synergy
http://www.dcs.napier.ac.uk/coil/
Newsletter of CoIL – The Computational Intelligence and Learning Cluster

About CoIL

CoIL is a partnership between NEuroNet, MLNet, ERUDIT and EvoNet. The four Networks of Excellence have joined forces to promote co-operation and cross-fertilisation between their respective fields of neural networks, machine learning, fuzzy logic and evolutionary computing.

CoIL aims to be more than the sum of its parts. By emphasising the potential synergy between these distinct but complementary paradigms, CoIL will work to extend the scope of computational intelligence and learning, and to foster the emergence of hybrid intelligent systems.

Coil is funded by the European Commission and has a combined membership of several thousand nodes and individuals – providing an unparalleled opportunity for cross-sector, inter-network and trans-European communication and collaboration.

During the year 2000, Coil will undertake a number of actions in pursuit of its goals of technical integration, collaborative research and technology transfer. These include producing a technological roadmap and tutorial materials, as well as organising workshops, a summer school and an annual CoIL competition.

For more information about CoIL activities, visit CoILWeb at:

http://www.dcs.napier.ac.uk/coil/

Lessons about self-learning

The record number of submissions attracted by this year's CoIL competition, and the wide range of approaches used to tackle the competition task, allow some important lessons to be drawn from the results.

The response to this year's CoIL Challenge was exceptionally high – double that of previous CoIL and ERUDIT competitions, and 50 percent higher than the reponse to KDDCup, the US counterpart data mining competition. In all, 147 participants registered, and 43 submitted a solution, with entries arriving from Europe, North and South America, Asia and Australia, as well as from both industry and academia.

'The Challenge was a tremendous group effort carried out by the participants in their spare time,' comments Peter van der Putten, one of the competition organisers and a consultant with the Dutch datamining company, Sentient Machine Research.

Based on a direct marketing problem, the goal of the competition was to predict and explain caravan insurance policy ownership on the basis of product usage and socio-demographic data supplied about the customer. While the results demonstrate that the CoIL community has some extremely effective approaches with which to address the target selection task, they also highlight a number of weaknesses.

'Typical problems are the focus on algorithms instead of methodologies and the lack of tools and efforts to explain the discovered patterns,' comments van der Putten.

The competition was divided into two categories: prediction and description. For the prediction task, participants were asked to identify the set of 800 customers in the test set that contained the most caravan policy owners. The maximum number of policy owners that could be found was 238 (random selection results in 42 policy owners) and participants' prediction scores ranged from 38 to 121 policy holders identified, with the best submissions offering a substantial return on investment had they been applied in practice.

A wide variety of methodological approaches were used, including derived attributes, recoding attributes, feature selection, feature construction and feature reduction, boosting, bootstrapping and cost-sensitive classification. Algorithms employed included standard statistics, neural networks, evolutionary algorithms, genetic programming, fuzzy classifiers, decision and regression trees, support vector machines and inductive logic programming.

'The results of the competition suggest that, at least on these kinds of problems, methodology is far more important than the algorithm that is used,' says van der Putten. 'The spread in the prediction scores is large, and the results of applying neural nets, for instance, vary strongly depending on the chosen methodology.'

He believes that the competition highlights the importance of trying a variety of different approaches. 'Participants who wanted only to use their pet algorithm and did not experiment with other algorithms and methodologies scored badly,' he says.

'Another point that arose during the post-Challenge discussions was that on real world prediction problems like this, one should try a wide variety of approaches, starting with the most simple ones.

'By simple I mean simple for a user to employ, with a robust, stable behaviour which can be easily explained – a kind of external simplicity if you like. Algorithms which are, say, internally simple, such as naive bayes and nearest neighbour, tend to be simple to use as well. However, sophisticated algorithms such as support vector machines can be automated so that they are more simple to use.'

(Continued on page 2)

Lessons about self-learning

(Continued from page 1)

He believes the competition results highlight the need for more research into automating the application of intelligent techniques for datamining. 'To turn datamining into a tool for end users, this expertise must be made explicit, formalised and where possible automated. This should include steps that go beyond the core algorithm, such as data preparation and evaluation.'

The purpose of the description task was to provide marketing personnel with a clear insight as to why customers owned a caravan insurance policy and how these customers differed from other customers. Descriptions could take any form, but were required to be comprehensible, useful and actionable for a marketing professional with no prior knowledge of computational learning technology.

As with the prediction solutions, a wide variety of approaches were employed, although there was a tendency to use rule based solutions. Submitted descriptions were evaluated by a marketing expert, Stephan van Heusden from MSP Associates in Amsterdam, who remarked that despite some excellent entries in this task, most lacked a good verbal description: 'Participants seem to forget that most marketeers find it difficult to read – and understand – statistics!'

'If the valuable but complex patterns which are detected by advanced intelligent techniques are not explained properly, end users like marketeers will still prefer the simple but crude and limited solutions,' observes van der Putten.

'In real world situations, the value of a blackbox model that simply provides predictions is very limited. User acceptance of such models is low, but more importantly mere predictions do not contribute to the incremental build up of knowledge about the problem. For instance, in this marketing problem, the scores can only be used to select prospects to mail or call, but they gave no insight as to what kind of message should be sent. There is no input for mass media campaigns such as commercials. Finally, there are no clues as to how to improve the model.'

Nevertheless, it's not all bad news for so-called blackbox techniques such as neural nets. 'There is no reason why you shouldn't use different algorithms for the description and the prediction problem,' van der Putten points out.

He believes that competitions such as the CoIL Challenge highlight the difference between the requirements of the laboratory and the requirements of the real world. 'I wouldn't be surprised if an experienced database marketeer outperformed the rocket scientists on this kind of problem,' he says. And it is for this reason that he suggests that competitions have a role to play in the field of meta-learning – learning to use self-learning techniques.

He draws an analogy with zoology where biologists perform experiments in controlled lab conditions in order to understand animal behaviour and abilities. 'Ethologists oppose this controlled approach and would rather perform field research where they observe animal behaviour in the wild. In meta-learning the current approach is to derive characteristics from datasets and algorithms and run a lot of controlled experiments to discover how features of the dataset relate to features of the algorithm, select the best algorithm for a given task and so on. This is very valuable research, but a more ethological approach to meta-learning might add new insights. Competition conditions are far from controlled, but perhaps they better represent how an average user interacts with these technologies. Algorithms that are very complex to use, for instance, would score well in a controlled environment, but badly in a competition.'

- **For more information, visit http://www.dcs.napier.ac.uk/coil/challenge/ – this site includes a link to *CoIL Challenge 2000: The Insurance Company Case*, a report by Peter van der Putten and Maarten van Someren with details of submissions and results.**
- **A detailed problem description and data sets are available at http://www.wi.leidenuniv.nl/~putten/library/cc2000/, homepage of The Insurance Company (TIC) Benchmark. This datamining benchmark dataset is ideally suited for testing algorithms or for use in lab sessions and is freely available for non-commercial use.**

THE WINNERS

The winner of the predictive task was Charles Elkan from the Department of Computer Science and Engineering, University of California, who used a naïve bayes approach. Apart from 'Purchasing Power Class', all socio-demographic variables were discarded and two derived variables were introduced that summarised the usage of car and fire policy products.

The winners of the descriptive task were YongSeog Kim and W. Nick Street, from the Management Sciences Department, University of Iowa, who used evolutionary local search algorithms (ELSA), chi square tests and association rules.

Evolutionary Computation Fuzzy Logic Machine Learning **Neural Networks**

Researchers at Sheffield University Department of Automatic Control and Systems Engineering have proposed a method that combines the predictive capabilities of neural networks with the search capabilities of genetic algorithms to find optimal inputs for steel alloy design. Although the mechanical properties of steel generated by the heat treatment process are well understood, those of a range of alloy steels are less easy to predict. However, recently-developed neural network models can predict mechanical properties for a wide range of steel alloys and have been shown to improve product reliability and process efficiency. Incorporating these models in the fitness function of a genetic algorithm (GA), the researchers demonstrated the GA's ability to adjust five variables (carbon, manganese, chromium, molybdenum and tempering temperature) to produce new alloy designs with pre-specified values for Ultimate Tensile Strength (UTS) and Reduction of Area (ROA). If no constraints were placed on the GA search, evolved solutions were commercially impractical. However, further experiments showed that when standard deviations relating to the predictions were included as a penalty term in the GA's fitness function, the GA produced a reliable and practical solution which was later verified by several metallurgists. Further research, which will include material costs in a multi-objective context, is planned.

