“...suitable for people new to the field who want to know how computational intelligence techniques may affect us and our businesses. There are ample opportunities to meet and have informal discussions with technology providers and more importantly with other customers, so we can quickly put ourselves in the position of being informed...”

It is the intention, therefore, that the themed day of the next meeting will be dedicated to fulfilling this part of NCAF's remit. Overlooking the beautiful and historic city of Bath, the University campus at Claverton Down will host the next meeting on the 21 and 22 of January 2004 (and not 14/15 as previously advertised). The overall aim is to encourage those new to the field to understand the practical benefits that natural computation offers and to consolidate the fundamental knowledge of existing members. This should enable those delegates who do not routinely use these techniques, to become more confident in applying them to suitable problems in their area.

To this end there will be a wide variety of tutorial-flavoured talks to cover all the main categories which we feel are encompassed by natural computing: Andy Webb (QinetiQ) will be discussing a broad spectrum of pattern analysis techniques, Tom Khabaza (SPSS) will address data mining and Steve Roberts (Oxford University) will talk about the intelligent use of neural networks. In addition, we are pleased to have Prof. Nigel Shadbolt (Southampton University), a leading light in the field of knowledge engineering, management and associated topics. The afternoon session will

include a workshop, during which it is hoped a free exchange of ideas will occur. Delegates are encouraged to put forward their own case studies for discussion. Hopefully, the expertise in the room will be able to provide additional insight into how real problems can be usefully tackled using natural computation.

To complement the fundamental material covered on the first day, the second day will be targeted at a more general audience. Speakers will include Dave Cliff (HP Labs), Carren Holden (BAE SYSTEMS) and Tom Musicka (University of Newcastle).

As well as contributing to the technical programme, BAE SYSTEMS is also generously supporting the Bath meeting, as they are keen to encourage the use of natural computing techniques in addressing practical problems.

The social event, to be held on the second evening, will be a pleasant and genteel soiree, involving dinner in the refined Georgian splendour of the city and maybe even with a surprise or two thrown in.

Thanks to the fundamental nature of the subject matter and the quality of the speakers that have volunteered to support this event, the Bath meeting will provide an excellent opportunity to introduce new starters to the field of natural computing. It should prove stimulating, varied and entertaining and we hope many of you will be able to join us in January.

Don't forget, the NCAF Annual General Meeting will also be held during the Bath meeting and NCAF is currently seeking nominations for vacant committee positions from existing members. For further information please contact Mark Cheeseman, NCAF Secretary.

'Bert Bullen
BAE SYSTEMS

INSIDE

- Cambridge in the Sun
- Why do AI based projects fail so often?
- Postscript to a furniture mover's nightmare
- Puzzle Corner No. 25

Back Page

- Data Rich. Information Poor?
- Diary Dates

Cambridge in the sun

Blender uses the solution to a suitable Poisson equation to clone and edit images without edge effects. For example, if you want to add the sun to a picture, this can be done without interfering with the sky around it.

In glorious sunshine, NCAF held one of its best ever meetings in Cambridge in early September. The theme, Human-Computer Interaction, meant that we heard a wide range of talks on new topics.

The first speaker, Bob Spence (Imperial College), spoke about designing new modes of interaction. He was behind MINNIE, a fully working CAD system with light pen interaction, that ran with 16Kb of RAM in the mid-70's. He is now researching new methods of presenting information and demonstrated the value of non-functional prototypes (i.e. mockups to show how the interaction would work if all the behind the scenes engineering were in place).

Patrick Perez (Microsoft Research) demonstrated two image surgery tools: Patchworks (a region tiling algorithm for removing objects) and Blender (guided interpolation for seamless cloning and editing). Patchworks intelligently copies neighbouring blocks into regions of an image, using auxiliary information (such as depth of field) to improve results. This automates tasks that can be performed manually but with extreme difficulty. For example, scratches, and overlaid text can all be removed. It is also possible to return photos taken at crowded tourist sites to their empty magnificence. Blender uses the solution to a suitable Poisson equation to clone and edit images without edge effects. For example, if you want to add the sun to a picture, this can be done without interfering with the sky around it. The tool is particularly useful for cloning objects with complex boundaries.

Ken Wood (Microsoft Research) described new user interfaces for managing digital photographs. Technically the most interesting aspect was using neural networks to classify coherent regions (c. 100-200 per image) to provide a search facility by visual context. However, users favour browsing over searching, and the most important feature of the interface was a good date-based layout. Peter Syakcek (Oxford) described algorithms for Brain Computer Interfacing (based on external EEG measurements) that can adapt to individual users. A sophisticated variational Kalman filter was used that could be carried out in real time.

