AISB Quarterly

No.124

The Newsletter of the Society for the Study of Artificial Intelligence and Simulation of Behaviour

On Body Posture as an Important Modality for Recognizing Affect

With technology becoming ubiquitous in our daily life, new application fields have emerged that are calling for renewed efforts from our community in improving the quality of the interaction between system and user. One area needing such improvement is the recognition, and exploitation, of the affective state of the user. By incorporating affective states in the loop, a more natural, and therefore more effective, interaction can be realized so that technology can augment, rather than impair, the sensitivity and cognitive abilities of the user.

In this respect, our research falls squarely in the scope of affective computing, a recent discipline defined as computing that relates to, arises from, or deliberately influences emotions.[1] However, whilst affective computing has been mostly about detecting affective states from facial expressions, vocal features, and bio-feedback, our research focuses on posture.

To address the lack of a commonly accepted set of descriptors such as Ekman and Friesen's Facial Action Coding System for example, we proposed a general description of posture based on the kinematics of the human and additional components, such as movement tendency based on Laban's work on dance. We used this general description to create an affective posture recognition system that mapped the set of postural descriptors into affective categories using a revised version of an associative neural network called CALM (Categorizing and Learning Module).[2] When kinematical only features were considered, the system yielded 71% of accuracy in recognizing four basic affective states (happiness, sadness, anger and fear) in 102 postures extracted from natural human motion capture data (see Figure 1 for a few samples). Adding a measure of the direction of the movement to the postural descriptor allowed for a significant improvement in the recognition rate (+8%), especially, by removing confusion between fearful postures and angry or happy postures.[3]

While these results were encouraging, they did not tell us anything about the saliency of the set of proposed features. In a series of follow-up studies, we tested the informational content of the posture descriptors. In reference 4 and related papers, we used discriminant analysis to build predictive models and measure the saliency of the set of features.

Nadia Bianchi-Berthouze
University College London
Continued on p. 3

Figure 1. Anthropomorphic avatars showing a range of emotions.

In this issue:

AI, the Olympics and the Force

The **AISB Quarterly** is published by the Society for the Study of Artificial Intelligence and Simulation of Behaviour.

ISSN 0268-4179

Editor
Colin G. Johnson.
Computing Laboratory,
University of Kent at Canterbury.
aisbq@aisb.org.uk
http://www.aisb.org.uk and click on AISBQ.

Editorial Assistant
Asim Rehman Shaikh

Page Layout
Dylan Banks

Advertising and Administration
Therie Hendrey-Seabrook
aisbq@aisb.org.uk
School of Science and Technology
University of Sussex
Falmer, Brighton
UK, BN1 9QH
Tel: (01273) 678448
Fax: (01273) 671320

It occurred to me a while ago that the presence of the Olympics in Britain in 2012 may provide some interesting opportunities to the AI community in the UK, and there may be ways in which the Olympics could be "leveraged" to help a number of ends. These include raising public understanding of AI, attracting more school-leavers and existing tertiary students into AI, inspiring funders and companies to provide more support for AI ... and even, who knows, inspiring new research in AI.

I'd like to get the Society at large thinking about this. Your Committee has been discussing some possibilities in brain-storming mode, and in this we have involved some of the Fellows of the Society. More pointedly, the Committee has decided that the best way forward is to seek to establish a Task Force charged with further exploration of possibilities, formation of proposals, lobbying the Olympian gods, etc. The Committee itself is not in a position to take on such a concerted, major task, and in case we want to try to engage on the Force the people who are most energetic, motivated and relevantly knowledgeable. The Committee would of course play a background supportive role. So, we are looking for volunteers for this Task Force.

To stimulate thinking, here are some possible ways in which AI could be related to the Olympics. There are no doubt others.

1) AI research helping sports science (e.g., via intelligent emulators and body sensors) in the run-up to the Olympics, and of course beyond.
2) AI involved in the conduct of human sport itself (e.g., intelligent video monitoring of games).
3) AI in virtual-reality sport for people.
4) AI competitions to be held in association with the Olympics: from robotic football through intelligent trading-agents competition to intelligent computer gaming.
5) AI research helping with the development of infrastructure for the Olympics: this could involve navigation systems, intelligent transport systems, intelligent buildings, security, automatic summarisation of news, real time translation, legal and financial reasoning, etc. etc.