For further information, contact M. Mahfouf: **M.Mahfouf@Sheffield.ac.uk**

CoIL Summer School

A springboard for collaboration

One of the main aims of the CoIL Summer School, held at Limerick in late August, was to act as a springboard for collaborations between new researchers from different backgrounds. With only seven days within which to achieve such an ambitious aim, the event was designed to focus on teamwork, co-operative problem solving and hands on learning.

During the first two days of the Summer School, 'seniors' – all of whom were experts in a particular area of computational intelligence – presented tutorials and posed key problems from a variety of application areas: robotics, image analysis, financial forecasting, control, bioinformatics, classification, network security, data mining and dynamic optimisation.

Participants then divided into teams of four and spent the remainder of the week attempting to solve one of the problems using one or more of the methods presented in the tutorials. The seniors remained on hand to provide advice and consultation, and by the end of the week each team had produced a draft paper describing their work.

'This summer school hit the bullseye. Not only did I meet wonderful new people, I'm sure that the work we did in that short time will be finished as a team and, hopefully, it will be published in the near future,' commented one of the participants.

'We were very pleased to see so much useful exchange of ideas in such a short time,' says Local Organiser Conor Ryan.

'The CoIL Summer School was an excellent opportunity for people from a variety of backgrounds and experiences to benefit from learning from each other. Working in teams, everyone contributed something but also gained ideas from their team partners and from the expertise of the Summer School seniors. In a very short time eight draft papers were produced and some of these are bound to go forward to publication.'

In a very short time eight draft papers were produced and some of these are bound to go forward to publication.

Winners of the 'Best Paper' award: Clarissa van Hoyweghen, Katja Verbeeck, Eduard Lukschandl, Jano van Hemert.

Inside this issue of Synergy

SNAPSHOTS of

Synergy talks to Marc Thuillard, the man behind the world's first mass produced fuzzy-wavelet device.

What benefits does wavelet analysis offer over those supplied by traditional Fourier analysis?
MT: In many of its aspects, wavelet theory can be regarded as an extension of Fourier formalism. Wavelet theory offers a very flexible and computationally efficient alternative to the short-time Fourier transform.

Instead of projecting the signal on sines and cosines, the signal is projected on a set of generally well-localised wavelet functions. Wavelets allow good resolution in both time and frequency, and, in contrast to the short-time Fourier transform, the wavelet transform is easily invertible.

What new activities/applications does it make possible?
MT: Multiresolution analysis has become a quite standard tool in signal processing. Wavelet theory has been applied to basically all scientific fields, including domains as different as quantum mechanics, econometrics and social sciences.

Despite the large variety of wavelet applications, the main domain of applications is still in image processing. The new standard JPEG 2000, for instance, is based on wavelet data compression schemes.

What new difficulties does it present for the practitioner?
MT: The main difficulty is understanding the mathematics behind wavelet theory in order to apply it correctly. Once the theory is mastered, multiresolution analysis is easily implemented. As wavelet analysis offers a greater flexibility than Fourier analysis, there is also more room for optimisation.

Is wavelet analysis less computationally expensive than Fourier analysis?
MT: The fast Fourier transform reduces from N^2 to $N \log_2 N$ the number of necessary operations for a Fourier transform of a signal with N values. A wavelet decomposition using the fast wavelet decomposition algorithm is slightly more efficient, as O(N) operations are necessary.

What was the motivation for combining fuzzy logic with multiresolution analysis in a commercial flame detector?
MT: Price and power consumption are important issues in sensorics. Combining wavelet theory and fuzzy logic in a single method furnishes a very efficient means to conduct both a spectral analysis and a classification task in a single step. This would not have been practically feasible using conventional methods.

Synergy at Work
Wavelet Analysis and Fuzzy Logic

A team at Cerberus, a division of Siemens Building Technologies, have developed the world's first commercial device to combine wavelet analysis and fuzzy logic. Their WaveRex flame detector represents a major advance in fire prevention, as it can reliably distinguish between fire and a common cause of false alarms – reflected sunlight on water or foliage. WaveRex exploits the characteristic flicker frequencies of flames, measuring a flame at three different wavelengths and subjecting the recorded signals to a 'fuzzy-wavelet' analysis. WaveRex is the result of fundamental research into the physics and dynamics of flames, research that led to a new understanding of flame pulsation. The exploitation of these research results required an efficient spectral analysis and classification method. Short time Fourier transforms in combination with a classifier were considered, but the necessary power was too high for the microcontroller of choice. A solution was found by combining fuzzy logic and wavelet theory. As a result, WaveRex can swiftly recognise the fingerprint of a flame and distinguish a signal of a fire from a potential source of false alarms.

For more information, contact Marc Thuillard: **marc.thuillard@cerberus.ch**

Marc Thuillard, Development Manager at Cerberus, a division of Siemens Building Technologies.

> multiresolution analysis and wavelet theory are a natural complement to soft computing methods

What does multiresolution analysis have in common with soft computing?
MT: Multiresolution analysis is of central importance in the mechanisms of perception and decision. Humans are particularly good at such tasks. For instance, image processing in the brain relies heavily on the analysis of the signals at several levels of resolution. Extracting details of importance out of a flow of information is an essential part of any decision process. Soft computing covers a range of methods that are somewhat tolerant of imprecision, uncertainty and partial truth. Hybrid methods combining soft computing methods with wavelet theory have therefore the potential to accommodate two central elements of the human brain: the ability to select an appropriate resolution to the description of a problem and to be somewhat tolerant of imprecision.

Multiresolution analysis and wavelet theory are a natural complement to soft computing methods. Soft computing deals with solving computationally intensive

(continued on page 6)

UNCERTAINTY

How wavelets provide a clearer focus on frequency and time

Beneath all the noise, discontinuities and incomplete data that make many real-world signals so difficult to unravel, lies a dark secret that can never be cracked. Quantum physics doesn't have a monopoly on uncertainty principles, there's one at the very heart of signal processing. Just as you cannot simultaneously know the momentum and the position of a moving particle, you can never know the precise frequency components that exist in a given signal at a particular instant in time.

Yet the bulk of real world signals are non-stationary: their frequency components change over time. Moreover, for many applications the most interesting information is hidden in the time/frequency ratio of the signal. Whether you want to pick up the small scale changes in an ECG reading caused by myocardial ischemia, or to identify an individual whale from his characteristic song, it's the time/frequency signature that you'll be looking for.

Traditionally, short-time Fourier transforms (STFT) have been used to slide a window across the different frequency and time bands within a signal and so glean information about what frequency bands exist at what time intervals. However, the problem with the STFT is that it gives a fixed resolution at all times. Once you have selected the size of the 'window' you wish to open on your data, you can't vary its time/frequency aspect ratio. It's a one-size-fits-all approach that frequently leaves engineers on the prongs of a dilemma: choose a wide window and gain good frequency resolution at the expense of poor time resolution, or choose a narrow window and gain good time resolution at the expense of poor frequency resolution.

For decades, researchers in fields as diverse as nuclear engineering and neurophysiology have been trying to box clever with uncertainty, to extract what trade-offs they can from a principle that dictates that more accurate time information can be obtained about higher frequencies, while more accurate frequency information can be obtained about lower frequencies. The result is multi-resolution analysis or wavelet theory.

In a sense wavelet analysis can be said to bow to the laws of physics, being specifically designed to offer good time resolution at high frequencies and good frequency resolution at low frequencies.

It can also be said to bow to the practicalities of signal processing in the real world, where high frequency components generally appear in intermittent bursts and low frequency components tend to be more persistent.

This ability to analyse data at multiple resolutions, together with the fact that wavelet decomposition is invertible (a signal can be perfectly reconstructed from its wavelet coefficients) is what chiefly distinguishes wavelet analysis from other methods. In effect, it acts as an adjustable zoom lens, resolving a signal at each scale in terms of differences and averages, so that one can focus on different levels of detail.

Wavelet analysis breaks with the sinusoids of Fourier analysis. Basis functions are no longer prescribed, rather new ones, suitable for different datasets or operations, are continually being described. The key idea is that any signal can be expressed as a linear combination of functions, all of which are simply dilations or contractions of a single mother function – the prototype wavelet.

Having selected a mother wavelet appropriate to your particular data, this becomes the prototype for every window opened on the time/frequency plane during the transformation process. Basically each window is a scaled (dilated or contracted) and shifted version of the mother wavelet, with temporal analysis obtaining more accurate results with contracted short time windows and frequency analysis obtaining more accurate results with dilated long time windows.

