

Applying Natural Computation to Real World Business - Selected Highlights from Cercia

Colin M. Frayn ¹⁾

1) Cercia, School of Computer Science, University of Birmingham
Edgbaston, UK, B15 2TT. <http://www.cs.bham.ac.uk/~cmf/>

Abstract: *Bringing cutting-edge natural computation research into industry is a challenging task, with numerous obstacles to overcome. In this paper I describe a few selected projects that we have developed at the Centre of Excellence for Research in Computational Intelligence and Applications (Cercia) and I discuss the promising future of Computational Intelligence techniques in industry.*

Keywords: Natural computation; industrial applications; optimization; data mining.

1. INTRODUCTION

Computational intelligence (CI) is an umbrella term which covers a number of diverse techniques such as Evolutionary computation, neural networks, agent-based models, simulated annealing and ant-colony optimisation. Most of these techniques also fall under the umbrella of 'natural computation' (NC) and, in particular, Evolutionary Computation (EC). [1]

These CI techniques distinguish themselves from conventional analytical techniques by several powerful properties:

- They are flexible
 - Which allows them to be applied to a very large range of challenges inside and outside of academia.
- They are adaptive and autonomous
 - Which means that they are able to learn from past experience with the minimal of human interaction.
- They are (often) decentralized
 - Which means that they often do not have a single point of weakness.
- They are easily parallelized
 - Many algorithms, especially evolutionary ones, can run perfectly in parallel with no loss of quality and a near-linear speed improvement.
- They are robust to errors
 - Which means that they are highly suitable for application to real-world problems, where data are very rarely neat, tidy or complete, and virtually never noiseless.
- They are simple
 - Though I appreciate that this might be controversial, and I also accept that the field of CI has a large number of potential pitfalls, I claim that the algorithms themselves are extremely

simple, which means that they are easy to teach and easy to implement. However, they are still capable of producing extremely complex and advanced results.

These properties make them ideally suited for many industrial applications, which very rarely conform to the ideal, theoretical applications on which such techniques are generally tested. Real-world data are very rarely well-behaved in the sense required by conventional algorithms.

A perfect example of this problem is with the Travelling Salesman Problem (TSP), which involves finding the shortest route between a number of fixed points by visiting each one once and once only. When dealing with the simplistic mathematical problem, CI techniques can be beaten by complex analytical techniques such as the cutting plane method [2][3].

However, once one considers real-world datasets, one encounters problems such as one-way streets, complex junctions, road gradients, fixed refuelling locations, etc. (see section 2). In this situation, it is relatively easy to extend an evolutionary algorithm to solve the TSP within a reasonable time, and moreover the CI method is an 'anytime' algorithm, producing a legal solution right from the first iteration and further refining this over time. However, conventional methods fail at this stage, as they cannot be modified in order to deal with such constraints.

Cercia

Cercia was funded primarily as a proof of concept for a host of computational intelligence technologies in order to assess their suitability for use in industry, as well as to educate the business community about the need to investigate, develop and adopt cutting-edge computational technologies. As such its goals are threefold:

Firstly, to continue the tradition of world-leading research that we have established in Birmingham, with a focus on developing technologies that will be of use to industry within a medium-long term timescale.

Secondly, to do work with industry applying current state-of-the-art techniques to real-world problems whilst guiding our own research and development towards more fruitful directions.

Thirdly, to carry out education within the sphere of natural computation, in order to ensure that the technological industries within the West Midlands are able to compete on a global scale in the 10-20 year future timeframe.

I share in this paper a number of examples of the above.

Firstly, two clear industrial applications, encountered as part of our work. Each one is the result of a collaboration between Cercia and one or more industrial partners, who are listed in the relevant sections.

Secondly, a selection of smaller, more classified or less-mature projects, each of which comprises of a brief description of the work undertaken together with selected further reading. These projects vary in direct industrial applicability from fully industrial projects right through to speculative research and teaching.

Any questions about the above should be addressed either to the author, or to the individual contacts listed in each section.

2. ROUTE OPTIMISATION FOR SALTING TRUCKS

In collaboration with Entice Technology Ltd, Birmingham.

The Challenge

Our partners, Entice Technology Ltd., have developed an advanced prediction model to calculate road surface temperature predictions based on climate models and geographical data, together with meteorological data gathered in real time from the government. They are able to predict which roads are likely to be covered with ice, especially after particularly cold winter nights. These roads then need to be gritted by the local council using a fleet of trucks especially designed for the purpose.

The challenge here is in designing a route for these trucks to follow, in order to optimise efficiency. That efficiency is in terms of distance travelled, grit spread, and time taken to cover a sufficient fraction of the road network. The road is also subject to a large number of strong constraints, including but not limited to the following:

- Certain roads will have more severe ice cover than others.
- Trucks only have a certain maximum fuel load.
- Trucks can only carry a maximum amount of grit.
- Some roads are one-way
- Some roads are more vital than others.
- All roads should be covered occasionally, just to

satisfy local tax payers!

- Different roads have different speed limits.
- Some roads are too narrow for the trucks to pass down.