There are special opportunities for getting children involved, and also for getting physically-disabled and learning-disabled people engaged.

By the way, I speak as someone who doesn't follow sport (apart from celebrity come-ice-skating).

So ... over to you. Contact any member of the Committee if even only dimly interested in being on the Task Force. Or, if you don't want to be that Forceful we'd be very interested in any ideas or comments you might have.

John A. Barnden
University of Birmingham
J.A.Barnden@cs.bham.ac.uk

A Welcome from the Editor

Welcome to the latest issue of the AISB Quarterly, my first as editor. I would like to thank my predecessor, Sunny Bains, and my new editorial assistant, Asim Rehman, for their help with this issue.

This issue contains articles on a number of exciting topics in AI — Nadia Bianchi-Berthouze on computer-based recognition of emotion through bodily posture, Mark Jago on resource-bounded agents, two articles on swarm intelligence, and my own article on genetic programming. In addition there is a report on the 2006 GECCO conference by Qijun Zhang, and AISB president John Barnden looks at the prospects for promoting AI and computing in the context of the 2012 London Olympics.

We are always interested in ideas for new articles — please look on the AISB website for details of how to submit ideas. In addition there are a number of books waiting to be reviewed — again, a list of these is on the web site, please contact me if you would be interested in writing a review of one of these for a subsequent issue.

I hope that you enjoy this issue.

Colin Johnson,
editor, AISB Quarterly
aisbq@aisb.org.uk

On Body Posture as an Important Modality for Recognizing Affect

continued from p. 1

These studies revealed that the models generally outperformed human observers in discriminating between 8 affective categories. More interestingly, they suggested an explanation to recent findings in neuroscience suggesting that the face fusiform area — the brain area responsible for facial processing — was involved in processing postural expressions of affect even when facial cues were removed.[5] Indeed, our statistical analysis showed that features related to head configuration were very important in discriminating between emotions and in particular between nuances of a particular emotion.

In reference 6 we further investigated the informational content of the features by looking at whether the features could account for different levels of three affective dimensions: arousal, valence and action tendency. The results (1% error on arousal, 20% on valence and 25% on action tendency) compared well with reported data on the recognition of affect from other modalities such as facial expression for example. This led us to investigate whether it would also be possible to quantify the influence of culture on postural affective displays. Using the same basic representation, we conducted recognition experiments with subjects from three cultures that covered the spectrum of Hofstede's cultural dimensions, and related the subjects' judgments and their differences to the postural features. Separate cultural models of affective posture recognition were built using Mixture Discriminant Analysis and Expectation Maximization, and their evaluation showed significant recognition rates.[7]

This body of work thus suggests that posture could be used, if not as an alternative to facial expressions, at least in conjunction with facial expressions to provide for finer grain appraisals and increased discriminatory power in the case of ambiguous or incongruent information. By providing the basis for culture-specific models, the studies open the door for the development of more believable animated agents. Finally, posture can also help to quantify the user experience itself. We are currently looking at the relationship between posture and engagement, immersion and even addiction.[8]

Nadia Bianchi-Berthouze
University College London
Interaction Centre
n.berthouze@ucl.ac.uk

References

1. R. Picard, Affective Computing, Cambridge: MIT Press, 1997.
2. N. Bianchi-Berthouze and A. Kleinsmith, A categorical approach to affective gesture recognition, Connection Science 15 (4), pp. 259-269, 2003.
3. A. Kleinsmith, T. Fushimi, T. and N. Bianchi-Berthouze, An incremental and interactive affective posture recognition system. International Workshop on Adapting the Interaction Style to Affective Factors, Edinburgh, UK, 2005.
4. R. de Silva and N. Bianchi-Berthouze, Modeling human affective postures: An information theoretic characterization of posture features, Journal of Computational Agents and Virtual Worlds 15 (3-4), pp. 269-276, 2004.
5. B. de Gelder et al., Fear fosters fight: A mechanism for fear contagion when perceiving emotion expressed by a whole body, PNAS 101 (47), pp. 16701-16707, 2003.
6. A. Kleinsmith, R. de Silva and N. Bianchi-Berthouze, Grounding affective dimensions into posture descriptions, International Conference on Affective Computing and Intelligent Interaction, pp. 263-270, Beijing, China, 2005.
7. A. Kleinsmith, R. de Silva and N. Bianchi-Berthouze, Cross-cultural differences in recognizing affect from body posture, Interacting with Computers (in press).
8. N. Bianchi-Berthouze et al., On posture as a modality for expressing and recognizing emotions, Workshop on the role of emotion in human-maching interaction, London, UK (in press).

