Networks 22 – March 1999

Condition monitoring at Birmingham 22-23 April 1999

The April NCAF meeting should be of considerable interest for a number of reasons. Not least because it is a joint venture with the Institute of Electrical Engineers (IEE). The IEE regularly arranges colloquia on a variety of topics; this colloquium is sponsored by professional group A9, Neural Computing, in the Informatics Division, and co-sponsored by B4, Intelligent Control Systems in the Control Division. The event can be attended by either NCAF members, or IEE members, at a members' rate. It will be held at IEE, Austin Court, Birmingham. As a result of combining with the IEE, it has been possible to arrange a more comprehensive programme of speakers than usual.

Like the January meeting, the topic is condition monitoring. During the two days, the focus will be on the use of neural computing and other computational intelligence techniques in the domains of condition monitoring and fault diagnosis of machines and external structures, and of health monitoring in medicine. These areas are usually treated quite separately, but they share a number of common issues and solutions and should benefit from a cross-fertilisation of ideas. The invited speakers (from USA, UK and Europe) will provide comprehensive reviews of recent research and techniques employed in the domains of condition monitoring, health monitoring and fault diagnosis; with the underlying aim of facilitating an exchange of ideas and solutions. Some of the common themes and problems which might emerge include those associated with; sensor fusion, avoiding certain kinds of errors, e.g. false positives; problems with unequal numbers of faulty and non-faulty examples, pre-processing methods, combining nets to improve performance, and dimensionality reduction.

Day One

Since the programme consists only of invited speakers, all the talks are of an hour's duration. Following a welcome and introduction by the familiar Peter Cowley (Rolls-Royce plc), the first scheduled speaker is Professor Czeslaw Cempel (Poznan University of Technology, Poland), who will provide an overview of the 'classical' approach to vibro-acoustic condition monitoring. In the second morning session, Dr Chuck Farrar will discuss the Los Alamos Health Monitoring Survey, and provide an outline of work on the monitoring of large structures, namely bridges. The Los Alamos survey was a very influential review in the area of fault detection using vibration data. Chuck Farrar is head of experimental mechanics at the Los Alamos National Laboratory in USA.

In the afternoon Professor James Penman (University of Aberdeen), will talk about 'Condition Monitoring of Electromechanical Plant and Civil Engineering Structures', Professor Lionel Tarassenko (University of Oxford) will discuss 'Novelty detection in Jet Engines', and Professor Ron Patton (University of Hull), will provide an overview of 'Artificial Intelligence Approaches to Fault Diagnosis for Dynamic Systems'. James Penman is head of the Condition Monitoring Research Group at Aberdeen. Lionel Tarassenko has researched novelty detection methods in domains ranging from that of breast cancer to jet engines.

And Ron Patton has been involved in a number of fault diagnosis projects.

Day Two

The second day contains two talks from industry (British Aerospace, and Siemens), and two from academia. A talk by Dr Paul Wells, (British Aerospace) on 'Acoustic Emission Location Using Four Sensors', will be followed by a talk on 'Fault Diagnosis for Closed-Loop Drug Infusion' by Professor Derek Linkens and Dr M.F. Abbod (University of Sheffield).

The medical theme of this talk is continued in the next talk by Dr Volker Tresp (Siemens) on 'Technical and Medical Consulting Systems using Bayes Nets'. The final talk is entitled, 'Why I am Not a Non-Bayesian', and will be delivered by Professor Mahesan Niranjan, a recent appointee to the University of Sheffield. The colloquium will be concluded with a panel discussion by the speakers; participants are encouraged to put questions to the panel, particularly questions which are relevant to a broad spectrum of applications.

Poster sessions will take place during scheduled breaks between the talks.

The event should be of interest to a wide ranging of audience, and as described above consists of a mixture of speakers from academia and industry, and a mixture of new and familiar faces. The audience too is expected to be a mix of familiar and new faces, in its combination of NCAF and IEE members, and of interested non-members. A number of enquiries, or expressions of interest from non-members have already been received, so don't leave it too late to register!

Further information about the event can be found at
http://www.iee.org.uk/Calendar and
http://www.dcs.shef.ac.uk/~amanda/program3.html.