A proposed route must optimise the following parameters:

- Minimise the total journey length, and thus minimise fuel cost
- Minimise the total journey time, and this grit the roads as quickly as possible. Related to minimising journey length, but also includes road types.
- Maximise the speed with which major roads are covered.
- Maximise the coverage of roads with the most severe ice risk, whilst minimising the coverage of roads with lower risk.

The Solution

Natural computation technologies, in particular evolutionary computation, provide an ideal solution for this complex challenge.

Capacitated Arc Routing Problems

Salting route optimization (SRO) can be regarded as an instance of the Capacitated Arc Routing Problem (CARP) [4][5]. Suppose that a graph is defined in which each edge has an associated cost. Additionally, a set of required edges and a demand is defined to each edge in the graph. There are several vehicles to fill the demands, where each vehicle has the predefined capacity of services for the demands. A depot is defined elsewhere in the graph. All vehicles must depart from this depot and return there at the end of their service tour. The problem is to find a set of tours which have a minimum total cost for all vehicles, ensuring the demands of all required edges are filled by at least one vehicle, whilst ensuring the total services capabilities of each vehicle are not exceeded.

Another overview of the project is found in [6]. More details about the exact evolutionary algorithm can be found in [7][8]. More details about the road surface temperature model can be found in [9][10].

Figures 1,2 show the road surface temperature for a marginal (near-freezing) night in the road network of the county of South Gloucestershire, England, together with its corresponding best-fit salting routes. The evolutionary method produced routes that were over 15% shorter in total than the conventional routes, designed by hand.

15% shorter routes equates to a corresponding saving in terms of time and cost, which also reduces environmental impact, and increases the speed with which vital arterial roads can become safe, potentially saving lives.

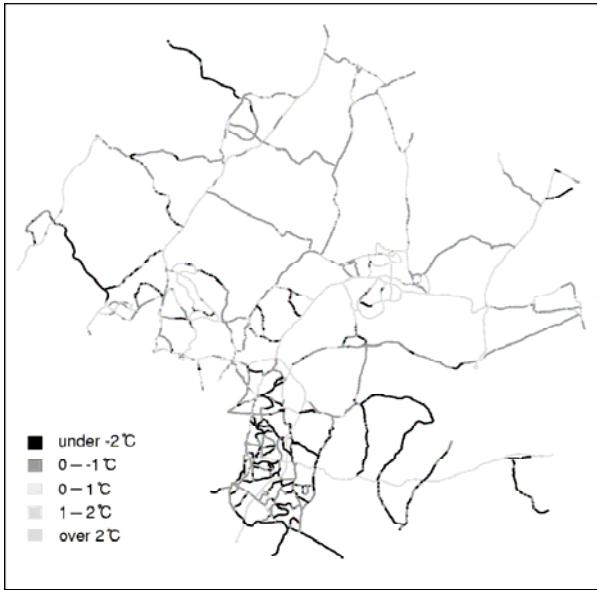


Figure 1 : Predicted road surface temperature graph for the county of South Gloucestershire, England.

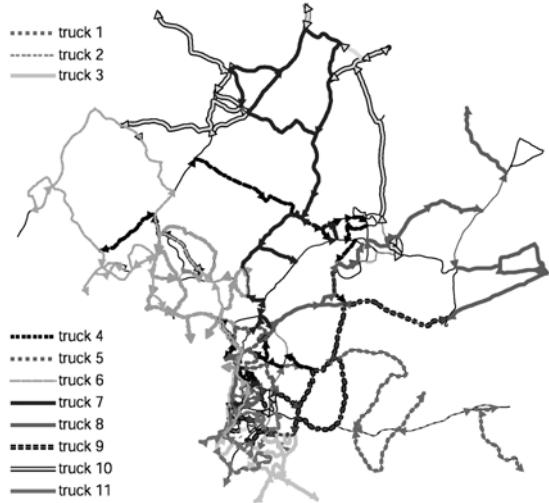


Figure 2 : Resultant salting routes for the network in figure 1.

3. STOCK FILTERING USING EVOLUTIONARY TECHNIQUES

In collaboration with the Investor's Chronicle (IC) magazine, Financial Times group, UK.

The Challenge

There are almost 2,000 individual companies listed on the London Stock Exchange. For each company, a large amount of financial information is available, including historical price movements, daily trading figures, earnings estimates, yield percentages. Together with these, one can also obtain figures for the aggregated index, allowing comparison with the

overall market sentiment for a given period. Also available are the interest rate and inflation values for the entire historical period.

Many stock analysis techniques have focused on analysing single stocks using signal processing techniques in order to estimate their likely forward prospects. However, this method seems largely fruitless, mainly because, at least with larger companies, one can assume that their stock price contains little or no useful information (the efficient markets hypothesis) and that, for smaller companies, price volatility caused by utterly external factors, such as related news or rumours, can make such methods fail.

There is much to be learned by instead producing a system that can take the entire list of all London quoted stocks, and then reduce that down to a smaller subset, around 10-20 stocks, which can then be thoroughly analysed by a human being. This way we combine the strengths of a computer in number crunching and data analysis, with the strengths of a human in textual investigation and instinct.