Jaz Kandola (Royal Holloway) told us how kernels can be constructed to analyse text. Antonio Criminisi (Microsoft Research) told us just how useful linear algebra can be in the study of visual art. He has worked alongside art historians who are interested in the techniques of Renaissance painters, particularly in the accuracy of their representation of the Euclidean world. He showed us virtual reality reconstructions of buildings from Masaccio and Piero della Francesca (the famous tiled floor in the background of the Flagellation of Christ) and topped it all with an analysis of the geometry of the convex mirror in the background of the Arnolfini Wedding by Jan van Eyck. Next time you visit the National Gallery, it will repay close study!

Dave Mackay (Cambridge) completed the programme for the day by demonstrating how it is possible to input text at near touch-typing speeds without a keyboard. A simple language model (based on text compression algorithms) is used to predict the next few characters; the user navigates amongst the alternatives with simple gestures, making the Dasher interface highly suitable for the disabled. Bravely, he

demonstrated inputting text using many different devices (and even in Japanese; one advantage of the approach is that it can learn from an individual's use of words and be quickly retrained for any language with a reasonable-sized alphabet), such as a mouse, touchpad, head-mounted light source, and even a simple belt to use breathing for control.

This was one of the most impressive demonstrations of technology seen at NCAF, and it is not only freely available, it is also in use through various charities.

The second day started with a tutorial from Chris Bishop (Microsoft Research) who, with his customary clarity, took us from the very basics of mixture modelling to sophisticated variational Bayesian methods for selecting the number of kernels. John Winn (Microsoft Research) then demonstrated the VIBES toolbox which allows the user to specify a probabilistic model graphically and then automatically carries out variational inference. This is a huge advance over the manual approach of deriving the variational equations and then implementing them in software, both highly error-prone tasks.

Peter Swann (Rolls-Royce) told us how they are using agent systems to improve job-shop scheduling. The particular advantage they bring is to improve the flexibility and make rescheduling, to deal with uncertainty and changes in the environment, much easier. The final talk, from Jon Rogers and Rory Hamilton (Royal College of Art) showed us how an understanding of human visual perception can be used to augment the artist's palette, particularly with perception of movement.

Thanks are due to lead Rezek (Oxford), who organised the scientific programme, and Microsoft Research, who very generously sponsored the meeting.

Dr. Ian Nabney
Aston University

Study – Delivering AI-based systems Why do AI based systems fail?

I am often asked this question by people sceptical about the value of AI. My response to them is usually defensive, as I patiently explain (as many of us in the AI community have done) that in fact AI-based systems are successful and out there and working. But in quiet reflective moments, I realise that the sceptics do seem to have a point. We go along to conferences and see innumerable papers on 'applications' but precious few products that actually make it to market

What is going on here?

So are AI-based projects really particularly prone to failure? And if so, why? Is it something that is inherent in the nature of AI, or is it something common to the introduction of new technologies? Is it related to the fact that many AI based projects have their roots in the academic world? Is it something to do with the design and definition of systems, failures in project management or decision making, a lack of understanding of the market or the management of suppliers? Or is it something to do with the culture of AI people?


Postscript to a furniture mover's nightmare

At the Autumn NCAF meeting in Cambridge, David MacKay won a richly deserved bottle of wine after a stunning performance that included the first ever on stage performance of input from a stomach mouse.

He generously returned the bottle and posed his own version of a mini puzzle corner:

"What is the largest object that can be passed round a right angled corner in a corridor of unit width?"

This problem has a fascinating history; it originates from Dr J.M. Hammersley (Trinity College Oxford). He delivered a splendid paper at the AGM of the Institute of Mathematics and its Application in 1967 titled 'On the enfeeblement of mathematical skills by "Modern Mathematics" and by similar soft intellectual trash in schools and universities'. I strongly recommend it to anyone who has not yet read it. It was later published in the journal of the Institute and had as Appendix 4 a set of 16 problems.



Problem 8 states:-


A long passage of unit width has a right-angled bend in it. A flat rigid plate of area A (always kept flat on the floor) can be manoeuvred from one end of the passage to the other. Prove that $A < 2\sqrt{2}$. Show that if the plate has a suitable shape (to be determined), we may have $A = \pi/2 + 2/\pi$. Is this the largest possible value for A ?