Amanda Sharkey
University of Sheffield

Peter Cowley reviews 'Bayesian Knowledge Discoverer'

Version 1.0 Beta, Authors: Marco Ramoni and Paola Sebastiani - Open University

e-mail: Bkd@open.ac.uk **web** http://kmi.open.ac.uk/projects/bkd

Note - Peter Cowley wrote this review last September so it is possible that changes may have taken place to address problems identified in the article. He reviewed the Windows95/NT version of this software. Apple and Sun versions are also available.

Bayesian Knowledge Discoverer (BKD) is a computer program designed to extract Bayesian Belief Networks (BBNs) from (possibly incomplete) databases. The aim of BKD is to provide a Knowledge Discovery tool able to extract reusable knowledge from databases, using sound and accountable statistical methods, without expecting any particular methodological background from the user.

The software allows the user to explore some of the properties of BBNs. It can be downloaded from the Open University Website referenced above. Modem

users need to be aware that the software is downloaded as an 8 Mb self extracting zip file which expands to over 20 Mb.

The software has a standard Windows user interface. The main window has three tabbed pages. These give a graphical view of the Bayesian network, a command line interface and a view of the database which is being analysed. I only used the graphical interface to the software. I approached the software as a complete tyro. I have modified some of my more naïve comments in the light of some corrections from the authors.

Windows help is provided with the software. The help files contain links to avi video clip files. These worked fine if the help file was opened as a stand alone application, but crashed on my machine (a 200 MHz Compaq Armada, 32Mb of RAM) if I attempted to open them from within the BKD software. [The authors suggest that 48 Mb of RAM is a minimum for full effectiveness.]

Trial and error

The remaining piece of information which is needed to run the software is the format of the database containing the data to be analysed. This information is not in the help files nor on the authors' web site. Trial and error led to the discovery that the database should be in plain ASCII with the field names in the first line, one record per line and each field space separated.

I tested the software with a 600 record database with 4 floating point numeric fields, 2 observations and 2 items of derived results. To check performance with grown up databases I copied the 600 records 50 times. The system did not crash with this size of database, but the results generated were not the same as those obtained with the smaller database which is worrying [if like me you don't understand the methodology, but I'm told fine if you do]. The processing times were acceptable - the larger database took around 2 minutes to generate a network. The user interface is quite friendly, but there are quite a lot of small bugs - non-operating keystroke commands, incorrect re-painting of child windows etc. - perhaps wrapping the underlying LISP in a modern GUI builder would help the authors with future versions of the software.

The software allows the user to define a BBN in terms of nodes, the discretisation of the node variables, their prior probabilities and depend-encies. Alternatively, the BBN may be generated from a database. In the case of our continuous data, the generation process discretised the data. There did not appear to be any way to control the discretisation process before performing the network generation - if a node was manually edited, then the generation process did not adjust the prior probabilities to match the new ranges. Indeed the editor was quite fussy - if an inappropriate entry was made for one discrete item range, none of the range values was subsequently displayed correctly.

Provided you are happy with the automatic discretisation process, then the automatic network generation seems to be effective. The process gave the correct relationships with our test data using any of three search algorithms (Greedy, Arc Inversion or Exhaustive).

Having built a BBN it can be queried by propagating individual items of evidence through the network and examining the posterior probabilities of the nodes. The prior probabilities of the inputs and the log likelihood of the model is calculated for the automatically generated network. Unfortunately this facility is not

available for manually entered ranges. [The authors thought that this should have worked, but I found that each discrete range was given an equal prior, which was patently the wrong answer, even after initialising the network and selecting network.

A feature that I expected and missed was the ability to analyse individual records having defined a model from the data. I would have expected that a significant use of a BBN would be to estimate the likelihood of new data given the model generated from a database of past cases. For example, what is the likelihood that this snapshot of a process is consistent with the model or that this customer will fail to repay a loan or even which records in a database are least likely given the rest of the database. [The authors with some justification think I am missing the point here. If so I stand corrected.]

Discrete interval ranges

The software is capable of estimating conditional probabilities for a manually entered model. This seemed fine when I allowed the software to discretise the data, but as mentioned before, it did not like my attempts to define the discrete interval ranges. This seems to substantially reduce the ability to combine prior knowledge with empirical data - surely this is one of the greatest benefits of BBNs? To be fair, the software never claims to do anything other than extract BBNs from databases and this is what it does. [Again, apologies to the authors if the problem is one of my own making.]

