Networks 19 – May 1998

Rolls-Royce treatment for NCAF Summer Meeting

NCAF is visiting Rolls-Royce in Derby for its Summer meeting. Rolls-Royce is a company with a long history and a well recognised name, operating in highly competitive markets. Through their insistence on excellence, commitment to strategic research, and close involvement with universities and organisations like NCAF, Rolls-Royce is keeping itself at the forefront of applied innovation, something which is essential in today's highly technological industries. This meeting, however, will not be dominated by aerospace interests, but will contain the usual mix of talks from industry and academia, with something for everyone.

Our invited speaker is Professor Gammerman, Head of the Computer Science Department of Royal Holloway College. The subject of his talk will be Support Vector Machines (SVM), a newly developed and exciting blend of Computational Learning Theory and Neural Networks. This method has been specially designed to prevent over-fitting in high-dimensional feature spaces, and results published by Vapnik et al show excellent performance.

While NCAF has neural computing in its name, our survey of the membership last year made it quite clear that there is wide interest in other computational intelligence techniques. Consequently we are delighted to welcome Rachel Pearce to give a tutorial on the use of Genetic Algorithms. The role of Genetic Algorithms as optimisers is widely recognised to be an excellent compromise between highly random search methods and more traditional gradient-based methods. Rachel will introduce the techniques, highlight the types of application most likely to benefit and give examples of the opportunities to take a different view of what we should be aiming to optimise, and how we should present the results to our customers.

The application talks are quite varied, ranging from case studies in finance, through wavelet feature extraction to the intriguing TARDIS with ANN. There is also the enigmatic 'What can neurons compute?'. Brian Kett of Neural Computer Sciences will be describing a recent departure for NCS, who now offer a bespoke applications service as well as their traditional line of products. And, of course, it wouldn't be Rolls-Royce if there weren't talks on aircraft data and condition monitoring somewhere in the meeting.

The social event for this meeting will involve a walk below the streets of Derby. Derby was originally a Roman settlement and over the intervening two millennia has built up layers of caves, dungeons, cellars and other places where unpleasant things were done in the name of law and order. There will be a tour of the most haunted parts of these nether regions of Derby, followed by an informal meal in the catacombs.

The final event of the meeting will be an opportunity to visit Rolls-Royce's interactive display area. This shows some of the technology that goes into Rolls-Royce's latest engines and some of the EDS technology that helps to make the technological information available.

As places will be limited, early booking is recommended. We look forward to seeing you in Derby.

Graham Hesketh
Rolls-Royce plc

Review of Sunderland Meeting April 1998

NCAF's second visit to Sunderland proved just as successful as the first in April 1996. The meeting was held in the impressive new School of Computing and Information Systems at St. Peter's Campus and hosted by John MacIntyre of the Centre for Adaptive Systems. As before, the social event took place at the Beamish Open Air Museum and was based in a recreated North East town complete with bandstand, shops and, most importantly, 'The Sun Inn', an extremely pleasant public house. The reviewer is shamefully unable to describe the majority of the attractions but can recommend the Sun Inn without reservation.

The programme for the meeting was based on the application guidelines given in Professor Lionel Tarassenko's book, and follows the Malvern meeting in adopting this theme. Proceedings began on the first day with the keynote presentation of Lutz Prechelt of the Universitaet Karlsruhe on 'Neural Computing Applications - How to Make Progress'.

Although the talk made some contact with Professor Tarassenko's guidelines, the main theme was concerned with communication. The importance of high quality publications was stressed for academics and applications developers alike. The talk also presented a limited but illuminating survey of the recent literature, which classified articles as 'good', 'acceptable' or 'not acceptable'. The rather low number of good articles found serves as an object lesson for us all.

The next session introduced a strong sub-theme of the meeting - error bars and confidence intervals. The first talk was by Ian Nabney of the University of Aston entitled 'Practical Assessment of Neural Network Applications'. Based around a series of case studies, the presentation stressed the importance of having a measure of reliability for model predictions.

The final talk of the morning complemented this perfectly; David Lowe, also of Aston, moved on from the question of 'why error bars' to the question of 'how'. The talk outlined eight methods of estimating confidence intervals with detailed description reserved for three: Bayesian error bars, Gaussian processes and Predictive error bars. The methods were illustrated using data from an automotive engine calibration. The question of estimating the reliability of the error bars was raised and provoked an interesting discussion.

