

Grand Challenges in Information Driven Health Care

Swansea University 20–21 May 2009

In the digital economy that we live in, it is possible to collect terabytes of data about each patient.

What is the equivalent of putting ‘Man on the Moon’ in the Information Driven Health Care field? Come to Swansea and take part in discovering it.

In the good old days, doctors had more time for one-to-one contact with their patients and were able to remember information regarding a patient’s health records. With just a phone call their expert opinion was available to other doctors as and when necessary – twenty four-seven. Today, a doctor is still only a phone call away, but you may see a different one each time you visit a surgery; communicating relevant information to the new doctor in the ever reducing consultation time has become a challenge. Doctors increasingly rely on each other’s writings in patients’ notes and records, and are under further pressure to meet performance figures. The ability of a doctor to continuously assimilate their patients’ data, convert it into useful information and knowledge, and provide it to other doctors as and when required has become a missing link in the healthcare chain.

In the digital economy that we live in, it is possible to collect terabytes of data about each patient. Today’s technology can convert all our paper notes into electronic format, and we can computerise all our health data, but this isn’t the Grand Challenge! The goal of the forthcoming NCAF meeting sponsored by Swansea University and the Engineering and Physical Sciences Research Council (EPSRC) is to help discover real Grand Challenges in Information Driven Health Care. How can we provide the same or better and personalised health care without relying on the same doctor (and their time) every time? How can we extract useful information and knowledge from the terabytes of data? How do we capture patient data without interfering with the doctor-patient consultation? The list continues...

The consultation process with a doctor, as shown in Figure 1, must remain at the heart of any scoping exercise in discovering the Grand Challenges. A doctor has to assess the patient, retrieve information from historical data, document/record findings, act on the decisions made, review the outcome and once again assess – a repeating cycle which is at the heart of the healthcare process.



Figure 1: A typical consultation process with a doctor.

In March 2008, EPSRC funded seven research groups in the UK to explore the Grand Challenges, and more funding could be on the way. The first day of this NCAF meeting has been devised to showcase results from the research groups along with other invited lectures. In January 2009, the Swansea group (one of the seven groups funded by EPSRC) organised a workshop with leading doctors in the UK in an attempt to scope Grand Challenges as they perceive it. At the meeting in May they will share their findings and use the opportunity to generate discussion.

The NCAF meeting programme includes panel discussions and breakout groups to shape up Grand Challenges that we – as data mining experts – would like EPSRC to fund, as well as opportunities to review what was funded in the past. There will be plenty of time for discussion and networking, extending research groups and collaboration. Representatives from EPSRC have been invited, and the meeting presents a chance to meet them and ask questions. Thanks to EPSRC and Swansea University sponsorship as well as NCAF support, there is no registration fee for this two day event. On the Wednesday evening (20 May) an Indian meal and a surprise recreational event is planned to enhance our team spirit!

On behalf of the Swansea group, it’s my pleasure to welcome NCAF members to participate in this exciting meeting.

Rajesh Ransing
Swansea University

INSIDE

- Intelligent Robotics comes to life at Sunderland
- NCAF online
- Change of Editor
- Puzzle Corner

BACK PAGE

- Open Source Software in Natural Computing
- Diary Dates

Intelligent Robotics com

POPeye is a robotic head that has been designed to mimic the basic abilities of the human head (neck pan and tilt, individual eye movement).

The winter NCAF conference was held on the University of Sunderland's St.Peter's campus in the north-east of England and hosted by the Centre for Hybrid Intelligent Systems (HIS). The focus for the meeting was 'intelligent robotics'.

After a welcome address by Alastair Irons, the Head of the Department of Computing, Engineering and Technology, Stefan Wermter (University of Sunderland) opened the conference with an overview of multi-modal communication and intelligent robotics, including current research at Sunderland. Some robots tend to be purpose built for repetitive pre-programmed tasks, with very limited capabilities. Of more interest to the HIS group, is providing robots with the capabilities to reason, learn complex tasks and the ability to interact with humans through different communication mediums.

Christo Panchev (University of Sunderland) presented his latest work on multi-modal visual, auditory and motor novelty and familiarity detection. The talk centred on using spiking neural networks to combine concepts from different modalities and integrate them into instructions. The implemented architecture was tested on a Khepera robot in simulation and real environments through interpretation and execution of simple instructions. Using a combination of available actions and manipulation of objects, the robot can be instructed to perform such tasks as 'find red ball' or 'turn left'. The architecture is able to associate and learn novel objects and actions. It can partly recognise familiar features in each modality and activate associated familiar features in other modalities.