The AISB Committee is delighted to announce that Austin Tate (University of Edinburgh) and Nick Jennings (University of Southampton) have accepted appointment as Fellows of the Society.

Nature-Inspired Approaches to Clustering

The identification of intrinsic classes in nominal or numeric data (often referred to as 'unsupervised classification') has been an important technique in biology and sociology for many years and these days plays a central role in the broad areas of knowledge discovery and data mining (KDD). Data clustering — the basis of unsupervised classification — consists of finding a 'natural' grouping in a data set, given some measure of similarity between each pair of data items.[1] One general, formal definition of the problem asks for a partition of the data into disjoint sets such that some measure of cluster homogeneity is optimized. This problem is NP-hard, however the prevailing difficulty of the clustering problem, in practice, is due in at least equal degree to the difficulty of deciding on an appropriate homogeneity measure to use (for several exist) in any given context. Different choices will lead to different results and *a posteriori* a good choice is one that leads to the identification of clusters (in arbitrary dimensions) that are analogous to what humans easily perceive, in two or three dimensions, as densely connected 'patches' or 'clouds' in the data space.

Recently, a number of novel approaches to clustering, drawing their inspiration from nature, have been investigated. A first category of such methods aims to model directly clustering processes that are observed in nature and attempts to modify these models to obtain algorithms suitable for the clustering of numeric data. Such algorithms comprise those imitating the flocking behaviour of swarms in nature,[2] as well as algorithms modelling the clustering and sorting activities observed in social insects (in particular ants, see Figure 1.[3,4]). All of these methods work through complex interactions between simple agents and their environment, and without an explicit formulation of the objective of the clustering task; this means that the difficult issue of choosing the clustering objective is side-stepped to an extent, and, moreover, the self-organizing operation of these algorithms leads to interesting phenomena, such as determining the number of clusters in the data, implicitly, as well as good scalability properties. However, these approaches also suffer from some drawbacks: in particular, the properties of these algorithms (essentially, the clustering objective that is implicitly optimized by these methods)

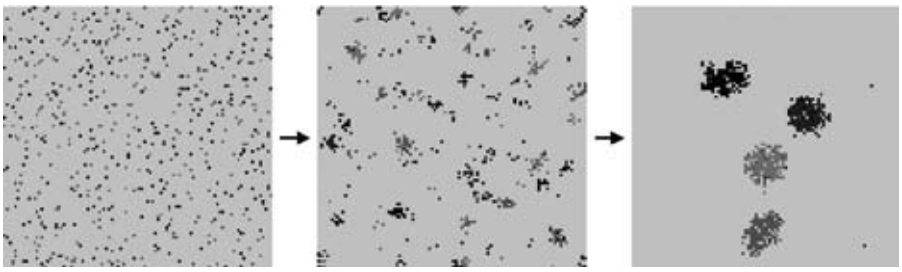

Figure 1. In ant-based clustering, the data items are initially distributed on a two-dimensional toroidal grid at random. The ants move on this grid and perform picking and dropping operations, which are biased by the density and similarity of the data located in the direct vicinity of a data item. In this way, a clustering of the data is obtained. In this example, a simple four-cluster data set has been correctly partitioned into four clusters. NB: the 2d grid used in this algorithm should not be confused with the data-space, which may be of arbitrary dimension.