The afternoon began with 'Cracking the Code' an interactive workshop based on putting the guidelines into practice. This featured Graham Hesketh in the role of 'God' with Iain Strachan as a sort of 'Holy Spirit', improvising solutions to problems using a laptop. The workshop was based around a card-game 'Eleusis' in which the audience was invited to discover the rules of the game, as ordained by 'God', from observational evidence. This made for a highly entertaining session, which nonetheless reinforced well the rules of 'best

practice'.

The first day concluded with 'From Project to Product, A Neural Based Cardiac Monitor'. Tom Harris and Lee Gamlyn, formerly of Brunel University and now of Cardionetics Ltd, described how they had carried a particular application from undergraduate project to marketplace. Alternating between technical and commercial aspects of the problem, the speakers managed to convey the excitement, and frequent frustration, involved in moving from an academic environment to a new company.

The second day began with a continuation of the error-bar theme. Julian Morris of the University of Newcastle presented joint work with Elaine Martin on 'Comparison of Two Novel Approaches to Confidence Bounds for Neural Network Representations'. The talk described how 'stacking' can be used to assemble results from multiple networks and generate statistics for their predictions.

The important question of input data density was also addressed. The approach was illustrated using data from process control applications. Throughout, Professor Morris stressed the need for stringent model validity testing in applications. The second presentation of the session was by Mihaela Duta of the University of Oxford on 'Neural Network Techniques for the On-Line Monitoring of Vigilance'. This provided a case study in the analysis of noisy multivariate data and stressed the importance of data visualisation.

Following coffee, Dan Bretherton of British Gas technology spoke about 'Neural Networks in Gas Demand Forecasting'. Neural networks are among the most successful of an array of techniques used by British Gas to estimate the short-term demand for gas. The presentation followed the guidelines in stressing the need for careful feature selection and extraction before passing to the neural network. The second talk of the session was by Odin Taylor of the University of Sunderland on 'Data Fusion in Complex Machine Monitoring', and described the ambitious 'Neural-Maine' project which seeks a structured solution to the general plant condition monitoring problem. Neural network novelty detectors would be used at machine level to report deviations from normality to 'concentrators' and thence to a decision-making 'overseer'.

The morning concluded with the solution to the last Puzzle Corner, 'The Good, the Bad and the Ugly'. Graham Hesketh led us through his solution assisted, in the best pantomime traditions, by David Lowe as Lisa and Julia Stegemann and Rachel Emmett (the evil Professor Zenner and Dr. Wolfram respectively).

The final afternoon comprised three presentations. In the first, Peter Mattison of the University of Sunderland spoke on 'Neural Networks for Steam Leak Detection'. The estimated losses for UK industry from this problem amount to some £260 million per annum. Current solutions are costly and suffer from false alarms. The presentation showed how neural networks can be used to eliminate the false alarms and build confidence in the conventional approach.

The second speaker was Andy Wright of British Aerospace's Sowerby Research Centre. In an entertaining departure from neural network practice, he outlined some thoughts on 'Flocks, Herds and Robots'. The main theme was concerned with how emergent properties result from collective behaviour; a simple physical model sufficed to produce quite complex behaviour.

The final talk of the meeting returned us to the applications arena. Jim Austin

and Ping Zhou of the University of York presented 'Neural Networks for Telephone Line Fault Detection'. Despite restrictions imposed by commercial confidentially, the presentation showed clearly, the benefits to industry from a principled approach to neural networks.

The meeting concluded with John MacIntyre presenting all the delegates with specially bottled NCAF souvenir brown ale.

Keith Worden
Sheffield University

Tarassenko's 'A Guide to Neural Computing Applications'

All NCAF members should have recently received a copy of Lionel Tarassenko's book 'A Guide to Neural Computing Applications'. Professor Tarassenko is well qualified to write such a guide, as he has been responsible for many successful neural computing applications, including two products: Sharp's neural network controlled microwave oven (the LogiCook), and the Questar sleep disorder diagnosis system sold by Oxford Instruments. The aim of this book is to educate its readers in the real benefits which neural networks can bring if properly applied. It does this by showing, through worked examples, both a sound way to train a neural network and the pathologies which can arise as a result of not doing so.