The discovery of the fast Fourier transform has had a profound impact on engineering. Similarly, the discovery of the fast wavelet transform marks the beginning of the wavelet era. This algorithm furnishes a simple and efficient way to make a wavelet decomposition. A fast wavelet decomposition is carried out with a cascade of two filters. The first filter acts as a low-pass filter (smoothing the signal) while the second high-pass filter furnishes the details of the signal (the wavelet coefficients). The details of the signal at a lower level of resolution are obtained by applying iteratively the same two filters to the smoothed signal. The overall picture that results could be called a stylisation of the original signal – a caricature, which highlights features of interest such as spikes, discontinuities and periodic components.

Not surprisingly, then, wavelet analysis is becoming an increasingly popular technique for denoising, data compression and pattern identification or matching. It's an approach that's rapidly finding a role in areas as diverse as earthquake prediction and biomedical signal processing. Wavelets have been used to watermark digital images and to increase the level of detail captured in satellite images. They are central to the compression capability of JPEG2000, and the FBI's central, searchable database of 30 million sets of fingerprints.

They make applications possible that were once far out of reach. Just how effective wavelets can be in homing in on a tiny signal amidst a sea of noise, has been demonstrated by engineers at the US Department of Energy's national laboratory, who used wavelet analysis within a new device capable of detecting the sound of a heartbeat even when a person is hidden inside a large vehicle – a technology that, had it been in use in Europe earlier this year, might have saved the tragic loss of 58 lives at Dover.

Wavelets on the Web

- **J. Altmann's Wavelet Tutorial**
  http://www. wavelet.org/wavelet/tutorial/ wavelet.htm
- **Amara Graps' Introduction to Wavelets**
  http://www.amara.com/IEEEwave/IEEEwavelet.html
- **Robi Polikar's Engineer's Guide to Wavelet Analysis**
  http://www.public.iastate.edu/~rpolikar/WAVELETS/WTtutorial.html

SNAPSHOTS of UNCERTAINTY

(continued from page 4)

problems with a limited amount of computing power and memory by giving up some precision. Multiresolution analysis can be used to determine how and where to give up the precision.

Where do you think the chief synergies between wavelet analysis and fuzzy logic lie?

MT: There are two very distinct domains in which fuzzy logic and wavelet theory complement each other. I have already mentioned the first: the combination of wavelet theory and fuzzy logic is a very efficient method for spectral analysis, as it furnishes a natural method to extend fuzzy logic to the frequency domain. An example of a fuzzy rule in the frequency domain may be: if high-frequency component of the signal is large then.... The degree of membership is computed from the wavelet coefficients.

The second chief synergy is in learning. Wavelet modelling can be used to develop a fuzzy system automatically from a set of data. It takes advantage of the connection between wavelet-based spline modelling and the Takagi-Sugeno fuzzy model. One starts from a dictionary of pre-defined membership functions forming a multiresolution. The membership functions are dilated and translated versions of a scaling function. Appropriate fuzzy rules and membership functions are determined either through a classical wavelet method or using a wavelet-network.

Multiresolution-based fuzzy methods furnish a new approach to the problem of transparency and linguistic interpretability. Linguistic interpretability is included per design by using predefined membership functions forming a multiresolution. Membership functions are chosen among a dictionary and describe terms such as „very small“ or „large“ that do not change during learning. A fuzzy system developed with this method consists of a number of rules using membership functions with clear linguistic interpretations.

Linguistic interpretability and transparency are slightly different concepts. Linguistic interpretability refers to the ease with which rules can be interpreted in natural language, for example: „if temperature is low then heater is on“. Transparency refers to the ease with which a system may be understood by the human operator. A preliminary condition for transparency is a natural linguistic interpretability of rules. A second condition is that the number of rules and the number of different levels in a hierarchical fuzzy system is still manageable by human experts. The fuzzy-wavelet approach furnishes innovative solutions to developing transparent fuzzy systems.

Do you know of any applications that combine these two techniques?

MT: Sensorics and image processing are two important application domains for multiresolution hybrid techniques. As the wavelet coefficients represent the fluctuating part of a signal, wavelet-based methods have been used for contour extraction, segmentation or time series analysis.

Applications range from computer tomography and anaesthesia control to fire detectors. As an example, Abbod and Linkens (1998) use the wavelet transform of the auditory evoked potential to infer the depth of anaesthesia. Heart rate and blood pressure measurements are combined with the wavelet features in a rule-based fuzzy system, developed in close collaboration with anaesthetists.

What future fuzzy-wavelet applications do you envisage?

MT: Besides control and sensorics, some promising fields are quality control (vibration monitoring), texture analysis, financial analysis and forecasts.

Where do you think the chief synergies between wavelet analysis and neural networks lie? Do you know of any applications that combine the two techniques and what future applications do you envisage?

MT: Over 1000 articles have combined elements of wavelet theory with soft computing. The vast majority of these use wavelet analysis in combination with neural networks, either for feature extraction or to reduce the dimension of the input space.

Wavelets have also been combined with neural networks in wavelet networks and wavenets. We have recently extended wavenets to biorthogonal constructions. This development allows wavenets to be used to develop fuzzy rules from data in online problems. The main advantage of this method is that wavelet theory furnishes a simple means to validate the fuzzy rules.

A number of interesting applications have taken advantage of the multiresolution properties of wavelet networks. Engine knock detection systems have been developed on the basis of wavelet networks. Wavelet networks have been successfully implemented to identify and classify rapidly varying signals, for instance to identify high risk patients in cardiology or for echo cancellation.

Forecasting and prediction of chaotic signals are two other promising fields. The multiresolution character of wavelets allows long term and short term variations to be captured. Applications in forecasting range from economics to prediction of short term load in power stations.

Synergy at Work

Wavelet Analysis, Fuzzy Logic and Neural Networks

Researchers at the University of the Aegean and MIT GmbH have developed a system based on wavelet transforms, neural networks and fuzzy logic in order to analyse the noisy time-series data typical of changing price values in equities on the Stock Market. Their aim is to construct a prediction and decision tool for investors, using very short-term rates of change to arrive at a short-term buy/sell/hold policy. In order to uncover trends in the daily rate of change of closing values of a selected equity, they used a filtering system based on wavelet decomposition, thresholding and reconstruction. This applies multi-resolution analysis to the primary inputs, so that only those parts of the initial signal's spectrum that might be considered periodic are preserved. The reconstructed signal was used to train a multi-layer perceptron in daily-trend prediction. Outputs from the trained neural net are fed into a decision system based on a fuzzy set, which makes use of the predicted trend to arrive at a final buy/sell/hold strategy. The researchers used learning techniques such as genetic algorithms and neuro-fuzzy techniques to select the most suitable membership functions and rule base for the fuzzy set.

For more information, contact Giorgos Dounias: **g.dounias@aegean.gr**

Does the choice of sample size and basis function mean that wavelet analysis presents practitioners with a search problem – to choose resolutions and basis functions capable of identifying the features of interest in the data?

MT: This is very much so. The most popular methods are the so-called best basis and the matching pursuit algorithm, a greedy algorithm.

> **Hybrid methods combining soft computing with wavelet theory have the potential to accommodate two central elements of the human brain: the ability to select an appropriate resolution to the description of a problem and to be somewhat tolerant of imprecision.**

Is this a problem suitable for evolutionary computing?

MT: Evolutionary computing has been used with success in some very specific problems. Some promising research results have been obtained by combining the matching pursuit algorithm and an evolutionary strategy for determining the best functions from a dictionary of basis functions.

How else might evolutionary algorithms be combined with wavelet analysis?

MT: Somewhat provocatively, one may claim that wavelet theory has already had a significant impact on evolutionary computing. Consider the analysis of deceptivity in terms of the Walsh partition functions. Walsh functions belong to the 'arsenal' of multiresolution analysis; the functions associated with the full tree Haar wavelet decomposition are the Walsh partition functions.

The combination of wavelet theory with evolutionary computing is probably an area of research with a high potential for future development. Indeed, multiresolution analysis is subjacent in several areas of evolutionary computing. Let us think of the building block hypothesis or the adaptive search methods implementing 'zooming' operators.

At COIL'2000 we presented a simple genetic algorithm that makes explicit some of the connections between Haar wavelets and genetic algorithms. The algorithm uses a single operator that tries to catch some of the main features of the crossover and mutation operators. The simplicity of the model allows the derivation of simple analytical results, a somewhat rare matter in evolutionary computing. In particular, the expected population can be computed in terms of the wavelet coefficients of the fitness function. Further, as wavelet theory can be expressed within the framework of filter theory, through subband coding, the results can be reformulated in terms of filters.
This simple model uncovers some important relationships between sampling theory, multiresolution analysis and filter theory.