In summary this is an interesting and educational piece of free software. For the non-Bayesian practitioner such as myself it gives some insight into the capabilities of the methodology. It is of limited practical value for the type of data analysis I do where I want to automate the interrogation of the models which are derived from data and I would think twice about the cost/benefit of downloading the software at home. However if your application is one of model selection then this could be the application you need.

The authors supplied the following quotation from an agricultural research laboratory which has used BKD: ***"I ran a network more of than a hundred nodes, with about 25,000 records. BKD worked overnight, and responded with a neat little network which was in accordance with the expert's expectations. Nice work by BKD."***

Peter Cowley
Rolls-Royce

Review of Sheffield Meeting

The first meeting of 1999 saw a return to Sheffield. After a certain amount of confusion caused by Granada Television's takeover of the Mappin Building, the meeting finally took place in the Henry Stephenson Building. The theme of the event was condition monitoring and fault diagnosis, fields in which computational intelligence techniques are playing an increasingly important role with each passing year.

The first morning began with a talk by Peter Lloyd of DERA Farnborough on the role of computational intelligence in structural health monitoring. The discussion concentrated on aerospace applications and gave the academics in the audience a valuable insight into precisely what industry expects and desires from neural network research. The presentation sparked the first of several lively

discussions. The second talk moved from the abstract to the concrete; Sophoclis Patsias of the University of Sheffield presented preliminary results from a monitoring study of a ball-bearing system. This was the first of several case studies.

After coffee, the meeting reconvened for the first of the keynote presentations. This was given by George Irwin of Queen's University, Belfast. The theme of the talk was Multivariate Statistical Process Control and this was illustrated using the Tennessee Eastman Process Benchmark. The benchmark is not well-known outside the process control community and it proved interesting to many of the audience to see how many complex issues in fault detection and isolation are actually addressed. The presentation also served as a tutorial introduction to MSPC, focusing on the techniques of Principal Component Analysis and Partial Least Squares.

System identification

After lunch, the first presentation of the afternoon was made by Andreas Kyprianou of the University of Sheffield. The subject was system identification using the Differential Evolution algorithm. The algorithm - essentially a Genetic Algorithm searching over real numbers - shows considerable promise in identifying systems which present problems for conventional techniques. The method was illustrated on a type of non-linear system of practical interest. The second talk was given by Andrew Starkey of the University of Aberdeen. This, another case study, presented a condition monitoring system for rockbolt ground anchorage systems which are used extensively to support Civil structures such as tunnels and bridges. The heart of the system was a Multi-Layer Perceptron which translated the results of an impulse test into a structural diagnosis.

The final event of the first day was a panel discussion on the role of computational intelligence in condition monitoring. This featured NCAF chairman Peter Cowley standing in, in the rather unusual role of devil's advocate. Andrew Starr of the University of Manchester joined him in friendly opposition to Chris Kirkham and John MacIntyre of Brunel and Sunderland Universities respectively. After the panel members introduced themselves, the audience took up the discussion, which broadened to consider questions of when condition monitoring is advisable in the first place to the problem of transferring the technology from academia to industry. The panel discussion served to fuel conversation throughout the evening.

The social event for the meeting was a 'traditional' conference dinner held in Halifax Hall of Residence where many of the members were staying. After dinner, the majority retired to the bar, while an intrepid few set out in search of the Fat Cat.

The second day began with a presentation by Trevor Holroyd of Holroyd Instruments overviewing some applications of acoustic emission. The speaker communicated some of his wealth of experience in the field and showed how the methods of computational intelligence had been incorporated via his collaboration with the Centre for Adaptive Systems in Sunderland. The second talk was given by Colin Wignall of AEA Technology on a business user's perspective on neural networks. This provided another valuable insight into the expectations and desires of industry.

After coffee, the second keynote speaker, Ulrich Rückert of the University of Paderborn, presented 'The Silicon Way to Artificial Neural Networks'. This departed from algorithms and case studies and discussed in some detail how to realise neural network architectures in microelectronics. A range of different paradigms were discussed from associative memories to radial basis function networks.