The book is aimed at the new user of the technology who has perhaps read about neural computing applications in their domain and has bought some software but who does not know how to use it to best advantage. After a brief technical description of the models used in the book (MLP, RBF and Kohonen map), the text starts with the planning and management of a project. It very clearly differentiates between the requirements of neural computing and 'conventional' system development. There is a particularly useful section on configuration management which lists the additional information that should be recorded to take the place of source code control. A whole chapter is then devoted to the question of identifying applications and writing a business case. While the advice of searching the literature for similar applications is good, more pointers to good sources of application papers would have been useful (although NCAF members are in the fortunate position of having their own journal to give them a start). A particularly important piece of advice is to use fast (non-neural) methods such as linear regression or nearest neighbour to predict the potential performance of a neural network solution. This reduces the risk of trying to develop an application where the data is not adequate to solve the task.

Chapter 6 is concerned with data collection and assessing data quality. Visualisation is strongly emphasised as a means of assessing the data and feature selection, though the methods for doing this are spread across several chapters. I would have liked to have seen more discussion of quantitative methods for feature selection, such as analysis of the correlation matrix, and more detail on PCA and related methods for dimensionality reduction. Variable coding and network training are explained very clearly and there is a useful section on common problems with symptoms and cures. Although no specific software tools are suggested, by analysing the approach used, it can be seen that Prof. Tarassenko uses MLP, RBF, visualisation methods (such as Kohonen), linear regression, nearest neighbour and PCA. This, then, is the checklist that the reader should use as the minimum requirements for a useful suite of

software. The most valuable part of the book is the case studies: sleep classification from EEG, and prediction of diabetes. Both of these are scientific/engineering applications, and it is generally true that the book shows a bias towards this type of application. There is also little discussion of regression or time series, which is a more serious limitation. The case studies are used to illustrate most aspects of the development life cycle (though not the business case or maintenance), and errors are deliberately introduced to show the effect on results of making poor decisions. The book concludes with pointers to more advanced topics, of which the most important are improved methods for visualisation and better training algorithms than gradient descent.

Overall, I believe that this book will help novice neural network users apply the technology in an appropriate way. The case studies and pathologies sections are particularly useful in this regard. I found the presentation of the technical material less satisfactory; it was rather piecemeal, tended to be repetitious, and didn't always give the definition of terms before using them. These issues mean that the book is better suited to readers who already have a basic understanding of the technology. There is a strong bias to engineering, rather than commercial, applications and the text assumes that the reader is familiar with undergraduate level mathematics, including partial derivatives and probability theory. This is justified by the need to understand the reasons behind the pathologies that can arise in practice. The book fills a useful gap and will, I believe, help many people to develop successful applications.

Ian Nabney
Aston University

New Secretary for NCAF

Welcome to Chris Kirkham, our new secretary. After many years, Tom Harris has stood down as the secretary of NCAF in order to concentrate on his new job as Research Director of Cardionetics. The saga of setting up this neural computing based company was, of course, the subject of Tom's most recent presentation to NCAF in Sunderland.

We owe a lot to Tom for his energy and business skills in making NCAF the stable and successful organisation it is today. Tom's will be a hard act to follow, but Chris is just the guy to do it. He comes from the same Brunel department as Tom and shares his enthusiasm for neural computing.

PUZZLE CORNER - Number 5

After the 'Bloodbath in Brooklyn', Lisa's reputation at the University of Hard Knox soared. She was immediately approached by the Contracts department to help with their annual funding negotiations. An external Agency was responsible for placing 100 fully-funded projects per year, and the only universities which competed for them were UHK and their arch-rivals UHP, the University of Hedonistic Pursuits. Each contract went to one or the other and was always for $100,000.

For each project, the Agency sent the first university Form A to complete and the second university got Form B. Each university had to reply immediately indicating if they thought the project was feasible or not (via a check box on the

form). When the Agency compared the replies, if both had indicated it was feasible then the project (and hence the money!) went to the university named on Form B. If there was a difference of opinion then the project went to the nominee on Form A, irrespective of whether they were the one who had said it was feasible! If both indicated it was not feasible then the Agency repeated the procedure but this time the first university got Form B and the second university got Form A (i.e. they inverted recipients). The new replies were then evaluated as before. If this still did not succeed in placing the contract then the whole thing was repeated from the beginning. In reality, the projects were always feasible, so the objective of each university was merely to maximise the number of contracts it was awarded by judicious (rather than honest) use of the check boxes.

'We did OK the first year,' said UHK's Chief Negotiator, 'but we got slaughtered ever since. We realised that our best strategy depends on our assessment of UHP. Unfortunately, our assessment is that they're smarter than us - which doesn't help. They're now so confident that they have agreed to let us start with Form A on all the projects this year. We know this should give us an advantage, but we've failed to exploit it in the past. Can you help us, please?'