Synergy at Work

Wavelet Analysis and Neural Networks

Researchers at Katholieke Universiteit Leuven have combined wavelet denoising with a neural network to isolate and classify myoelectric signals indicative of driver fatigue. Electromyogram signals contain transient signals related to which muscle groups have been activated and to other physiological characteristics such as the presence of fatigue. Such signals are extremely noisy due to the superposition of several muscle activations filtered through different transfer paths of the surrounding tissues and the skin itself. However, the team, who performed electromyography (EMG) measurements on different muscle groups of a volunteer driver, were able to show that wavelet based denoising can very efficiently isolate the myoelectric activity bursts from the rest of the signal and hence isolate coordinated muscle activity. The reconstructed, de-noised signals were used to train a self-organising map to detect the presence of fatigue in the volunteer driver. The research team believes that the proposed techniques may be of use within a fatigue detection and prevention system. Their initial project was funded by the European Community (project TRANSWHEEL TIDE 3013), and further research is to be supported within the EC Fifth Framework (project SAFEGUARD).

For more information, please contact George Papaioannou:
george.papaioannou@flok.kuleuven.ac.be

About ERUDIT

Most of the concepts we deal with as human beings exist on a continuum. Warm shades off into cold; tall shades into short; true shades into false. The boundaries between the concepts we use to express things to one another are fuzzy and imprecise.

Computers on the other hand inhabit a binary all-or-nothing world where boundaries are sharply defined and things are either true or untrue.

Fuzzy logic provides a bridge between the continuous world of our perceptions and the digital world of computers. It is a precise mathematical approach for translating the fuzziness of linguistic concepts into a form that computers can understand and manipulate.

Because fuzzy logic can transform linguistic variables into numerical variables without jettisoning partial truth along the way it allows us to construct vastly improved models of human reasoning and expert knowledge. Fuzzy technology is already becoming an essential component of advanced control and decision support systems, and features in myriad electronic consumer products.

To ensure that Europe is at the leading edge of fuzzy application and research, the European Commission has funded ERUDIT, The Network of Excellence in Fuzzy Technology and Uncertainty.

ERUDIT's mission is to foster innovation, training, technology transfer and the exchange of best practice. The Network's website (http://www.erudit.de/) provides a contact point for institutions interested in collaborative projects, as well as information about its member nodes, applications, products, conferences and training events.

To find out more about ERUDIT, or to join the Network, contact:
ERUDIT Service Center, Promenade 9
52076 Aachen, Germany
Tel: +49 2408 6969
Email: erudit@mitgmbh.de

COMPUTATIONAL FLU
— *an intelligent approach*

How HONDA is using soft computing techniques to power the next generation of jet propelled aircraft.

The FX-20 turboprop engine, which Honda is currently developing for the MH-02 prototype aeroplane.

Working within the Future Technology Research Division at Honda R&D Europe, the Evolutionary and Learning Technology group (EL-TEC) combines neural networks, fuzzy systems and evolutionary computing methods with research in computational neuroscience and theoretical evolution. Their aims reflect Honda's commercial interests in engine development and mobility in the broader sense.

'In biological systems evolution is the designing process for neural systems while fuzzy systems are, perhaps, more appropriate to model higher level cognitive behaviour.'

'Reducing fuel consumption or, in more general terms, sparing resources is a major research effort,' EL-TEC Team Leader Bernhard Sendhoff explains. 'However, there are many objectives, such as building intelligent and adaptive systems that are environmentally friendly, both with respect to their operation as well as to their production, and that offer maximal safety.'

One of the group's projects has been to develop new designs for a transonic compressor cascade – a series of aerofoil layers, consisting of rotor and stator blades, found within the gas turbines that power jet propelled aircraft.

Essentially this is a fluid dynamics optimisation problem: the streaming conditions around the fixed 'stator' blades (which optimise the air flow angle), and the rotor blades (which move and propel the aircraft) determine the efficiency of the whole turbine. However, reflections from one blade have an impact on the properties of the air stream around following blades, so all the variables – shape, size, position and angle of the blades – are interdependent.

'Given the computation time required for the fluid dynamics calculations, optimisation only became feasible a few years ago,' comments Sendhoff. Yet even with Moore doublings, computation time is still at a premium and any device that saves time brings more complex problems within reach. The device in this case is the artificial neural network.

Rather than use computational fluid dynamic calculations in combination with some sort of turbulence model to numerically solve the Navier-Stokes equation and so evaluate each potential new design, a neural network was trained on data calculated by the CFD simulation program and then employed for the evaluation of the individual solutions.

'We use neural networks as meta-models to simulate the streaming conditions and as a part of the fitness function of an evolutionary algorithm,' says Sendhoff. 'In this way, computation time can be saved – although in this case it is not merely a matter of saving some time, it is absolutely necessary in order to be able to optimise at all.'

The decision to use evolutionary search to flush out potentially good designs was due partly to the global nature of the problem – a small change to the value of one variable may have a considerable knock-on effect on the system as a whole.

Synergy at Work

Evolutionary Computation **Fuzzy Logic** Machine Learning **Neural Networks**

Scientists at the Casaccia Research Centre in Italy have developed a system that uses fluid dynamic models, combined with a neural network and fuzzy logic decision maker, to measure the mass flow rate of three-phase (oil/water/gas) flows in oil pipelines. In the oil industry, knowledge of mass flow rate helps in predicting well exhaustion, detecting pipeline leaks and managing transportation and storage. However, measurement is difficult as the mix of oil, gas and water varies between wells and during the lifetime of the same well. This leads to variations in liquid viscosity, density and flow pattern so that no single computational fluid dynamic (CFD) model can be used to estimate flow rate over the full range of operating conditions. The classical approach has been to use several CFD models and to switch between them depending on the situation. However, the team at Casaccia designed a system in which four CFD models co-operated with a neural network using a meta-decision maker based on fuzzy theory. The resulting system, which has been installed in the AGIP oil field in Trecate, is able to cope with the highly non-linear interactions of the target data, giving higher precision results with an average error of between 3% and 7% compared to the 5% to 11% error of the best classical system. Indications of measurement reliability are also produced and the fuzzy rules allow the system to work in conditions far from those under which the neural net was trained and in the presence of sensor failures.

For more information, contact Stefano Pizzuti: **stefano.pizzuti@casaccia.enea.it**

UID DYNAMICS

'Compared to gradient focused search techniques (where the gradient has to be estimated, which is in itself extremely time-consuming), evolutionary algorithms offer the possibility of being more global and due to their stochastic nature leave more room for design surprises,' says Sendhoff.

The starting point was to optimise a 2D description of the blades. With this under their belts, the four German-based members of the EL-TEC team and their colleagues in Tokyo have turned their attention to the formidable task of 3D optimisation.

'As a first approximation the fact that the real blade is 3-dimensional does not matter, because the variations in the third dimension are much smaller,' Sendhoff points out. 'However, the more serious you get about the optimisation the more important it is to encode the full blade. The practical benefits are twofold: firstly you have more freedom in the design and secondly the CFD calculations are much more precise – although at the expense of very long computation times!

'In any case, the goal always is to design the most efficient blade with respect to all the various constraints.'

He believes that when it comes to complex problems such as this it can be extremely beneficial to hybridise different soft computing techniques, building on their strengths to construct a complete solution.

'Firstly, from a technical point of view, different techniques have specific strengths, for example, by combining neural networks and fuzzy systems, it is possible to integrate knowledge-based (for example from engineers) system design with data-based system design for which neural networks, in my belief, are more appropriate.

'Secondly, neural systems can be utilised for the fitness evaluation of individuals during evolutionary optimisation. Thirdly, and perhaps most importantly, there is a deeper connection, because in biological systems evolution is the designing process for neural systems while at the same time fuzzy systems are perhaps more appropriate to model higher level cognitive behaviour.'

The transonic compressor cascade is only one of a number of projects that the Evolutionary and Learning Technology group is pursuing – and perhaps not the most sensational. In collaboration with colleagues at the Honda Robotics Department in Tokyo they will apply soft-computing and methods from computational neuroscience to the control of the Honda humanoid robot in order to make it more autonomous and more human-like in its decision and interaction capabilities.