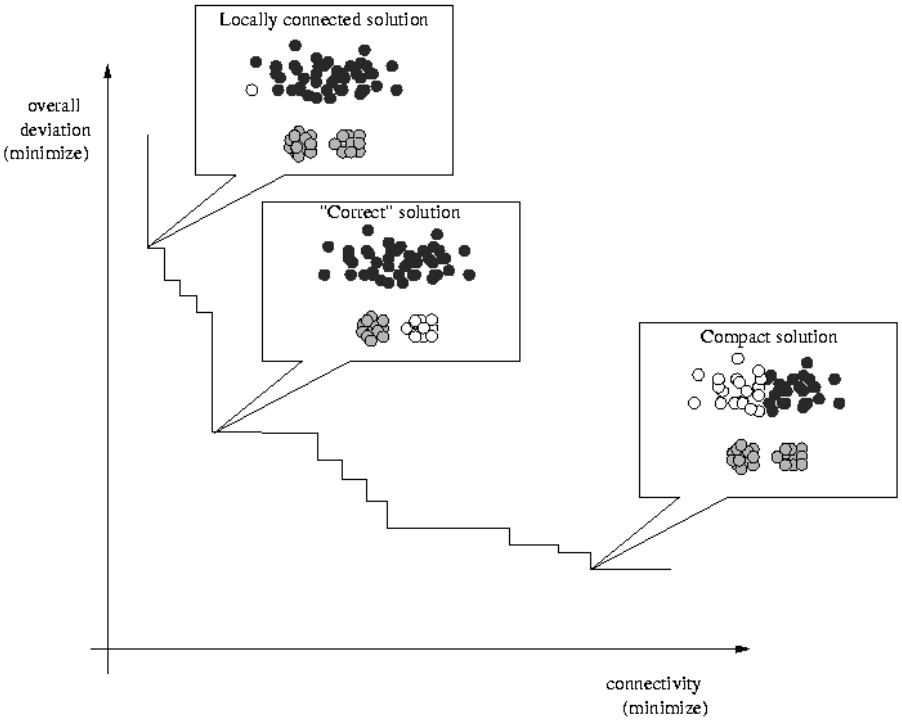

Figure 2. Illustration of the principal idea behind multiobjective clustering. Here, two different objectives are optimized, one of which (overall deviation) assesses cluster compactness, and the other one of which (connectivity) assesses the correct identification of neighbourhood relationships. Single-objective algorithms like k-means and single link hierarchical clustering will return clustering solutions that are optimal with respect to one of the two objectives (shown at the bottom right and the top left of the Pareto front). However, for the data set shown, the clustering solution that is the most intuitive to a human observer corresponds to a trade-off between the two objectives and can only be identified by a multiobjective approach.

Julia Handl and Joshua Knowles
University of Manchester
Continued on p. 7

Modelling Resource-Bounded Agents

Agents are individuals or systems that are capable of deliberating and acting autonomously. All agents, whether they are people, robots or software systems, are subject to limited cognitive resources. Agents only have so much memory and computational power available to them and they often have set periods of time in which they must complete their tasks. These are consumable resources that an agent needs to allocate to particular reasoning tasks in order to behave intelligently. The bounds on an agent's resources can affect how effectively the agent reasons and the conclusions that the agent can draw, which in turn affects how the agent behaves. The agent's reasoning abilities, such as the rules it uses to form plans or infer new conclusions, can also be thought of as resources that affect the agent's performance.

Agent-based technology is playing an increasingly important role in artificial intelligence. As with any relatively new technology, tools are required to verify that agent-based programs perform as they should. Logical models of agents have proved to be extremely powerful tools in this domain but make strong assumptions about the reasoning powers that agents have. In general, these approaches cannot factor in an agent's resource bounds, instead treating agents as ideal reasoners. In order to provide tools to model cases in which an agent's resources are crucial, different approaches are required.

An approach that is well suited to modelling artificial agents takes Dennett's intentional stance, treating agents as rational systems that have beliefs, desires and intentions.[3] A logical model can be built by considering the rules that the agent uses to come about new beliefs from old. An early approach along these lines deductively closes the agent's set of beliefs under these rules but this ignores the agent's time and memory bounds.[7] Our approach is to consider how the agent's set of beliefs can evolve over time, given the rules it has with which to reason. Each time the agent applies a rule to derive a new belief, the system moves into a new state.

The corresponding logical model represents every possible such state, each related to the states that could be reached from it by applying one of the agent's rules. The resulting relational structure can be described by a simple modal logic. Modal logics allow us to easily extend the model to incorporate multiple agents that communicate with one another. In this framework, it is easy to model an agent's time bounds by assuming that each application of one of the agent's rules requires a unit of time. It is possible to represent an agent's rules such that this assumption produces accurate results.

We then model an agent with a fixed time bound by restricting paths in the model to an appropriate length. There are a number of interesting logical results when particular classes of agents are considered.[4] The logic is decidable and has a complete axiomatization.[2] An interesting open problem is how best to extend the approach to deal with non-monotonic reasoners. Modelling memory bounds is done by thinking of each state of the agent as containing a fixed number of slots to which beliefs can be assigned. When all slots are full, the agent cannot add to its beliefs without first forgetting or deleting old ones.[1] This approach allows properties of an agent program to be verified using efficient model checking techniques. This provides an accurate way of modelling agents at a high level, without abstracting away from the crucial aspect of the agent's resource bounds.