Black Arts

The NCAF AGM followed immediately after lunch. After a review of the year by the chairman, treasurer and secretary, three new members of the committee were elected. The first talk of the afternoon was by Robert Milne of Intelligent Applications Ltd on their TIGER system for gas turbine monitoring. Apart from the details of the extremely effective system, the talk again raised issues about transferring technology from the research environment to the marketplace. John Brandon of Cardiff University continued the expert system theme with his presentation 'Black Arts, Black Boxes and Forgotten Arts', but ranged over broader issues including the incorporation of computational intelligence techniques into the overall techno-managerial system. Puzzle Corner followed this session with the solution acted out in the now traditional pantomime. The usual entertaining battle of wits ensued between Graham Hesketh with his carefully crafted script and the actors who had no intention of following it.

The final session of the day began with Nick Lieven of Bristol University who presented his work on model updating using neural networks. The updating task is a notoriously ill-conditioned inverse problem and neural networks are showing some promise in overcoming the sensitivity of the problem to noise. The day and meeting closed with a presentation by Lindsey Jack of Strathclyde University on the monitoring of rotating machinery. This focused on the choice of features for construction of an effective diagnostic: spectra, wavelets and higher-order statistics were among those discussed.

The theme of fault diagnosis will be carried forward to the next meeting - Birmingham in April.

Keith Worden
Sheffield University

PUZZLE CORNER Number 8

Lisa decided to get her hands dirty with some real condition monitoring problems, and was hired to inspect a Neocognitron. This huge machine had four See-Mac Thread-U-Like bearings arranged in a line. Each bearing could be in one of ten different states of wear ranging from 0 (brand new) to 9 (call this a bearing!). To monitor the machine health, a single accelerometer was placed on the housing. The vibration signature it recorded was unique for each possible permutation of wear level and position. However, it was not possible to diagnose the wear states of the different bearings directly from the signal because of the effect of day-to-day environmental changes.

The only way to deduce the true machine state was to hypothesise a particular set of wear levels and synthesise the expected signal with a feedforward computer model, incorporating the current environmental conditions. A trained engineer could compare the synthetic trace with the actual accelerometer reading and could say with certainty which hypothesised wear levels were

actually present, and of those correctly identified, which were also correctly located. It will come as no surprise that the computer model was complicated, slow and expensive, so the engineers were very keen to limit the usage.

Today, Lisa was being trained by an engineer. He sucked his finger, put it in the air, and fed the following wear level set into the computer model: [9 7 6 0]. Comparison of the traces indicated that one of the wear levels was present, but it was not correctly located (e.g. if one of the bearings was brand new, it was definitely not bearing 4). Continuing, he fed in [4 2 9 1], and the comparison yielded a similar conclusion. Undeterred, he fed in the following sets, each one returning the same conclusion - one wear level present but not correctly located: [0 5 4 3], [1 6 5 8], [8 4 7 9]. Turning to Lisa he said, 'It's never taken me more than 6 simulations before to diagnose the state, but I can't figure this one out.' Lisa said, 'Try [3 0 1 3].'

'That can't possibly be right', said the engineer. 'I know', she replied, ' but it guarantees to find the answer without further simulation.' He tried it and (guess what) got the same conclusion as for all the other tries. Lisa now confidently identified the correct wear levels for all four bearings.

What were the actual levels? Was Lisa's solution optimal? Could the engineer have done better?
The answers will be given at the next NCAF meeting (22-23 April 1999, IEE Birmingham)

Fenella the Rottweiler

Data Mining Update

Data mining is now building up a head of steam. Publicity has reached national radio and the national press. In addition to well-established specialist companies, large database suppliers are now competing to supply a desktop data exploitation capability, building on the widespread application of data warehouses. With rapid progress in OLAP (On-line analytic processing, a.k.a. multi-dimensional tables), large numbers of users are able to 'drill-down' into data, and organise their data so as to show clearly hierarchical structures and clusters. It is often at this point that an interest arises in taking exploratory analysis further. Desktop data mining tools such as Business Miner from Business Objects are making the results of years of machine learning R&D available at a reasonable price, with user-friendly interfaces and wizards. Whilst a little knowledge can be a dangerous thing, a user community aware of the possibilities of data mining makes the task of introducing concepts such as neural networks and unsupervised learning somewhat less hard work than it used to be.