The Solution

Together with collaborator Dan Oakey at the IC, I obtained a dataset for the top 100 companies in the FTSE index, including all the above information on a daily basis back to 1990, or when the company was formed (whichever was most recent). To this, I applied two separate techniques from the field of evolutionary computation, namely genetic programming (GP) and evolutionary conjunctive rules (ECR) in order to develop stock screens – that is, sets of rules which can be applied to the full list of companies in order to reduce them to a more manageable number, with a strong bias towards those companies which are expected to outperform in the short term.

I will explain these methods in turn.

Genetic Programming for Stock Filtering

Genetic programming is a process whereby programs, or formulas, may be evolved through an evolutionary competition. Conventionally, programs take the form of a tree, with crossover and mutation following the standard form for tree-based evolutionary operators.

Within the realm of stock fitting, genetic program (GP) trees were evolved in order to combine the various fundamental financial statistics in a way that would create innovative and potentially insightful rules for stock filtering. The benefits of a GP individual over a conventional stock filter are threefold:

- GP filters can explore a much larger search space than humans can
- GP filters may test ‘non-humanlike’ filter terms

that a human would never think of, but which might provide a substantial improvement over the standard results.

- GP filters are not constrained by the human bias of knowing what to expect.

The drawbacks of using GP in a substantial search space are primarily concerned with over-fitting, and the vast explosion of search complexity for a modest increase in node counts.

The search complexity issue can be at least partially solved by using a parsimonious fitness function, which rewards filter trees which achieve a good result with the smallest size possible. The problem here lies in configuring the balance between these two quantities without overly penalising novelty in the evolutionary process. The more strictly one penalises tree complexity, the more rapidly one arrives at an answer, but the less likely one is to find a **good** answer.

Figure 3 shows an example GP stock filter, derived from a database of FTSE top 100 stocks over the period January 1990 – November 2004. This example is one of the more complicated ones, containing 39 nodes.

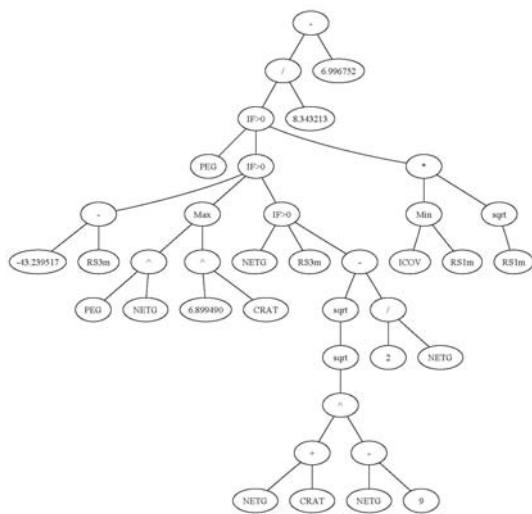


Figure 3 : An example GP stock filter

Evolutionary Conjunctive Rules Algorithm for Stock Filtering

Conjunctive rules are simply sets of binary statements joined together by an AND operator. That is to say, a data point passes the test if and only if it satisfies every single one of a series of rules.

The genome of an individual therefore consists of an open ended set of rules, each of which is specified in a numerical manner including the rule type and its associated values.

Minor mutations in a conjunctive rule set are far less destructive than mutations within a GP code, so

therefore this method is more robust to evolutionary change, and is also less susceptible to falling into local minima, due to its vastly lower search complexity.

Several extensions remain to be studied. For example, the conjunctive rules algorithm can be adapted to include a ‘mostly correct’ pass criterion, whereby a data point (individual stock) may fit all but one of the rules, and still be accepted.

Results

Work is still ongoing with this project, but the aim is to be able to test GP and ECR filters against those that are conventionally used in finance, such as the ‘Graham screen’, which looks for value stocks.

Example GP and ECR portfolios seem almost always to prefer ‘momentum investing’. That is, to pick stocks that already seem to be outperforming the index, and then to choose these with a certain set of provisos, usually based around debt levels and a ‘braking’ effect to avoid choosing those stocks which look as if they are growing too quickly.

4. SELECTED FURTHER PROJECTS

There isn’t space in one paper to cover all of the various research and development projects that we have undertaken in Cercia, so I include a selection of other projects here in brief detail. Further information can be provided on request, either from the author or from the lead researchers cited at the beginning of each section.

Many of these projects were carried out in collaboration with an industrial collaborator or sponsor. In some cases, this restricts the amount that can be disclosed about the projects themselves.

Adaptive Scheduling for Warehouse Management

In collaboration with a major Birmingham-based database & software company. Lead researcher : Yong Xu (yxx@cercia.ac.uk)

Warehouses are managed by extensive computerised databases, which record the current warehouse state and the current list of tasks to be done. Tasks may be stock moves, stock checks or stock manipulations (e.g. splitting or re-forming of component parts).