This logic can be extended to a model of the information available to an agent.[5] Unlike traditional theories of information, this approach takes an agent's resource bounds into consideration. An interesting area of future work involves extending this logic to incorporate knowledge and justification as well as belief and information.

Mark Jago
University of Nottingham
mark.jago@nottingham.ac.uk

References

1. N. Alechina, P. Bertoli, C. Ghidini, M. Jago, B. Logan and L. Serafini, Model checking space and time requirements for resource-bounded agents, Proceedings of the Fourth Workshop on Model Checking and Artificial Intelligence, 2006.
2. N. Alechina, M. Jago, and B. Logan, Modal logics for communicating rule-based agents, Proceedings of ECAI 06, 2006.
3. Daniel C. Dennett, The Intentional Stance, MIT Press, 1987.
4. Mark Jago, Rule-based and resource-bounded: a new look at epistemic logic, Proceedings of the Logics for Resource-Bounded Agents workshop, 2006.
5. Mark Jago, Imagine the possibilities: Information without overload, Logique et Analyse (49), December 2006.
6. Mark Jago, Logics for Resource-Bounded Agents. PhD thesis, University of Nottingham, 2006.
7. K. Konolige, A Deduction Model of Belief, Morgan Kaufman, 1986.

On the Performance Analysis of Particle Swarm Optimisers

Approximately one decade ago a new stochastic optimisation algorithm, called Particle Swarm Optimisation (PSO), was proposed by Kennedy and Eberhart.[1] Since then, PSO has captured the interest of a growing research community and, as a consequence, several different variants are now available.

Over the years, the research on PSO has focused on different aspects including the introduction of algorithmic variants that improve over the initial PSO algorithms, applications of PSO for solving practical optimisation problems, the study of PSO algorithms for the solution of dynamic or multi-objective problems, and so on. However, at this very moment, it is unclear which of the many variants of PSO algorithms should be chosen for tackling a specific problem. In fact, the field lacks a clear definition of the set of algorithmic variants that could be considered the state-of-the-art. This can be considered an obstacle for the advancement of the field. For this reason and as a starting point for our work, we embarked on the task of comparing some of the most influential PSO variants in order to identify the state-of-the-art in the field. We have done this considering the application scenarios we might encounter in the real world.

When studying stochastic optimisation algorithms, one should always have in mind that, ultimately, they will be used to solve practical problems. Therefore, one of our tasks as researchers is to design algorithms capable of finding high quality solutions under the conditions that are likely to be found in real-world scenarios. Of course, it is always desirable to find high quality solutions in a timely fashion and without wasting computing power; however, because of the stochastic nature of these algorithms, this only happens with a certain probability. Formally, a stochastic optimisation algorithm A applied to a problem instance $\square$, will find a solution of quality q in time t only with a certain probability $P_{A,\square}(q,t) = P(RT_{A,\square} \leq t, SQ_{A,\square} \leq q)$, where $RT_{A,\square}$ is a random variable that represents run-time and $SQ_{A,\square}$ is another random variable that represents the solution quality achieved. $(RT_{A,\square},SQ_{A,\square})$ is a bivariate random quantity describing the run-time ad solution quality behaviour of an algorithm A when applied to problem instance $\square$; the probability distribution of this random variable is also known as the run-time distribution of A on problem $\square$.[2] An example of a (bi-variate) run-time distribution