Matthew Casey (University of Surrey) gave an overview of his current research in which parts of the brain are modelled through the use of Hebbian-based neural networks. Several models have been developed to explore particular behaviours and

test hypotheses through the use of topographical conditioning maps (TCMs) to combine senses that form a multi-sensory representation of an environment. The model has data representing a stimulus location in different modalities using topographical representations of the visual and auditory space. The visual and auditory map outputs can then be associated using Hebbian learning and multi-modal training data. Although this initial model is limited, Matthew noted that improvements would enable many potential applications, such as the ability to prime our bodies for action upon sensing something threatening.

Robotic head

Heidi Christensen (University of Sheffield) highlighted their three-year European project, POP (Perception On Purpose). A key aspect of the project is the audio processing binaural sound source localisation system that is able to run in real-time on the audio-visual head POPeye. POPeye is a robotic head that has been designed to mimic the basic abilities of the human head (neck pan and tilt, individual eye

movement). POPeye is able to focus upon the single most dominant sound. Not only is POPeye able to identify a particular sound source, it is also able to do so when there are multiple sources. Time-frequency regions are identified that belong to a particular source in order to be able to distinguish between multiple sources.

Jindong Liu (University of Sunderland) described research within the MiCRAM project, studying human auditory mechanism. This assists robots operating in noisy environments and the evaluation of biological mechanisms for auditory implants to aid the hearing-impaired. Two different models are combined to find the best estimate of the sound source. Experimental results so far involving noisy and voice sounds, demonstrated that the developed model is able to perform sound localisation that approaches biological performance.


Mark Ebdon (University of Oxford) presented a method for dynamic coalition formation (DCF) in scenarios where a long-term prediction of a target requires many sensors over a long time span. Effective and efficient systems have already been designed; but they are limited to small-scale networks with little or no target dynamics. In complex situations, it is necessary for each sensor to be assigned to a specific target from the set of moving targets, while maintaining high performance and optimal resource usage. Each sensor has a neighbourhood that includes the other sensors that can observe the same targets. Mark stated that forecasts of target movements, with high or low confidence, are incorporated when choosing sensor states. It is suggested that the DCF algorithm could also be appropriate for certain problems in job scheduling and operations research.

'Fenella the Rottweiler' rounded off a day of presentations with the solution to 'Puzzle Corner', after which we took a brisk walk in the cold over to the nearby National Glass Centre. There we were treated to an intriguing and very warm glass blowing demonstration, followed by an enjoyable meal. For those brave enough to venture on into the cold early hours of the morning, a few games of 'Blackjack' at the casino were in order. 'Fenella' quickly proceeded to trade up an initial free £10 bet into £100 before slowly returning the capital to the house and exiting cash neutral.

Simulating sound

Harry Erwin (University of Sunderland) got us back underway on the second day describing the use of robots to apply realistic models and the emulation of natural intelligence, in particular, simulating the detailed processing of the Inferior Colliculus (IC). The IC is a large neural module that interfaces to the thalamus. It seems to play a major role in localising and classifying sound sources. In order to localise and classify sound sources and to control a robot's behaviour in real-time, there is a need for simple models of the IC principal neurons. The models should refer to a general library of channel models in order to be validated. Sodium and potassium channels have been modelled adequately using this procedure; however calcium channels are unable to be modelled sufficiently.

Mark Elshaw (University of Sheffield) highlighted issues relating to the difficulty and current limitability of automatic speech recognition systems, even



The host Stephan Wermter, his student Jindong Lui and their binaural robot.

es to life at Sunderland

in 'noise-free' environments. He proposed a new biologically inspired approach through the use of a recurrent self-organising memory model for emergent speech representation that was developed during the ACORNS (Acquisition of Communication and Recognition Skills) project. The model is described as an approach for associating speech signal representations and semantic features with interactions between working and long-term memory. Speech is processed using a combination of basic self-organising maps (SOMs) each operating at a different level. Different units or regions of the SOM are combined over time in the representation of different word sounds to produce a chain of active units, the semantic features which are found to be critical in the representation of words.

John Murray (University of Hertfordshire) discussed the European project 'FEELIX - Growing' that studies the social incorporation of robots. It will be important that robots can interact in everyday environments. Particularly of interest is the development of emotional expression and non-verbal interaction. Humans and animals convey information by expressing their internal states through emotional expressions. This helps determine the interaction taking place. ERWIN (Emotional Robot With Intelligent Networks) is a multi-modal robotic head that is able to learn and interact through vision and speech. During the experiments, several people interact with ERWIN, with some deciding to play and others not. ERWIN learns and recognises the voices of those who play and builds an internal emotional likeness towards them. After the interactions, ERWIN would respond more favourably towards those participants who had previously interacted than those who had chosen not to interact.