For several years it was assumed that $A = \pi/2 + 2/\pi$ ($= 2.2074\dots$) was indeed the optimum solution and was the area of a shape known as the Hammersley Sofa. It consists of two quarter-circles on either side of a 1 by $4/\pi$ rectangle from which a semicircle of radius $2/\pi$ has been removed.

As recently as 1992, Gerver produced an even better solution with an area of $2.2195\dots$. The shape is non trivial being bounded by 18 sections and its solution can be found in (J. L. Gerver, On moving a sofa around a corner, *Geometriae Dedicata* 42 (1992) 267-283).

The more astute will have noticed that David MacKay's paraphrase of the question omitted the word rigid – hence Ian Nabney's (Aston University) solution of an infinite 2D beanbag seems to be the optimum – can anyone improve on this?

Peter Herdman
Arjo Wiggins



tems

AI-based projects fail so often?

How can we find out?

A research study has been established to explore these questions. This study intends to gather information on real AI projects (successful and otherwise), supported by discussions with AI developers, researchers and managers. Information gathered will be analysed to explore the underlying issues, and to try and see where improvements could be made in our approach to delivering AI-based products.

It is hoped that the results will identify whether there are real issues that are specific to AI based systems, and/or issues that apply to the development of all IT based systems that are exacerbated by the use of AI techniques. The study aims to identify approaches and techniques that can help developers avoid the worst pitfalls in future developments.

How can you contribute?

The study team are actively seeking input, and would welcome your contribution. If you have views on the practical aspects of developing AI based systems, then visit the study website (www.ai-research.org.uk) and contribute by contacting the organisers or by

completing one of the online enquiry forms. All those participating will be kept informed of progress.

The study will be also supported by a workshop to be held as part of AI-2003, the Annual Conference of the BCS Specialist Group on AI, to be held in Cambridge, UK in December 2003. NCAF members would be particularly welcome. For further details see the AI-2003 website (www.bcs-sgai.org/ai2003).

Where can you find out more?

Full details can be found at www.ai-research.org.uk. The website provides details of the study, and also has forms for feedback on specific projects and general views on the development and delivery of AI-based systems. Or feel free to contact me if you have any questions.

Richard Ellis

Stratum Management Ltd, NCAF and BCS SGAI

www.ai-research.org.uk

richard.ellis@ai-research.org.uk

richard.ellis@stratum-management.co.uk

PUZZLE CORNER

Number 25

One of Lisa's old colleagues from the Big Ball Bearing Company (BBBC) of Leeds contacted her to describe a challenge they had been given by a rival, Bearings For Gourmets (BFG). The BBBC had been sent an opaque box which they were told they must not open. The box contained a number of identical ball bearings, each one weighing a whole number of brockleys, greater than one brockley. They were told that since the box was virtually weightless, if they knew the weight of the whole thing (in brockleys) they would instantly know the number of bearings in the box.

In order to do this, BFG told them they would need a pair of old-fashioned scales and a particular set of identical weights, each a whole number of brockleys, with the whole set weighing in at 2 kilobrockleys. They were assured that this would be sufficient to calculate the total weight and hence the number of bearings.

However, rather than sending this set of weights, BFG said that if the BBBC was simply told the weight (in brockleys) of an individual weight, then it would still be possible work out the number of bearings. Without providing any further information, BFG immediately challenged the BBBC to divine the number of bearings forthwith.

Needless to say the BBBC were somewhat stumped. However, Lisa was able to give them the answer straight away.

What is the number of ball bearings in the box?

The answer will be given at the next NCAF meeting 21-22 January 2004 Bath University.

Fenella the Rottweiler

(with acknowledgement to Peter Harrison)

COMMITTEE NOTES

Chairman

Graham Hesketh
Rolls-Royce plc

Secretary

Mark Cheeseman
Rolls-Royce plc

Treasurer

Dr Ian Nabney
Aston University

Editor of Neural Computing & Applications Journal

Professor John MacIntyre
University of Sunderland

Managing Editor of Networks

Rick Adderley
AE Solutions

Publicity Officer

Richard Ellis
Stratum Management Ltd

Mandy Bradley
Technical Forecasts Ltd

Dr Richard Everson
University of Exeter

Dr Ilead Rezek
Oxford University

'Bert Bullen
BAE Systems

Dr Elaine Martin
University of Newcastle

Please contact NCAF through Mark Cheeseman Secretary – 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>

Edited and Produced by:
Chris Hawthorne
Forum Communications
Eastgate House
Eastgate Street
Winchester SO23 8DZ
Tel: 01962 877833
Fax: 01962 877988
e-mail:
chrishawthorne@forum-pr.co.uk

NEXT EDITION

Preview of Spring Meeting

Data Rich. Information Poor?