'Using techniques which are adaptive and robust in changing environments seems computationally prohibitive in this domain, and it is a fascinating research area,' comments Sendhoff. 'The Honda robot demands a very advanced and complex control architecture. We believe that our techniques might help to develop such a system.'

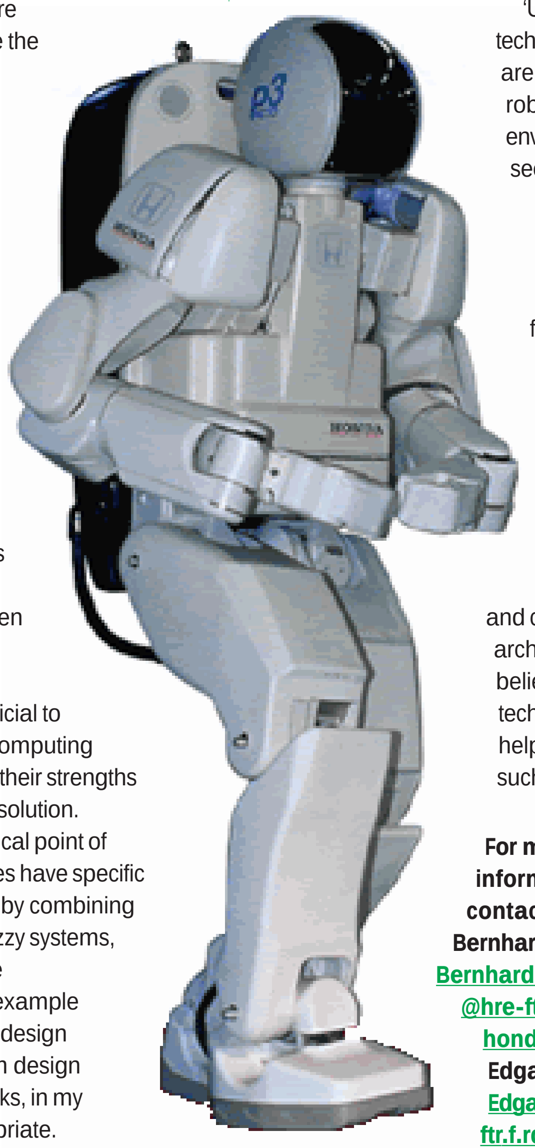

The Honda robot, billed as the most advanced humanoid robot in the world.

For more information contact:

Bernhard Sendhoff
Bernhard.Sendhoff@hre-ftr.f.rd.honda.co.jp

Edgar Körner
Edgar.Koerner@hre-ftr.f.rd.honda.co.jp

About EvoNet

Evolutionary computing harnesses the power of natural selection to turn computers into automatic optimisation and design tools.

The three mechanisms that drive evolution are reproduction, mutation and survival of the fittest. In the biological world these enable lifeforms to adapt to a particular environment over successive generations. The eagle's eye, the dolphin's sonar, the crafty human brain itself – all these solutions to environmental problems were generated by evolution. All bear witness to its power as a universal optimiser.

An evolutionary algorithm consists of a population of candidate solutions and a 'fitness function' which measures how well each individual within the population performs in terms of a designated task. Because the rules for generating offspring solutions from parent solutions mimic the mechanisms that drive evolution in nature, evolutionary computing can be used to breed progressively better solutions to a wide variety of complex problems.

The European Commission has recognised it as one of the important new technologies of our time and has funded a Network of Excellence in Evolutionary Computing, EvoNet, to assist in the transfer of knowledge and expertise to the manufacturing and service sectors.

The Network's website (http://evonet.dcs.napier.ac.uk/) hosts a comprehensive collection of training resources and an electronic repository of research papers. It also provides information about educational opportunities, conferences, applications, projects and recruitment.

To find out more about EvoNet, or to join the Network, contact:

EvoNet, School of Computing, Napier University, 219 Colinton Road, Edinburgh EH14 1DJ, United Kingdom
Tel: +44 (0) 131 455 4219
Email: evonet@dcs.napier.ac.uk

Sound*bytes* from

Representatives from MLNET, NeuroNet and ERUDIT took part in a panel debate on Intelligent Technologies at the COIL 2000 Symposium in Chios, Greece. The participants were: Maarten van Someren (Machine Learning), Christer Carlsson (Fuzzy Logic) and Georg Dorffner (Neural Networks).

What are the achievements of the COIL approaches?

Christer Carlsson

Let me give you the broad perspective: in my work, which is management decision-making and strategic planning, fuzzy logic has been the tool we move to when we come to the end of the road with standard tools. For instance, when we run into multi-criteria decision problems with conflicting criteria, we move on to fuzzy multiple criteria methods. When standard tools cannot make sense of large data sets, we use fuzzy logic based tools. It's a good way to carry on systematic decision making despite the fact we have spotty data or imprecision in our understanding of a problem. So it's something we move to when we run out of steam with the standard approaches.

Fuzzy control has become standard for control applications, so that points to one area where FL has made a significant impact, and it has proved useful in a multitude of applications where standard logic based approaches do not work any more.

As an overall comment, I don't believe in contrasting various intelligent technologies with each other. The real world is large enough for all these approaches to work in parallel, to learn from each other and to create synergies with each other. One context where we can work with all these approaches, and all combinations of them, is the virtual organisations created through new telecommunication technology. It's become apparent that in order to manage virtual organisations we need intelligent systems.

Virtual organisations are the way of the future; they consist of teams of people working together, over large geographical distances. In order to manage a virtual organisation you can work with two types of information. So called 'thin information' gives you numbers to work with. It's often said that thin information is like the body count MacNamara used in the Vietnam war. In contrast to that, 'thick information' is the reality faced by the soldiers in the jungles of Vietnam.

In order to manage virtual organisations, we need to develop tools and methods to handle large amounts of thick information. I have seen management models where a manager in a virtual organisation will have to handle 200 subordinates. Of course we need to stretch our imaginations if we want to figure out how to manage 200 people successfully with something called 'thick information'. So now we're looking for good metaphors to translate these technologies to management tools and I think this is challenging enough for all of us, with all our approaches, to address.

Maarten van Someren

The initial motivation of Machine Learning and Artificial Intelligence was to reproduce human learning. They looked at it as a functional problem: people are able to learn to recognise certain objects, to learn to perform tasks – so how can we achieve this functionality? AI used to believe that the way forward was to try to model the higher cognitive processes going on during learning and use these to develop algorithms and computer systems that can perform the same actions. This was the start of the field, and it has gradually developed a more technological orientation – at the moment there are very few people in ML who are trying to model human learning in the way that I just sketched. Almost everyone is addressing technical problems of data analysis.

If you look back to earlier years, the methods used were fairly specific and would not scale up to larger problems and to general classes of problems. This had changed by the end of the '80s and over the past ten years there's been an increasing range of methods for different types of learning problems.

In the past symbolic rather than numerical data was the main focus, but it's moving more towards combinations of numerical and symbolic data, textual data, sequences, more complex patterns and structures. People in ML think they have been successful in that they have developed effective algorithms for a wide class of problems of this data analysis type. The successes that people point to are the many industrial applications that have proven to be both simple (they produce fairly simple models for relatively complex data sets) and are still very effective (they predict very well).

Recent issues include modelling relational structures in which an object is modelled not as features but as relations between values. Other developments that people expect much from are hybrid methods combining different forms of representation. There are many cases in which you can demonstrate that these hybrid models are even better than the basic methods from which they are built. The other thing that people are proud of in ML is their methodology, which is characterised by a strong emphasis on cross validation, testing your methods on real data and publishing only if you can prove that they are generally effective and better than other methods.

One success story involved predicting how well the NASA space shuttle would land under certain weather conditions. NASA tried to build an expert system, based on a model of the space shuttle and a model of the weather conditions, that would allow it to predict if the operation would go well. After a year, the team of six knowledge engineers and technicians gave up. By coincidence two of them took part in a ML course a few months later and the teacher suggested that they should apply ML to the problem. Because they were able to simulate the shuttle's behaviour under different weather conditions, they created examples, running the simulator under different conditions and measuring what happened. Within a few weeks they solved their problem.

CoIL 2000

Georg Dorffner
The biggest achievement of neural networks is that they've made statisticians think about non-linearity – one strength of neural networks is they allow you to build models when most of the knowledge is in the data and that is where the non-linearity lies.

Another strength of neural networks is their probabilistic nature. Personally, I believe in probabilities, which is why I'm a little cautious of fuzzy systems because although both may deal with uncertainty, I trust more in probabilities than in other formulations of uncertainty. For instance, once you go into the Bayesian framework you have only to accept a few axioms and then everything comes naturally and you can prove it.