The inspiration for Michael Knowles' (University of Sunderland) research is concerned with Perceptual Symbols Systems. For a robot to interact effectively

with humans, it is necessary for their cognitive processes to be grounded in the same perceptual elements as humans. The PSS model provides a method for capturing symbolic properties of senses and integrating them into higher cognitive processes. Experiments have assessed this approach in a scenario where a robot learns about objects through interaction and reinforcement. Results so far have shown that the use of human perceptual elements combined with interactive reinforcement leads to intuitive learning and interpretable knowledge structures.

Heidi Christensen concluded the conference with her second presentation, concentrating on research regarding the inference of self-movement and tracking moving sources from binaural data as part of the POP project. The human ability to navigate in both familiar and unfamiliar environments is extremely proficient and highly accurate. Sound sources can be identified despite a lack of visual information. This can be achieved while stationary, even while navigating through an unknown environment. The purpose of the system is to localise sound sources in acoustic conditions similar to navigating in real world environments with stationary and active sources. The framework involves five talkers based around a table, taking turns to speak while the perceiver is blindfolded and asked to look at the person who is currently speaking. Initial results have indicated that the ability to detect and track a talker have largely been successful. However the model is capable of tracking independent movements of a perceiver and multiple acoustic sources.

Thanks must go to the local organiser, Stefan Wermter, for making the conference so successful, and to all the speakers for providing us with excellent presentations.

Martin Page
University of Sunderland

Ben had a cracking business idea. His oasis produced an unlimited supply of pineapples and there was a ready market for them in the nearest town. Unfortunately, the only way he could transport his wares from the oasis to the town was on the back of Marmaduke, his trusty but idiosyncratic camel.

It was a good 100 miles to the town, but 'Marmy' was capable of making the trek so long as he was allowed to eat one pineapple for every mile. The problem was that he couldn't carry any more than 100 pineapples at a time. Ben knew that every surplus pineapple he could get to the town he could sell for £10. Unfortunately, 'Marmy' was hooked on pineapple and wouldn't eat anything else, so there was no chance of using any of the money to buy cheaper 'fuel'.

Ben reasoned he could eventually make money out of this business, but he was disappointed by how long it was likely to take, since 'Marmy' refused to do more than 10 miles a day. Lisa, munching down yet another pineapple, asked Ben how much he wanted to make before retiring, and was told "only £4,500". Lisa then assured Ben it would take no longer than 3 years with her foolproof strategy.

What was Lisa's plan and what is the shortest time needed for Ben to achieve his retirement target?

The answer will be given at the next NCAF meeting (20-21 May 2009, Swansea University).

Fenella the Rottweiler

NCAF Online

We can now continue our meetings between meetings thanks to an online discussion group (a.k.a. mailing list) that has been set up for NCAF hosted at: <http://uk.groups.yahoo.com/group/ncaforum/>

You can subscribe online or send an email to ncaforum-subscribe@yahoogroups.co.uk (the subject and body of the email do not matter) after replying to the confirmation email you can then post messages to the group using ncaforum@yahoogroups.co.uk and take part in the online discussion related to NCAF and natural computation, for instance: post questions, answers, discuss 'Puzzle Corner', inform members of upcoming seminars, conferences, workshops and courses, present novel natural computing applications, review books, papers, and have pre- and post-meeting discussion and anything else deemed relevant.

The Yahoo Group also offers facilities for uploading files and photographs and has a calendar which you may like to make use of.

Take advantage of it, it is there for you.

Change of Editor

After many years of sterling service, Nick Granville has reluctantly stepped down as Managing Editor of *Networks*. Nick has been a long time supporter of NCAF, and since taking over the reins as Editor in 2004 he has hardly missed a meeting. A recent change of job, however, has meant that he will find it very difficult to attend meetings in future, and, while he still wants to stay in touch, he felt that it would not be possible to do justice to the newsletter in those circumstances. On behalf of the NCAF community, I would like to extend our warmest thanks to Nick and wish him all the best in his new job.

As Nick steps down, James Hensman has stepped up to the plate as the new Managing Editor. James, a protégé of Keith Worden at the University of Sheffield, has already distinguished himself as the organiser of the September 2008 meeting. With youth on his side (sorry Nick), we look forward to many more years of continued excellence in our principal publication.