In a warehouse there are a fixed number of people, and a fixed number of vehicles. Each person has a set of abilities, allowing them to use, or prohibiting them from using each of the items of machinery. Similar permissions are also defined for access to physical locations within each warehouse (e.g. restricted secure areas, freezer storage), access of each vehicle to each area and permission of each person to be able to move

each product type. Finally, each vehicle has a certain capacity, maximum height, physical size and speed which affect the tasks that it can do efficiently.

The challenge is to generate a dynamic task allocation system which will distribute, on demand, the pending tasks in such a way as to minimise the number of tasks completed past their deadline and to maximise the throughput of tasks through the system with the fewest resources possible.

We created two components for this particular feasibility study. Firstly, a scheduler which took the current warehouse state, permission matrices and task list, and generated on demand a next best task for each of the available staff members. Secondly, a complete warehouse simulation environment in which our proposed scheduling algorithms would be tested.

We investigated many technologies on request from this company, in particular the application of neural networks to task duration estimation, and also the use of evolutionary algorithms to provide a powerful on-demand scheduling tool.

ChessBrain – The World's Largest Chess Supercomputer

Lead Researcher : Colin Frayn (cmf@cercia.ac.uk)

The ChessBrain project currently holds an official Guinness World Record for the largest number of computers used to play one single game of chess.[11][12][13]

ChessBrain was formed as a collaboration between Carlos Justiniano (CFC Inc., USA) and Colin Frayn (Cercia) in order to investigate the feasibility of massively-distributed inhomogeneous speed-critical computation over the internet. Distributed search gives huge speed benefits on parallel problems, though chess tree search is only partially parallel so this introduces a number of new challenges. In addition, we had to deal with the effects of varying CPU speed, O/S and connection bandwidth, all of which made the synchronization of search results extremely difficult.

In the two years since ChessBrain played its first match, we have been working on a second generation framework into which we can host the same chess-playing AI structure, but which will enable us to make far better use of that same AI and will permit efficient access to a far wider range of contributors, including locally networked machines and dedicated compute clusters.

Figure 4 shows the configuration of the next generation ChessBrain II system, which utilises a hierarchically distributed model in order to cut down the issue of connection requests clogging the central server. The intelligence sits in the middle, distributing new work units as and when they arise, and updating its current

memory state with any new data returned from the connected contributors.

By the time this paper is delivered, ChessBrain II will have played its first two demonstration matches in Los Angeles, California and Reno, Nevada. Both of these will be taking place in early May, if all goes according to plan.

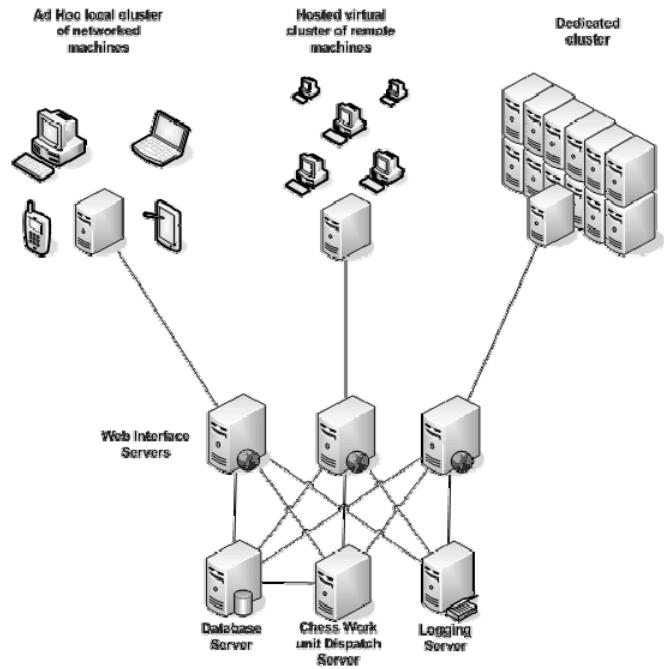


Figure 4 : ChessBrain II configuration

Energy Usage Monitoring Systems

In collaboration with the University of Birmingham Estates Office. Lead Researcher : Thorsten Schnier (txs@cercia.ac.uk)

The University of Birmingham has over 500 electricity meters, remotely read, throughout its campus. These readings arrive at a frequency of one every 30 minutes from each of the 500 meters, which makes monitoring them a very difficult task except in an extremely naive manner. Currently the best available solution is either to set a fixed maximum limit for each day, per meter, or else to specify a manually-configured warning profile. Our system takes historic data and learns past trends in order to warn when usage deviates significantly from the expected values for a given location, time of day and time of year.

Analysis of EEG Signals

In association with the Medical School, University of Birmingham. Lead Researcher : Xiaoli Li (xxl@cercia.ac.uk) [14][15][16][17]

This work covered several aspects of the analysis of EEG data using wavelet techniques. The four main areas of interest were:

- Prediction of Epileptic Seizure;
- Computation Neuronal Oscillations;
- Measure of Depth of Anaesthesia via EEG recordings;
- Multiple time series for Intracranial EEG Recordings;

In the first project, which has most real-world significance, wavelet coherence techniques were used to monitor the signals from EEG measurements in epileptic patients. Initial studies showed how wavelet techniques incorporating chaos theory could be used to predict an upcoming epileptic seizure significantly ahead of time. This is now being developed into a Matlab toolkit, and potentially into a commercially available hardware and software package for sale to medical centres (initially) across the UK.