Suppliers of relational databases are wary of plunging into development of advanced data mining techniques themselves, and typically are looking to partner with small, specialist companies, or buy them up. Development of application programming interfaces (APIs) to support data interchange and OLAP provides an opportunity for fusing areas of hierarchical, multidimensional tabular analysis and the neural network/rule induction/fuzzy logic approach. Maybe the next generation of tools will take advantage of this? There is a challenge here to identify learning algorithms which are able to work with such structures.

Another issue arising is the design of databases to facilitate knowledge discovery. In many cases data warehouses designed by pure database theorists have a highly normalised structure which makes extraction and cleaning of data for a data mining exercise difficult.

Parallel to these software developments, the pace of improvement in hardware performance continues to be fast and furious, resulting in the advent of companies who specialise in mining very large databases without using sampling or abstraction. What is the impact of this for the data mining community? Rules arrived at by induction or design, and the input-output models developed with neural networks are based on a concept of statistical generalisation. Progress in hardware and OLAP performance is enabling a more individual treatment of data records. Whilst there will still be a strong need to make generalisations about parts of a population, the new approaches may lead to more focused and targeted manipulation of, for example, point-of-sale or individual financial data.

Simon Cumming
British Airways

God's a Bayesian.....but is anybody else?

The Bayesian statistician has uncertainty about his prior. This is demonstrated in the Bayesian approach in practice, in which a prior is specified, then analysis is done on that basis, then the analysis gives stupid results and then the prior is modified.

It has been argued that Bayesian statistics offer the only coherent method for handling uncertainty, in fact there is a proof to this effect. Call 'strong Bayes' the proposition that uncertainty must always be handled using Bayesian statistics, in particular representing our ignorance by means of a prior.

According to strong Bayes, a problem involving uncertainty must be addressed by specifying a prior. Then, the uncertainty concerning the prior should be dealt with by specifying a meta-prior. But this meta-prior must also be uncertain, so a meta-meta-prior should be specified. And so on to infinity.

Bayes is a tool

My argument does not prove that Bayesian analysis is a bad idea. I imagine that a meta-prior across priors is sometimes useful, just as the more conventional prior is sometimes useful. The regress is cut off by the same common sense that avoids Bayesian analysis altogether for some applications. Bayes is a tool, not a rule.

It is always open to the Bayesian, to express doubts about his prior by making a more comprehensive, but single, prior. (I would call this 'advanced main-stream Bayes'.) But the fact remains that the prior, however modified, is uncertain, and this uncertainty is not managed using Bayesian statistics.

We seem to have a choice between incoherence and infinite run-time. Psychological studies show that human beings err on the side of the former, and this seems wise.

The Bayesians might counter that the infinite regress of priors, can converge; so it can be approximated with a finite number of layers. But then the approximation problem is also open to uncertainty. Data analysis that terminates is not strong Bayesian.

Andrew Swann
Rolls-Royce

WELCOME TO NEW MEMBERS

The Committee is pleased to welcome the following new members to NCAF:

Mr Keith Copsey, DERA, Malvern
Dr Leslie Malkin, Centre for Process Analytics and Control Technology
Dr Paul Corcoran, University of Derby
Mr Lindsay Jack, University of Liverpool
Mr Tshilidzi Marwala, University of Cambridge
Mr Colin Turnbull, Axeon Limited

Diary Dates

22-23 April NCAF/IEE Joint Meeting, Birmingham.
Contact: Events Office, IEE, Savoy Place, London WC2R OBL. Tel: +44(0)171 240 1871 ext 2206 or 2205, or e-mail: events@iee.org.uk.

21-23 April ESANN'99: 7th Conference. European Symposium on Artificial Neural Networks in Bruges, Belgium. http://www.dice.ucl.ac.be/esann/

2-4 June IWANN'99. International Work-Conference on Artificial and Natural Neural Networks (IWANN) in Alicante, Spain. http://iwann99.umh.es/

10-16 July IJCNN'99. 10th International Joint Conference on Neural Networks in Washington, DC, USA.
http://www.cas.american.edu/~medsker/ijcnn99/ijcnn99.html

13-14 July NCAF Meeting. Bristol University

Contact: Sally Francis. Tel: +44 1784 477271 or +44 1784 431341 ext 270, fax: +44 1784 472879, e-mail: ncafsec@brunel.ac.uk

Members' news and views
Deadline for the next edition is 7 May 1999.

Next Edition
Review of the NCAF/IEE Joint Meeting.
Preview of the Bristol Meeting.