Lisa said she had a strategy for him which would work equally well no matter what UHP did, and, for the coming year at least, should net him significantly more than half the projects.

What strategy did Lisa suggest (in terms of how to reply to Forms A and B) and what fraction of the contracts did she expect to win? Full answers will be given at the next NCAF meeting (30 June/1 July 1998, Rolls-Royce, Derby).

The Rottweiler

Visit NIPS Colorado this Winter

NIPS is the Neural Information Processing Systems conference, which happens every year in early December in Denver, Colorado. It is as an interdisciplinary conference, bringing together workers in the biological side of 'neural networks', the mathematical theorists, and those who turn the ideas into real applications. The main conference in Denver has a single stream and attracts about 450 people, mostly from the USA, Canada, the UK and Germany, but with representation from many countries around the world. Attendance is predominantly academic, but small numbers of researchers from large companies are also present. NIPS is accompanied by an excellent tutorial day before the main conference and the (now famous) workshop sessions afterwards at one of Colorado's top-rank ski resorts.

Downtown Denver is a pleasant city and even in December the weather is often dry and warm (though it can be very cold at night). Denver is situated at 5,280ft above sea level, so is literally a mile high (and that's at the bottom of the mountains!), while Breckenridge enjoys 9,500 ft of elevation (at valley level) and some premium ski territory in the midst of the Rockies.

The main themes at the 1997 conference were Support Vector Machines, Gaussian Processes and Probabilistic Graphical Models, which subjects also formed the applications oriented side of the tutorial sessions. In many ways attending NIPS was a useful continuation from the applications Week at the

Newton Institute, as many of the same issues were being addressed, and further ideas being brought to bear on them.

The support vector machine is an idea from Vladimir Vapnik, and represents really the first usable method to come out of computational learning theory. The idea is to map the input vector into a high-dimensional feature space and construct an optimal separating hyperplane in the feature space, so that the system will generalise well even if dimensionality is high.

It is difficult to do justice to all the papers at the main conference, but some which I found noteworthy were Freund (AT&T) on adaptive boosting, Maron and Lozano-Perez (MIT) on multiple instance learning, Atkeson (GeorgiaTech) on learning from demonstration, and Baluja (Justsystem, Pittsburgh) on using expectation to guide processing, taking advantage of temporal coherence or predictability in applications such as autonomous driving, hand tracking and semi-conductor etching.

On the workshop days at Breckenridge, I attended sessions on Graphical Models and Dynamical Data Structures. The former brought together many of the world experts in this rapidly moving domain, and the pace was fast and furious. Graphical models are a general framework incorporating Bayesian belief nets, some kinds of neural nets and hidden Markov models. Some significant progress on algorithms which were first introduced in the 1960's and 1980's has been accompanied by realisations of links with other disciplines, for example coding theory.

The 'dynamical data structures' workshop addressed various aspects of getting neural networks to work with inputs and outputs which are not fixed-length vectors. This area fascinates me, but the theory of how to handle variable or structured inputs is in its infancy. One approach is Folding Architecture Networks, which can be seen as a natural extension of the concept of a recurrent network. Much of the work on this is going on in Germany. One idea being tried is 'back-propagation through structure' with tree-based structures as inputs and outputs. Another approach is Recursive Autoassociative Memory (RAAM). Other themes which were explored in this wide-ranging but rather theoretical workshop were neural learning of spatiotemporal dynamics and learning of dynamical Bayesian networks.

Add to this the informal networking on the slopes and in the hot-tub, it made for an exciting and intensive conference. Some of you will be submitting papers for NIPS 98. I would encourage anyone who can to go. This year's NIPS is from 30 November until 5 December, again at Denver and Breckenridge.

Simon Cumming
British Airways plc

Welcome New Members

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

Mr Kevin Swingler, Neural Innovation Limited
Miss Rachel Emmett, GKN Westland Helicopters
Mr Weera Kompreyarat, University of Northumbria at Newcastle
Mr Jahangir Haque, Mektronika Systems Limited

Professor Shigekazu Ishihara, Onomichi Junior College, Japan
Dr Mark Lauder, DERA, Chertsey
Professor J R Wright, University of Manchester
Professor D A Linkens, University of Sheffield
and all members of the ERA SMART Software Technology Club