Financial Planning and Strategy Development for the Deregulated Water Industry

In association with a large utility company within the UK. Lead Researcher : Jin Li (jxl@cercia.ac.uk)

This particular client was interested in exploring computational intelligence tools in order to build a comprehensive model of competition within a deregulated water industry. Deregulation in the UK is only just over the horizon, and is expected to be introduced in the near future.

In this environment, companies must be able to identify where they are most likely to make substantial gains in terms of customer numbers, and in which areas their competition will be fiercest. This forewarning will allow our clients to direct their advertising spend towards the places where it is most necessary, and also to configure their tariff layout accordingly.

Cercia developed a model which could be incorporated with our client's existing databases and prediction models, which would simulate the effect of various pricing strategies and assumptions about competitor policies and thus predict the expected market share on a geographical and temporal basis.

Exploiting Ensemble Diversity For Automatic Feature Extraction

In collaboration with Honda Research Laboratories, Germany. Researchers : Vineet Khare, Gavin Brown, Xin Yao, Jeremy Wyatt (Birmingham) [18][19]

This work is an effort towards developing a system that would automatically discover natural decompositions

of complex problems, while simultaneously solving the sub-problems. In most cases such decomposition relies on human experts and domain analysis. A system that can produce modules, which solve a subset of a big problem, can save us from manually crafting them. Other than the design and analysis of such a system, knowledge about how and when modularity is useful in automatic problem decomposition also constitutes an essential part of this work.

Multi-network systems, i.e., multiple neural network systems, can often solve complex problems more effectively than their monolithic counterparts. Modular neural networks tackle a complex problem by decomposing it into simpler sub-problems and then solving them. Unlike the decomposition in modular neural networks, a neural network ensemble usually includes redundant component nets and is often inspired by statistical theories. We have studied different types of problem decompositions and discuss the suitability of various multi-network systems for different decompositions. A classification of various multi-network systems, in the context of problem decomposition, was obtained by exploiting these differences. Then we investigated a specific type of problem decomposition, which gives no information about the sub-problems and is often ignored in literature. A novel modular neural network architecture for problem decomposition is presented. With the help of this modular network structure we assessed the usefulness of modularity in Artificial Neural Networks for the type of problem decomposition chosen.

Various factors, including learning algorithm (batch and incremental; first and second order gradient-descent and evolutionary learning algorithms), type of network (Radial Basis Function and Multi-Layer Perceptron) and type of task (linear and non-linear; static and dynamic) required of the network, were considered for this assessment. Based on this assessment we concluded that the usefulness of modularity, if judged on the basis of the performance of the system on a static task, depends on the learning algorithm used and by using a more sophisticated learning algorithm it is possible to achieve similar, if not better, performing non-modular structures as against a modular structure. However, we found modular structures at an advantage for a couple of dynamic-task-scenarios. These dynamic environments might have given modular structures the much debated selective advantage over monolithic systems during evolution.

A co-evolutionary model was created, which we used to design and optimize such modular neural networks with sub-task specific modules. The model consists of two populations. The first population consists of a pool of modules and the second population synthesizes complete systems by drawing elements from the pool of modules. Modules represent a part of the solution, which co-operate with each other to form a complete solution. Using two artificial supervised learning tasks,

constructed from smaller sub-tasks, we showed that if a particular task decomposition is better than others, in terms of performance on the overall task, it can be evolved using the co-evolutionary model. The co-evolutionary model is assessed on various tasks, previously found favourable to modular structures, to check for the emergence of corresponding modularity. Further, the co-evolutionary model is also used to support arguments presented as possible reasons for the abundance of modularity in natural complex systems.

Diversity creation in local search for the evolution of neural network ensembles

Researcher : Pete Duell (Birmingham)

One approach to the design of accurate and diverse ensembles is the EENCL algorithm (Evolutionary Ensembles with Negative Correlation Learning) [20]. Unlike many other ensemble methods, the individual networks are trained in parallel, rather than independently or sequentially. Individual networks learn by Negative Correlation Learning (NCL) [21] and evolutionary learning [22]. Diversity amongst the final population is encouraged by the negative correlation of the networks' outputs and through speciation by implicit fitness sharing [23][24].

Both accuracy and diversity are important for the creation of good ensembles. EENCL proved successful on some problems. However, little work has been carried out to analyze why EENCL is effective. Specifically we are interested in what contribution NCL makes to the performance of the algorithm, since it introduces an extra parameter and additional complexity. This paper uses additional datasets to analyze how the two learning mechanisms of global evolution and local search interact. Surprisingly, we find that NCL is not an essential component of EENCL for the datasets tested, and that a comparable performance can be achieved with a much simpler local search technique: Back propagation. Our experiments show that by replacing NCL in EENCL with Back propagation, we can achieve comparable classification accuracies, and also produce ensembles that are just as diverse in terms of the joint correct sets of the networks and also in terms of correlation of outputs.