Shaun Marriott is based in a small company in Scunthorpe called Multon Process Technology Ltd (MPT) on secondment from the University of York under the teaching company scheme. MPT specialises in integrated systems solutions – both hardware and software – for the metals industry. The core product is an optimisation system, known as OPTIMAT, which allows optimal blending of scrap and virgin metals to give blends within specified physical and chemical parameters. The OPTIMAT software, together with supporting systems (such as inventory control and management reporting) leads to savings both in materials and energy costs. He writes...

My job is to investigate the possibility of applying neural network solutions to types of issues which have surfaced during work with existing MPT customers. For example, MPT installs a system and everything is working fine and a customer might say "It would be nice if we could predict...as well." If the 'problem' is of the type solvable by a neural network method then a suitable solution could be found, implemented and integrated with the core OPTIMAT system. The technology does not necessarily have to be based upon the 'latest and greatest' algorithm hot from academia. Indeed, a tried and tested architecture tailored to the specific problem will possibly suffice. Industrial applications usually have more fundamental issues to be dealt with before necessitating the development of a brand new algorithm or architecture.

A core theme of NCAF meetings has been the availability, quantity, and quality of data. Variations on the utterance "Yes, we have loads of data collected over a number of years" can cause even a well-experienced 'netter' to tremble (not with excitement, I hasten to add!). Papers with nice, 'clean' problem domains cannot prepare you for dealing with your first encounter with real data.

Often, it has not been collected with any clear purpose in mind; at the very least, a clear purpose disciplines the data collection and increases the chances of a positive outcome.

Where does the whole neural network solution development process begin? Right there, sleeves rolled up, digging about in the data. In addition to data analysis of various levels of sophistication (even 'eyeballing' can be useful!), rapid prototyping of possible neural network solutions can be helpful. In this context, I feel that I must mention the usefulness of the freely available Netlab neural network suite by Ian Nabney based at Aston University. It is written in MATLAB and provides 'potted' versions of a number of useful algorithms for data analysis and model building. It is based upon the seminal 'classic' text of Chris Bishop (Microsoft) and can be found at <http://www.ncrg.aston.ac.uk/netlab/>. The integrated nature of the Netlab suite embodies the philosophy of a principled framework for neural computing methods; the field is not simply about a disparate bunch of algorithms and theories.

To sum up, I'm looking at neural network (or natural computing!) type problems at MPT and currently dealing with the first hurdle, data. Who said we are "data rich and information poor"? Often, the data is poor too. When handed a large archive, ask yourself, how much of it is relevant data?

A final thought: perhaps natural computing methods would penetrate more into the IT mainstream if the emphasis is changed to one commensurate with daily experience. A nice topic of debate for an NCAF meeting...

Shaun Marriott
University of York /
Multon Process Technology Ltd.

DIARY DATES 2004

21–22 January 2004 NCAF – Fundamentals of Natural Computing, Bath University (sponsored by BAE SYSTEMS).

For more information contact Mark Cheeseman e-mail: enquiries@ncaf.org.uk Tel: +44 (0) 1332 246989

5 – 7 April EuroGP2004
7th European Conference on Genetic Programming
Coimbra, Portugal

20 – 22 April ACDM 2004
6th Evolutionary / Adaptive Computing in Design and Manufacture

University of the West of England, UK
9–10 June 2004 NCAF – Intelligent Fault Diagnostics, Sunderland University.

For more information contact Mark Cheeseman e-mail: enquiries@ncaf.org.uk Tel: +44 (0) 1332 246989

12 – 14 July CIMCA 2004
International Conference on Computational Intelligence for Modelling, Control and Automation
Gold Coast, Australia

13 – 17 July SAB04
The Eighth International Conference on the SIMULATION OF ADAPTIVE BEHAVIOR (SAB'04)
Los Angeles, CA, USA

AI-2003

23rd SGA1 International Conference on Innovative Techniques and Applications of Artificial Intelligence. Cambridge, England, 15 – 17 December 2003

AI-2003 offers a wealth of natural computing interest, including a half day tutorial on Evolutionary Computation and Genetic Programming. NCAF members get a £60 discount on the conference fee. See you there!

Full details at <http://www.bcs-sgai.org/ai2003/>

MEMBERS' NEWS AND VIEWS

Deadline for contributions for the next edition – 1 February 2004. Please send to: Managing Editor – Rick Adderley, e-mail: RickAdderley@A-ESolutions.com