Another point I want to bring out is the inter-disciplinary nature of the field. As we all know, the original motivation behind neural networks was to model the brain. Nowadays most people are solving problems and analysing data, so the field doesn't really have much to do with neural science any more. Nevertheless you get neural scientists at neural network conferences and a lot of inspiration has come from that. For instance, now we have spiking neuron models and some first applications in signal and image processing. Since we have this inter-disciplinary character, we can hope for more inspirations and further achievements in the future.

If a funding body offered you a five year grant, what would you use it for?

Christer Carlsson
I would like to explore the possibility of creating value added services for third generation mobile technologies. This is an area where all the technologies we have been talking about today will make significant inroads. Third generation mobile technology is going to be a matter of everyday usage for more than 200 million people in Europe by 2003, if we calculate that the number of people in Europe is about 330 million and the penetration rate of mobile technology is going to reach more than 70% by 2003. Remember that third generation mobile technology includes business to business applications, business to consumer applications, electronic data interchange between companies, medical support services for people with diseases that are easily monitored; new tools for face to face education. As we understand it, third generation mobile technology is going to be something huge – the volume is going to be billions of Euros.

Georg Dorffner
I would use half of the grant to explore issues of modelling uncertainty. I think the Bayesian inference framework is very promising, although with respect to applications it is still very computationally intensive so you need some improvements. For many applications, it's really important to give an estimate not only of your output, but also of your uncertainty. So, for instance, if you have a diagnostic support system, it shouldn't only diagnose a disorder, it should also say how certain it is that its answer is correct.

I would spend the other half of the grant on investigating models that are closer to the brain and seeing whether they could be used effectively in industrial applications. We have made some mistakes in the past. We said, 'Wow, the perceptron! This is like the brain, it learns.' Then it turned out that it is just a linear discriminator – nothing new. So we have to put work into finding out whether, for instance, a spiking neural model brings in something new or is just the same old thing phrased in new terms.

Maarten van Someren
About a year ago we organised a competition and received a number of submissions and solutions. I thought it would be a good idea to write a scientific paper explaining why one method worked and another didn't. The methods that people used were fairly simple: decision tree learning, regression, neural nets... a range of things. There were a few for which I could find an explanation as to why they would or

(Continued on page 12)

About MLNET

The term 'machine learning' encompasses a variety of approaches that can derive general problem solving strategies from example problems.

Computer programs that learn automatically from experience promise far-reaching benefits both for industry and society as a whole – the ability to extract valuable information from jungles of data, increased reliability, improved human-computer interaction, the automation of monotonous or difficult tasks and that holy grail of computing: adaptable systems capable of tuning themselves to changing requirements. Already machine learning applications have proliferated in sectors as diverse as banking, medicine, astronomy, marketing, telecommunications and robotics.

In recognition of the importance of this technology to the future competitiveness of European industry the European Commission has funded MLNet, the European Network of Excellence in Machine Learning, Knowledge Acquisition and Case-Based Reasoning.

MLNet aims to extend the scope and applicability of computer learning by improving co-ordination between academia and industry, and by raising awareness of applications, technologies and scientific results between different communities.

The Network's online information service (http://www.mlnet.org) provides details about research projects, educational opportunities, tools, software, datasets and available expertise. MLNet also hosts a number of scientific, educational and industrial events.

To find out more about MLNet, or to join the Network, contact:
Maarten van Someren (MLNet Co-ordinator), Universiteit van Amsterdam, Roetersstraat 15, 1018 WB Amsterdam, The Netherlands
Email: maarten@swi.psy.uva.nl

(Continued from page 11)

wouldn't work, but for most I couldn't – and I did not have the conceptual model that would help me to answer this question. This makes me feel like a magician's apprentice. And I think this question must be addressed if we are to make real progress. We can give people more tools and some observations and experiments, but we will never be able to tell them which tool to use, when, if we cannot find some theory – some underlying explanation – as to why the methods work. That is where I would put my money.

How does one validate the success of a COIL approach to a problem?

Christer Carlsson
Validation is a standard part of the methodology of FL based applications. Success can be evaluated in different ways. The first step concerns description and explanation: does the FL based application help us to describe and explain things so that somebody gets an understanding of something that they could not comprehend before? The second step is: are we able to solve problems or make decisions that are in some sense better than before? Thirdly, if the end result of these activities shows up on the bottom line of a corporation, or in a social index showing that people are better off than before, that's the final step of validation.

Georg Dorffner
Neural networks provide a way to validate accuracy. Once you operate in a probabilistic framework you have ways of estimating how good your model is and how accurate it is. Bayesian approaches are just one way of doing so. I would say it is possible, although it is not always done, to not only build an estimate of your relationship in your data but also build an estimate of how accurate you are or, to put it another way, of how uncertain the model is.

To validate success you must take your basic statistics seriously. Never trust a single application of your model – if it worked once it won't necessarily work a second time. You must always have multiple runs, do averaging and see whether this confirms what you observe and if your performance level goes down, report that and not the performance level you get with one perhaps lucky run.

Maarten van Someren
There are some methods that allow real statistical analysis of the modelling error, but not many. This means that in practice you estimate. When you talk about validation you talk about the uncertainty in the model – how well will it predict on new data – and the standard procedure for this is cross validation.

Another point is the issue of comprehensibility. In many cases the model has an additional use; it's not just used for prediction, people want to learn from it, and in that case other forms of evaluation, other criteria, come into play.

One lesson we learned in ML is that there are many problems that the problem owner thinks are difficult but which are actually quite simple. I think everybody who has done this type of data analysis has the experience that people come with a large number of variables and in the end you find that maybe only 20% of the variables are necessary and adding more does not help at all. So although the owner of the problem believes that all these things are relevant, many problems are simpler than they seem to be. Simplicity of a model is a quality measure for two reasons: because in many cases simpler models are easier to understand and because, for statistical reasons, a simpler model that explains the data is often a better model than a more complicated one.

Are there situations where one or more COIL approaches are doomed to fail?

Christer Carlsson
Experience tells us that we're bound to fail from time to time and that is going to happen no matter what techniques we are using. The only thing that is really stupid is to go on making the same mistakes again and again. I think we should start looking for frameworks that are more flexible, so that if there appears to have been a mistake, we can quickly change the approach and go in and make corrections.

Georg Dorffner
You can answer this question by looking at the conditions that should be fulfilled when you apply a neural network. The first condition is to have a sufficient amount of data samples in relation to the amount of input features you have and the potential amount of degrees of freedom you will introduce into your neural network.

Another condition is that there should be evidence that simpler, traditional methods cannot fully capture the data complexity. Also, if you have knowledge and rules then use them; you should not try to use a neural network to replace them. The last condition refers to comprehensibility: if you want an understandable solution again you might fail with a neural network.

Maarten van Someren
The ML task is generally to construct a model from a set of data describing some phenomenon. However, you cannot say that ML fails because it does not refer to a single method: it's a whole range of methods. So the question from a ML view point must be, 'When do some methods fail and why?' One reason that they sometimes fail is because the class of models that can be generated by the method does not include the best model. The other enemy is overfitting, where the method constructs a model that fits well on the original data but is too specific and fails the next time it is applied.

This debate is summarised in a FAQ supplemented with digital audio files at: http://fuzzy.iau.dtu.dk/debate2.nsf

CoIL Roadmap

The CoIL Roadmap – an overview of the current status of CoIL technologies with perspectives for the future – will be completed before the end of the year. Please check http://www.dcs.napier.ac.uk/coil, where a link will be provided as soon as the Roadmap becomes available.

Obituary – Professor Agnessa Babloyantz

March 2000: Professor Agnessa Babloyantz of the Université Libre de Bruxelles has died. She was the Belgian Managing Node representative of NEuroNet since 1994.

Agnessa Babloyantz was born in 1932 in Tabriz, Iran. She obtained a degree in Chemistry and a PhD in Physical Chemistry from the Université Libre de Bruxelles (ULB). In 1991, on retirement from her position at the Faculty of Sciences of the ULB, she became a member of the research team of the International Solvay Institutes for Physics and Chemistry.