MSc in Natural Computation

Lead by Xin Yao, Birmingham. The University of Birmingham School of Computer Science, in collaboration with several Cercia researchers, has been running the UK's first natural computation masters course for several years, with excellent results. Several of our graduates have gone on to work for major multinationals, study for MBA degrees and continue their studies at the PhD level and beyond [25].

Continuous Spatial Iterated Prisoner's Dilemma with Kin Markers

Lead researchers : Colin Frayn, Siang Yew Chong, Andy Pryke (Cercia).[26]

The iterated prisoner's dilemma (IPD) has been used as a model for investigating cooperation in nature. In this work, we presented an analysis of the evolution of reciprocal cooperation in a dynamically simulated environment in which individual agents were free to move in space, interacting with their nearest neighbours in fixed-length IPD games. [27]

Agents aim to avoid those against whom they score poorly, and to seek out those against whom they score highly. Individuals are modelled using finite state machines, allowing us to extend previous work on kin group markers. Though they had no direct effect on an individual's strategy, such markers did lead to the emergence of coherent, mutually-cooperating sub-populations.

Our conclusions, resulting from over 5.4 trillion individual IPD games, include the following:

- (1) We confirm the prediction made by McElreath [28] that individuals prefer to cooperate with, and remain near, their own kin compared with others.
- (2) Increasing state complexity or kin diversity increases the degree of inter-group cooperation and leads to less frequent extinctions.
- (3) Semi-stable predatory niches with mimicry do exist within this environment.
- (4) The introduction of errors increases the degree of inter-group cooperation.

A Memetic Algorithm for VLSI Floor Planning

Researchers : May Tang, Xin Yao (Birmingham)

For more information, please see [29]

Adaptive Nonlinear Dimensionality Reduction with a Local Metric

Lead Researcher : Huanhuan Chen (Birmingham)

With the development of data collection and storage capabilities, researchers working in domains as diverse as engineering, astronomy, biology, remote sensing and economics, have to face larger and larger data sets and higher and higher dimensional data. Traditional statistical methods break down partly because of the increase in the number of data points, but mostly because of the curse of dimensionality

Understanding the structure of multidimensional patterns, especially in the unsupervised case, is of fundamental importance in data mining, pattern recognition and machine learning. Several algorithms have been proposed to analyze the structure of high dimensional data based on the notion of manifold learning. These algorithms have been used to extract the intrinsic characteristics of different types of high dimensional data by performing nonlinear dimensionality reduction. Most of these algorithms rely on Euclidean metric and manually choosing neighborhood size. They can't recover the intrinsic geometry of data manifold automatically and accurately for some hand data sets.

In this work, we have developed an adaptive version of ISOMAP[30] that integrates the advantages of local and global manifold learning algorithms. Faster convergence rate of this algorithm is obtained by replacing the Euclidean metric with an arc length metric as the second-order approximation of geodesic distance. Our experiments on synthetic data as well as real world images demonstrate that our proposed algorithm can achieve better performance than ISOMAP and solve some harder data sets that are impossible for global methods. Paper submitted April 2006.

Environmental Monitoring and Protection

Lead researchers : Yong Liu, Xin Yao

Environmental monitoring and protection are crucial to the quality of our life and the ecological system in the world. For example, a major outbreak of the blue-green algae in a fresh water lake can have a devastating impact on the fish and other lives in the lake.

We have used evolutionary artificial neural networks to predict the activities of Chlorophyll-a in a lake in Japan [31]. We were able to produce much more accurate prediction than other methods. The more accurate prediction would enable us to take an appropriate preventive action and reduce the risk of a major outbreak.

Traffic Flow Prediction in Telecommunications Networks

Lead researcher : Yong Xu (yxx@cercia.ac.uk)

In a telecommunications network, it is very important to know when the network is busy and when it is not. Such information will enable a carrier (company) to make an informed decision on the necessary capacity between two cities as well as set a pricing policy that encourages off-peak use of certain lines. However, it is very difficult to predict traffic flow in a conventional telecommunications network. We have used our newly developed neural network ensembles to predict traffic flow in an Austrian telecommunications network among 32 regions [32]. We have shown that our negative correlation learning algorithm can train an

ensemble successfully to solve this problem.

Although the project was carried out for a telecommunications network, the techniques we used can equally be applied to flow prediction in an electricity network, a water network, or a gas network. We will be able to predict the demand within a period based on historical data.

Further work in this area can be found in [33][34].

Credit Card Assessment

Lead researcher : Xin Yao (xin@cercia.ac.uk)

Credit cards are widely used in the world, even in developing countries like China and India. However, issuing a credit card carries some risk because the card holder may or may not pay back the money. It is important for any bank to assess the risk of a potential card holder before issuing a card. We have applied our evolutionary artificial neural networks and neural network ensembles for credit card application assessment in an Australian bank. Excellent results have been obtained in comparison with other existing methods.

The techniques we have developed for the Australian credit card problem can be applied to other problems, such as insurance fraud detection, risk assessment for loans, premium setting for insurance, etc.