Graham Hesketh
Chairman

COMMITTEE NOTES

Chairman/Treasurer

Graham Hesketh
Rolls-Royce plc

Secretary

Jo Thomas
Rolls-Royce plc

Editor of Neural Computing & Applications Journal

Professor John MacIntyre
University of Sunderland

Managing Editor of Networks

James Hensman
University of Sheffield

Dr David Clifton
University of Oxford

Dr Mark Ebdon
University of Oxford

Professor Ian Nabney
University of Aston

Dr Andy Pryke
The Data Mine Ltd

Dr Rajesh Ransing
Swansea University

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

Please send contributions to the Managing Editor – James Hensman,
e-mail: James.Hensman@sheffield.ac.uk

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 Swansea meeting
Preview of September meeting

Open Source Software in Natural Computing

Readers of *Networks* will be familiar with NETLAB, which is a set of Matlab® files which implement various neural networks and other learning algorithms. For years, NETLAB has been a pivotal tool in the applied Neural Networks community. Based on the influential work of Chris Bishop (Bishop 1995), and with its own accompanying book (Nabney 2001), NETLAB provides a basis for researchers, teachers and practitioners to start using its powerful algorithms. One might argue that NETLAB owes at least some of its success to its openness: users are free to read and modify the code, subject to a small set of restrictions.

Many users have taken advantage of NETLAB's open code, using it as a starting point for their own algorithms and as a means of understanding the algorithms it implements. Also, users have submitted corrections and extensions to the original code which are incorporated into later versions. These two-way benefits form the basis of the open source ideals, which simply state that users and programmers can read, modify and redistribute the source code of a piece of software.

The sharing of source code is advantageous to any research field which deals with the invention, refinement and implementation of algorithms; the field of Machine Learning is no exception. Having open source implementations of published algorithms would improve the reproducibility of results and allow one to make much fairer comparisons of algorithms. Having more open source software implementations of algorithms would allow people to build on each other's work, as so many have upon NETLAB, avoiding duplication of effort and allowing for easier collaboration.

In 2007, 16 authors from the worldwide Machine Learning community published a paper in the *Journal of Machine Learning Research (JMLR)* entitled 'The Need for Open Source software in Machine Learning'. In it, they argued that the use of open source software would be of great benefit to the Machine Learning community, and proposed that the JMLR had a new track specifically for the publication of Machine Learning Open Source Software (MLOSS). This, they argued, would bring many benefits to both

the software author and the community as a whole: authors could receive academic credit for their software (since software published in the JMLR can be cited as any normal publication); the standards of the software would be kept high due to rigorous peer review; open licensing of the software would ensure that the algorithms could be continually developed by the community.

Since its creation in response to this paper, the MLOSS track of JMLR has six publications, and many more software projects are housed at the MLOSS.org website. These include algorithms for regression, classification, clustering, feature extraction and more, written in a variety of languages and for a variety of operating systems. The success of the open source ideal is evident in software projects such as the Linux operating system and the Firefox web browser, but what impact can open source have on the NCAF community?

The ultimate value of an algorithm should surely be judged on its ability to solve real world problems. As the JMLR MLOSS track and MLOSS.org increase the access to quality implementations of machine learning algorithms, there is an increasing opportunity for NCAF members to assess and use the algorithms in our various applied fields, critiquing and improving them, and so playing a part in the open source movement.

Open source software is only a small part of open science movement. As the methods for accrediting scientific contributions catch up with the ease of distribution and collaboration that the Internet provides, we are beginning to see a trend toward open standards, formats, and publications, as well as software. We are starting to see the publication of methods (JMLR, jove.com), open pre-print repositories (arxiv.org), public data sets (amazon.com/publicdatasets) and even collaborative theorem proving (gowers.wordpress.com). Of course, open software and open publication are nothing without a public forum for discussing and disseminating the subject, which is a role that NCAF plays so very well.

James Hensman, University of Sheffield
Michael Dewar, University of Edinburgh

DIARY DATES 2009

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 – BIOCAMP'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/>

7–11 September – ECML PKDD 2009, The European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases, Bled, Slovenia.
<http://www.ecmlpkdd2009.net/>

14–17 September – ICANN'09, 19th International Conference on Artificial Neural Networks, Limassol, Cyprus. <http://www.kios.org.cy/ICANN09/>

September – NCAF meeting – Venue and theme TBA. For information email enquiries@ncaf.org.uk or telephone +44 (0)1332 246989

11–15 October – IROS 2009, IEEE/RSJ International Conference on Intelligent Robots and Systems, St. Louis, Missouri, USA.
<http://www.iros09.mtu.edu/>