Professor Babloyantz's research was in the area of thermodynamics applied to basic biological systems, biopolymers and morphogenesis. She showed that patterns may arise in morphogenetic systems and investigated spontaneous temporal and spatio-temporal structures arising in neural tissue. She showed that epileptic seizure could result from a coherent behaviour of a large ensemble of neurons. Applying the techniques of non linear dynamics to heart and brain waves, she also demonstrated the presence of deterministic chaos both in the heart and several stages of brain waves. Agnessa Babloyantz used her knowledge of brain activity and multicellular assemblies to make progress in the field of artificial intelligence. She guided a large number of PhD students during her career and also influenced research at an international level in her fields of interest. Her book on non linear dynamics has been translated into several languages.

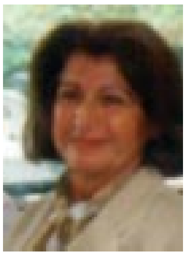

Neuroinformatics for 'living' artefacts
EC CALL FOR PROPOSALS

To stimulate research at the interface between the life sciences and information technology, the European Commission has launched a proactive initiative 'Neuroinformatics for living artefacts'. The aim of the initiative is to study basic methodologies and technologies that enable the construction of hardware/ software 'artefacts that live and grow'.

Proposals will be expected to use models or methods from the life sciences, and neuroscience in particular, as a catalyst for developing new methodologies that allow design, analysis and implementation of artefacts. The rationale for the initiative is two-pronged: to enable a leap in performance in terms of constructing, deploying and using intelligent entities and to project developed models back into the domain of life sciences to provide new insights and hypotheses.

The initiative is structured around three main themes, which may act as catalysts for individual proposals.

Theme 1: Adaptation of the 'living and growing'
Living systems adapt at multiple levels and time-scales to changes in their environment. Likewise, living artefacts should capitalise upon interactions among different levels of adaptation and growth. Research supported by this call should go beyond conventional adaptive algorithms taken in isolation and tackle issues such as: development and physical growth; learning and knowledge growth; open-ended evolution.

Theme 2: Embodiment and integration
A growing artefact has a body that allows development and adaptation and complies with the constraints of real-time operation, robustness and scalability. It is important that proposals go beyond the traditional assembly approach where the artefact is broken down into components leaving integration to a later stage.

Theme 3: Behaviour and organisation
A living artefact is an autonomous, self-sufficient and self-motivated system characterised by finite resources and responsive to interaction with the environment. Challenges which could be addressed by the design of the artefact include: emotion/self-motivation/curiosity; conflicts and inconsistency; finite resources; interactions/communication.

The deadline for receipt of proposals is 11 October 2000. For more information, please visit: http://www.cordis.lu/ist/fetni.htm

About NEuronet

Artificial neural network research takes its inspiration from the most impressive fielded application of a learning system available on Earth – the brain.

Animals learn through iterative adjustments to the synaptic nodes interconnecting the complex web of neurons in their brains. So, by analogy, does an artificial neural network. Implemented in hardware or software, it can be trained by exposing it to example inputs and target outputs, and iteratively adjusting the connection weights between a number of simple processing units.

The result is a quick-to-build, robust and adaptable problem-solver capable of generalising from previous examples when presented with new inputs.

Artificial neural networks are ideally suited to pattern recognition, signal processing, classification and prediction tasks. Their ability to interpret complex real world data means that they are already proving an invaluable tool in industry and commerce. To stimulate further research and accelerate technology transfer, the European Commission has funded a Network of Excellence in Neural Networks – NEuroNet.

The NEuroNet website (http://www.kcl.ac.uk/neuronet/) provides information about training materials, case studies, courses, products, projects, resources, contacts and events. To provide companies with direct access to information about best practice and available expertise, NEuroNet has also established industrial clubs in Denmark, Italy, The Netherlands, Spain, Finland and Greece.

To find out more about the Network, or to become a member, contact:
NEuroNet,
Department of Electronic Engineering,
King's College London, Strand,
London WC2R 2LS, United Kingdom
Tel: +44 (0)171 848 2388
Email: neuronet@kcl.ac.uk

CONFERENC

For further information and links to the events

19–20 October 2000
10th Annual GMA-Workshop Fuzzy Control
Dortmund, Germany
Initial deadline passed

25–27 October 2000
ADVIS 2000 1st Biennial International Conference on Advances in Information Systems
Izmir, Turkey
Initial deadline passed

25–27 October 2000
SEAL2000 3rd Asia-Pacific Conference on Simulated Evolution and Learning
Nagoya, Japan
Initial deadline passed

27–29 October 2000
Learning'00
Madrid, Spain
Initial deadline passed

30–31 October 2000
PAKM 2000 3rd International Conference on Practical Aspects of Knowledge Management
Basel, Switzerland
Initial deadline passed

3–5 November 2000
AAAI Fall Symposium 2000 Learning How to Do Things
Cape Cod, Massachusetts, USA
Initial deadline passed

3–5 November 2000
SIA2000 Socially Intelligent Agents – The Human in the Loop
North Falmouth, MA, USA
Initial deadline passed

5–8 November 2000
ANNIE 2000 Smart Engineering Systems Design Conference
St Louis, Missouri, USA
Initial deadline passed

6–11 November 2000
ACM CIKM'2000 9th International Conference on Information and Knowledge Management
Washington, DC, USA
Initial deadline passed

13–15 November 2000
ICTAI 2000 International Conference on Tools with AI
Vancouver, British Columbia, Canada
Initial deadline passed

14–16 November 2000
ICSB 2000 1st International Conference on Systems Biology
Tokyo, Japan
Initial deadline passed

14–18 November 2000
ICONIP-2000 7th International Conference on Neural Information Processing
Taejon, Korea
Initial deadline passed

16–18 November 2000
Workshop on Advances and Trends in Artificial Intelligence for Problem Solving (at SCCC 2000)
Santiago, Chile
Initial deadline passed

18–21 November 2000
ICCI2000 10th International Conference on Computing and Information
Kuwait
Initial deadline passed

19–22 November 2000
COMPLEX SYSTEMS 2000 5th International Conference on Complex Systems
Dunedin, New Zealand
Initial deadline passed

22–25 November 2000
SBRN'2000 6th Brazilian Symposium on Neural Networks
Rio de Janeiro, Brazil
Initial deadline passed

5–8 December 2000
ICARCV 2000 6th International Conference on Control, Automation, Robotics and Vision
Singapore
Initial deadline passed

11–13 December 2000
PKAW 2000 The 2000 Pacific Rim Knowledge Acquisition Workshop
Sydney, Australia
Initial deadline passed

12–14 December 2000
InTech'2000 International Conference on Intelligent Technologies
Bangkok, Thailand
Initial deadline passed

12–15 December 2000
BIS'2000 International ICSC Symposium on Biologically Inspired Systems
Sydney, Australia
Initial deadline passed

12–15 December 2000
ISA2000 International Congress on Intelligent Systems and Applications
Wollongong, Australia
Initial deadline passed

13–15 December 2000
IDEAL 2000 2nd International Conference on Intelligent Data Engineering and Automated Learning
Hong Kong
Initial deadline passed

14–15 December 2000
PlanSIG 2000 The 19th Workshop of the UK Planning and Scheduling Special Interest Group
Milton Keynes, UK
Initial deadline passed

16–20 December 2000
AISAT'2000: International Conference on Artificial Intelligence in Science and Technology
Hobart, Tasmania
Initial deadline passed

E CALENDAR

s listed here, please visit the CoIL website at http://www.dcs.napier.ac.uk/coil/

17–19 December 2000
KBCS 2000 International Conference on Knowledge Based Systems
Mumbai, India
Initial deadline passed

3–6 January 2001
8th International Workshop on Artificial Intelligence and Statistics
Florida, USA
Initial deadline passed

30 January – 3 February 2001
ASP-DAC 2001 Asia and South Pacific Design Automation Conference 2001
Yokohama, Japan
Initial deadline passed

11–15 February 2001
WSES International Conference on Fuzzy Sets and Fuzzy Systems, Neural Networks and Evolutionary Computing
Tenerife, Canary Islands
Deadline (papers) 30 October 2000

7–9 March 2001
EMO'01 First International Conference on Evolutionary Multi-Criterion Optimisation
Zurich, Switzerland
Deadline (papers/abstracts) 16 October 2000

11–14 March 2001
SAC 2001 16th ACM Symposium on Applied Computing
Las Vegas, USA
Initial deadline passed

20–23 March 2001
ISI'2001 International Congress on Information Science Innovations (including **IAM2001** Intelligent Automated Manufacturing 2001 and **ENAIS2001** Engineering of Natural and Artificial Intelligent Systems 2001)
Dubai, UAE
Initial deadline passed