Clustering and its Application to Rail Maintenance

Researchers : Derek Bartram, Xin Yao (Computer Science), Michael Burrow (Engineering)

Railway track intervention planning is the process of planning what maintenance to perform, where and when. Until approximately the early 1990's the process was performed solely by the track engineers using their judgment. While the track engineers are typically very experienced, there was a concern that the maintenance planning was not consistent, and certainly not optimal.

For this reason, the European Rail Research Industry (ERRI) commissioned EcoTrack; a system for producing intervention plans. The complexity of railway track intervention planning is down to three main factors; firstly railway track is comprised of a large number of parts (with many variations) with highly complex interactions secondly, the deterioration of each of the various components is unique and again complex, and thirdly, track (including its components) can fail in a number of ways due to many different factors, many of which are external.

Producing a single uniform model for rail deterioration is a highly complex process, however by adopting a divide and conquer approach (using clustering as the process for dividing the complexity of the problem),

the process of producing a model can be simplified. Instead of producing a single highly complex model for rail deterioration, for each type of failure a less complex model can instead be produced.

As part of the clustering process, rule extraction is used for the determination of which cluster a new data element belongs to. Two main techniques are used, decision tree generation, and evolutionary algorithms. The rules generated can be also be used to verify the correctness of the clusters, especially in cases where the concept of closeness is less clear, as in the case of rail deterioration systems.

Paper currently in preparation.

5. CONCLUSION

Computational intelligence technologies are now being used in industry to an increasing degree. Technology transfer centres like Cercia are helping by developing new CI technologies, and also by tailoring their research and development around what industry wants and needs.

With increasing levels of data and complex numerical processes within industry, it is more important now than ever before that we encourage the adoption of 21st century techniques in areas where 20th century techniques will surely fail.

The greatest benefit of natural computation techniques in industry is that they are so easy to sell – with businesses interested in the bottom line, all it takes is a well-crafted demonstration which shows the percentage saving, either in time or money, that a natural computation technology can bring. Especially in fields like journey optimisation, packing, scheduling or stock cutting, it is particularly easy to convert savings in time, distance or space directly into money.

Natural computation technologies are no longer the exclusive realm of academic research, and the best way to secure future years of funding is to publicise the technologies not just as academic toys, but also as the revolutionary, powerful, flexible technologies that they are. By pushing them towards commercial prominence we can guarantee their continued support within academic circles.

The easiest way to expand the knowledge and use of NC technologies within industry is simply to expand the education of these technologies at University undergraduate level. Our highly successful Natural Computation MSc programme has provided a large number of highly intelligent NC graduates to industry and academia, and now several other UK Universities are following suit.

It is at these grass roots levels that we must focus our effort in encouraging the uptake of NC technologies. Our experience with dealing with a top-down approach has been largely fruitless except within those large companies who already possess a

substantial research department. Attempting to introduce cutting-edge CI techniques has met a huge amount of resistance at all levels within industry, primarily on grounds of ignorance, fear and resistance to change. In order for NC techniques to become indispensable to industry, their use must be considered right from the beginning of any project, and incorporated into the fabric of good coding practice.

Our experience has shown that NC technologies can provide a very significant saving in many areas for companies of all sizes. Their flexibility, robustness and simplicity are all highly valuable selling points, and they are the exact reasons why industry is finally beginning to sit up and listen.

6. ACKNOWLEDGMENTS

CMF acknowledges financial support and assistance from Advantage West Midlands, UK. Also from the companies and colleagues named in this report. Some of these projects were carried out with financial support from VIN Technology services, University of Birmingham, UK.

7. REFERENCES

- [1] Back, T., Hammel, U., Schwefel, H-P., “Evolutionary computation: comments on the history and current state”, (1997), IEEE Trans. Evol. Comp., Vol 1., No. 1, pp3-17
- [2] Applegate, D., Bixby, R.E., Chvatal, V., Cook, W., “On the solution of travelling salesman problems” (1998) Documenta mathematica – Extra Volume, ICM III pp645-658
- [3] Applegate, D., Bixby, R.E., Chvatal, V., Cook, W., “Finding cuts in the TSP” (1994), Math. Prog. Symp., Ann Arbor, Michigan.
- [4] Golden, B., Wong, R.T., “Capacitated Arc Routing Problem (1981) Networks, Vol. 11, pp305-315
- [5] Lacomme, P., Prins, C., Ramdanicherif, W. (2004) “Competitive Memetic Algorithms for Arc Routing Problems,” Annals of Operations Research, Vol. 131, pp. 159–185
- [6] Handa, H., Chapman, L., Yao, X., “Robust route optimisation for gritting/salting trucks: A CERCIA experience”, (2006) IEEE Computational Intelligence Magazine, Vol. 1, No. 1, pp6-9.
- [7] Handa, H., Chapman, L., Yao, X., “Dynamic Salting Route Optimisation using Evolutionary Computation”, Proc. CEC (2005).
- [8] Nagata, Y., Kobayashi, S. Edge Assembly Crossover: A High-power Genetic Algorithm for the Traveling Salesman Problem,” (1997) Proc. 7th Int'l Conf. on Genetic Algorithms, pp. 450–457
- [9] Chapman, L., Thornes, J.E., Bradley, A.V. (2001) “Modelling of road surface temperature from a geographical parameter database. Part 1: Statistical,” Meteorological Applications, Vol. 8, pp. 409–419
- [10] Chapman, L., Thornes, J.E., Bradley, A.V. (2001) “Modelling of road surface temperature from a