18–20 April 2001
EuroGP2001 and **EvoWorkshops2001** 4th European Conference on Genetic Programming and EvoNet Workshops on Evolutionary Computation
Lake Como, Italy
Deadline (papers) 16 November 2000

22–25 April 2001
ICANNGA 2001 5th International Conference on Artificial Neural Networks and Genetic Algorithms
Prague, Czech Republic
Deadline (draft papers) 20 September 2000

25–27 April 2001
MOSIM'01 Industrial Systems Design, Analysis and Management
Troyes, France
Deadline 30 September 2000

14–18 May 2001
SNRFAI'2001 9th Spanish Symposium on Pattern Recognition and Image Analysis
Castellon, Spain
Deadline (papers) 30 November 2000

28–30 May 2001
ISATP 2001 4th International Symposium on Assembly and Task Planning
Fukuoka, Japan
Deadline (full papers) 15 October 2000

13–15 June 2001
IWANN'2001 6th International Work Conference on Artificial and Natural Neural Networks
Granada, Spain
Deadline (papers) 28 February 2001

19–22 June 2001
CIMA2001 Computational Intelligence – Methods and Applications
Bangor, Wales, UK
Deadline 31 October 2000

26–29 June 2001
SOCO/ISFI 2001 Soft Computing/Intelligent Systems for Industry 2001
Paisley, Scotland, UK
Deadline (abstracts) 30 November 2000

26–29 June 2001
EUROSIM 2001
Delft, The Netherlands
Initial deadline passed

4–6 July 2001
NOLCOS Nonlinear Control Systems
St Petersburg, Russia
Deadline (papers) 15 November 2000

25–28 July 2001
Joint 9th IFSA (International Fuzzy Systems Association) World Congress and 20th NAFIPS (North American Fuzzy Information Processing Society) International Conference
Vancouver, Canada
Deadline (abstracts) 15 December 2000

27–30 July 2001
CONTEXT'01 3rd International Conference on Modeling and Using Context
Dundee, Scotland
Deadline (papers) 2 February 2001; (workshop proposals) 2 March 2001

2–7 September 2001
15th IFAC Symposium on Automatic Control in Aerospace
Forli, Italy
Deadline (papers) 1 October 2000

4–7 September 2001
ECC2001 European Control Conference
Porto, Portugal
Deadline (draft papers) 1 October 2000

29 November – 2 December 2001
ICDM '01 The 2001 IEEE International Conference on Data Mining
Silicon Valley, California, USA
Deadline (papers) 15 June 2001

2–5 December 2001
FUZZ-IEEE 2001 10th IEEE International Conference on Fuzzy Systems
Melbourne, Australia
Deadline (workshop proposals) 17 November 2000; (tutorial proposals) 1 June 2001; (papers) 2 March 2001

CONTACTS

CoIL is structured around six *Actions*, each of which is designed to raise awareness about computational intelligence and learning, and to promote co-operation between the four participating networks.

Action C1 Roadmap

The goal of this Action is to produce a technological roadmap providing an overview of the state of the art in computational intelligence and learning, and indicating likely avenues for future development.

Erudit contacts

- Kauko Leiviskae kauko.leiviska@oulu.fi
- Hans-Juergen Zimmermann zi@or.rwth-aachen.de

EvoNet contacts

- Wolfgang Banzhaf banzhaf@tarantoga.informatik.uni-dortmund.de
- Colin Reeves srx014@coventry.ac.uk
- George Smith gds@sys.uea.ac.uk

MLNet contacts

- Lorenza Saitta saitta@al.unipmn.it
- Maarten van Someren maarten@swi.psy.uva.nl
- Lenie Zandvliet lenie@swi.psy.uva.nl

NEuroNet contacts

- Georg Dorffner georg@ai.univie.ac.at
- Stan Gielen stan@mbfys.kun.nl
- Erkki Oja erkki.oja@hut.fi

To email all listed contacts for Action C1 use the alias coil-roadmap@dcs.napier.ac.uk.

Action C2 Workshops

This Action aims to foster collaboration between networks by supporting joint activities and organising common workshops.

Contact

- Lenie Zandvliet lenie@swi.psy.uva.nl

About CoILWeb

Website manager: Chris Osborne

URL: http://www.dcs.napier.ac.uk/coil

CoILWeb provides information about all CoIL activities, together with contact information, training materials and useful links.

To contribute material for publication on CoILWeb, please use the alias: **coil-contribute@dcs.napier.ac.uk**

Action C3 Competitions

This Action aims to increase the awareness of available techniques through an annual competition.

Erudit contacts

- Neil Cade neil.cade@gecm.com
- Jens Strackeljan tmjs@odin.itm.tu-clausthal.de

EvoNet contacts

- Marc Shoenauer marc.schoenauer@polytechnique.fr
- Andea Tettamanzi genetica@tin.it

MLNet contacts

- Ashwin Srinavasan ashwin.srinavasan@comlab.ox.ac.uk
- Maarten van Someren maarten@swi.psy.uva.nl
- Lenie Zandvliet lenie@swi.psy.uva.nl

NEuroNet contacts

- Mark Plumbley mark.plumbley@kcl.ac.uk
- Peter van der Putten pvdputten@smr.nl

To email all listed contacts for Action C3 use the alias coil-competitions@dcs.napier.ac.uk.

Action C4 Education

This Action aims to provide easy access for members of one network to the main concepts, theories and techniques of the other networks. Each participating network will provide tutoring materials in the form of web-based distance learning courses and/ or lecture notes with demonstrations.

Erudit contacts

- Mariagrazia Dotoli dotoli@poliba.it
- Jan Jantzen jj@iau.dtu.dk

EvoNet contacts

- Gusz Eiben gusz@wi.leidenuniv.nl
- Zbigniew Michalewicz zbyszek@daimi.aau.dk
- Ben Paechter benp@dcs.napier.ac.uk
- Jim Smith im@ics.uwe.ac.uk

MLNet contacts

- Katharina Morik morik@kimo.informatik.uni-dortmund.de
- Maarten van Someren maarten@swi.psy.uva.nl
- Lenie Zandvliet lenie@swi.psy.uva.nl

NEuroNet contacts

- Daniele Caviglia daniele@dibe.unige.it
- José Dorronsoro dorron@vera.iic.uam.es

To email all listed contacts for Action C4 use the alias coil-education@dcs.napier.ac.uk.

Network Co-ordinators

MLNet: Maarten van Someren, Universiteit van Amsterdam, Roetersstraat 15, 1018 WB Amsterdam, The Netherlands

Erudit: Hans-Juergen Zimmermann, ERUDIT Service Center, Promenade 9 52076 Aachen, Germany

EvoNet: Terry Fogarty, School of Computing, South Bank University, 103 Borough Road, London E1 0AA, UK

NEuroNet: Mark Plumbley, Department of Electronic Engineering, King's College London, Strand, London WC2R 2LS, UK

Action C5 Communication

To promote information exchange between networks, this Action will focus on integrating and extending the member networks' existing electronic communication facilities.

Erudit contacts

- Sabina Heck erudit@mitgmbh.de
- Karl Leiven kl@mitgmbh.de

EvoNet contacts

- Juan Julian Merelo jmerelo@kal-el.ugr.es
- Chris Osborne c.j.osborne@dcs.napier.ac.uk

MLNet contacts

- Mathias Kirsten mathias.kirsten@gmd.de
- Stefan Wrobel wrobel@iws.cs.uni-magdeburg.de
- Lenie Zandvliet lenie@swi.psy.uva.nl

NEuroNet contacts

- Terhi Manuel-Garner terhi.manuel-garner@kcl.ac.uk
- Mark Plumbley mark.plumbley@kcl.ac.uk

To email all listed contacts for Action C5 use the alias coil-communication@dcs.napier.ac.uk.

Action C6 Coordination

This Action involves financial administration, preparing and reporting steering committee meetings, reporting to CEU and general management.

Erudit contacts

- Karl Leiven kl@mitgmbh.de
- Hans-Juergen Zimmermann zi@or.rwth-aachen.de

EvoNet contacts

- Terry Fogarty fogarttc@sbu.ac.uk
- Jennifer Willies jennifer@dcs.napier.ac.uk

MLNet contacts

- Maarten van Someren maarten@swi.psy.uva.nl
- Lenie Zandvliet lenie@swi.psy.uva.nl

NEuroNet contacts

- Mark Plumbley mark.plumbley@kcl.ac.uk

To email all listed contacts for Action C6 use the alias coil-coordination@dcs.napier.ac.uk.