- geographical parameter database. Part 2: Numerical,” Meteorological Applications, Vol. 8, pp. 421–436
- [11] Frayn, C.M. & Justiniano, C., “The ChessBrain Project – Massively Distributed Inhomogeneous Speed-Critical Computation”, Proceedings IC-SEC, Singapore, 2004
- [12] Justiniano, C. & Frayn, C.M. “The ChessBrain Project: A Global Effort To Build The World's Largest Chess SuperComputer”, 2003 ICGA Journal, Vol. 26, No. 2, 132-138
- [13] Justiniano, C. “ChessBrain: A Linux-Based Distributed Computing Experiment”, 2003 Linux Journal, September 2003
- [14] Li, X., Polyzgiannakis, J., Kapiris, P., Peratzakis, A., Eftaxias, K., Yao, X. “Fractal spectral analysis of pre-epileptic seizures in terms of criticality”, (2005) Journ. Neural Engineering, Vol. 2, No. 2 pp11-16.
- [15] Li, X., Ouyang, G., Yao, X., Guan, X., “Dynamical Characteristics of Pre-epileptic Seizures in Rats with Recurrence Quantification Analysis”, (2004) Phys. Lett. A., Vol. 333, No. 1-2, pp164-171.
- [16] Kapiris, P.G., Polyzgiannakis, J., Li, X., Yao, X., Eftaxias, K.A., “Similarities in precursory features in seismic shocks and epileptic seizures”, (2005) Europhys. Lett. Vol 60., No. 4 pp657-663
- [17] Li, X., Ouyang, G., Yao, X., Guan, X., “Dynamical characteristics of pre-epileptic seizures in rats with recurrence quantification analysis” (2004) Phy. Lett. A., 333 pp164-171
- [18] Brown, G., Yao, X., Wyatt, J., Wersing, H., Sendhoff, B., “Exploiting Ensemble Diversity For Automatic Feature Extraction” (2002) Proc. ICONIP, pp1786-1790.
- [19] Khare, V., Yao, X., Sendhoff, B., “Multi-network evolutionary systems and automatic problem decomposition” (2005) IJGS (accepted)
- [20] Liu, Y., Yao, X., Higuchi, T., “Evolutionary ensembles with negative correlation learning”, (2000) IEEE Trans. EC Vol. 4 No. 4 pp380-387.
- [21] Liu, Y., “Negative correlation learning and evolutionary neural network ensembles.” (1998) PhD thesis, UNSW.
- [22] Yao, X., “Evolving artificial neural networks”, (1999) Proc. IEEE, Vol. 87 No. 9.
- [23] Holland, J.H., “Adaptation in natural artificial systems”, (1975) Uni. Mich. Press.
- [24] Goldberg, D.E., Richardson, J., “Genetic algorithms with sharing for multimodal function optimisation.” (1987), Proc. ICGA ed. Grefenette pp41-49.
- [25] Yao, X., “A Research-Led and Industry-Oriented MSc Program in Natural Computation”, (2006) IEEE Computational Intelligence Magazine, Vol. 1, No. 1, pp39-40.
- [26] Chong, S.Y., Yao, X., “Behavioural Diversity, Choices, and Noise in the Iterated Prisoner's Dilemma” (2005) IEEE Trans. EC, Vol. 9, No. 6 pp540-551.
- [27] Axelrod, R., “The Evolution of Cooperation” (1984), Basic Books, New York
- [28] McElreath, R., Boyd, R., Richerson, P.J.: “Shared Norms and the Evolution of Ethnic Markers”, Current Anthropology, Vol. 44 (2003) 122-130
- [29] Tang, M., Yao, X., “A Memetic Algorithm for VLSI Floor planning”, 92005) IEEE Trans. Syst. Man. Cyber. Part C (accepted)
- [30] Tenenbaum, J.B., de Silva, V., Langford, J.C., “A global geometric framework for nonlinear dimensionality reduction”, (2000) Science, 290 pp2319-2323
- [31] Liu, Y., Yao, X., “Evolving Neural Networks for Chlorophyll-a Prediction” (2001) Proc. ICCIMA pp185-189, IEEE CS Press
- [32] Yao, X., Fischer, M.M., Brown, G., “Neural network ensembles and their application to traffic flow prediction in telecommunications networks” (2001) Proc. IJCNN pp693-698. IEEE Press
- [33] Xu, Y., Salcedo-Sanz, S., Yao, X., “Metaheuristic Approaches to Traffic Grooming in WDM Optical Networks”, (2004) PPSN workshop
- [34] Au, W.-H., Chan, K. C. C., Yao, X., “Data Mining by Evolutionary Learning for Robust Churn Prediction in the Telecommunications Industry,” IEEE Transactions on Evolutionary Computation, Vol. 7, No. 6 pp532-545