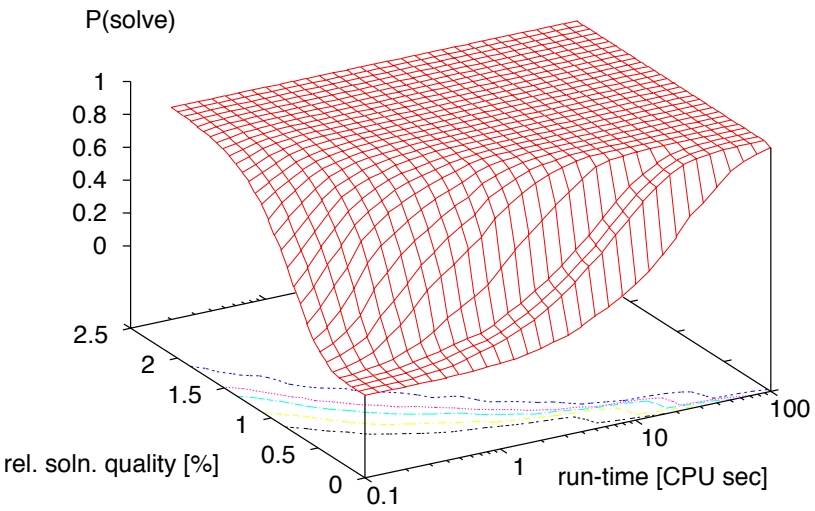

Figure 1. A typical run-time distribution of a stochastic optimisation algorithm applied to a particular problem. The independent variables, time and solution quality are on the x and y-axes respectively. The z-axis shows the empirical probability of finding a solution of quality q at time t.

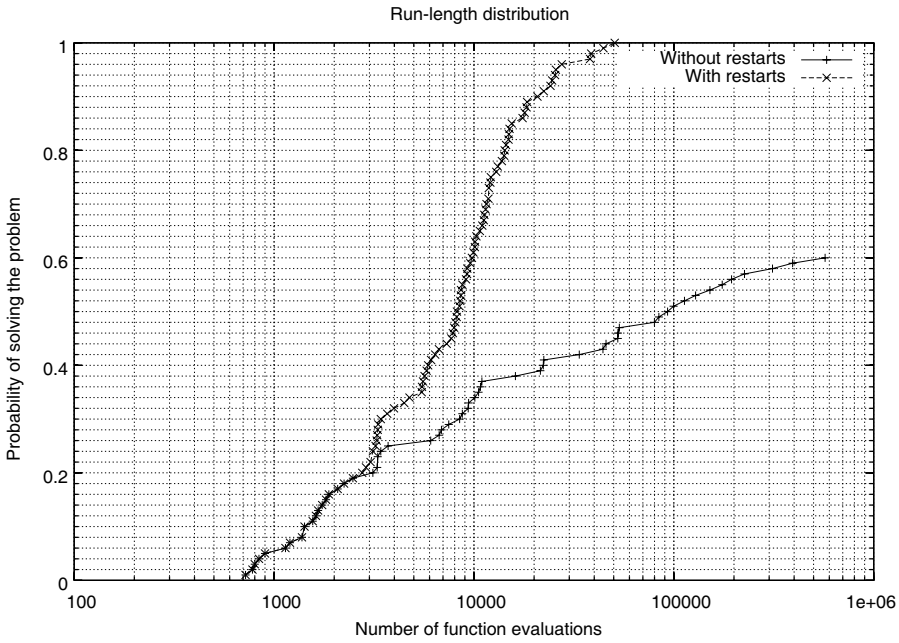

Figure 2. Qualified run-length distributions of one of the compared variants with and without dynamic restart mechanisms. With a restarting mechanism, the probability of finding a solution of a given quality is forced to approach

Marco A. Montes de Oca
Université Libre de Bruxelles
Continued on p. 7

On the Performance Analysis of Particle Swarm Optimisers: *continued from p. 6*

can be seen in Figure 1. Other commonly used performance statistics such as the average solution quality after a certain number of function evaluations, or the so-called "success rate", that is the ratio between the number of runs finding a solution of a certain quality and the total number of runs, only report the behaviour of the algorithm in a particular moment in time. The run-time distribution completely characterises the behaviour of algorithm *A* on problem □ since the probability of finding a solution of a certain quality at any point in time is given. We can thus draw conclusions about the behaviour of the studied algorithm during the optimisation process and not just at certain points in time. In continuous optimisation, it is common practice to measure run-times in terms of objective function evaluations. In this case, we talk of run-length distributions rather than run-time distributions.

We have recently compared seven PSO variants using run-length distributions on some commonly used benchmark functions.[3] In this study, we focused on measuring qualified run-length distributions which are cross-sections along the computing time axis of a full run-length distribution. The information contained in such qualified run-length distributions reveals some strengths and weaknesses of the algorithms that cannot be seen using other statistical measures. In particular, it is evident that, for some problems, a dynamic restarting mechanism can greatly improve the performance of the compared variants. Figure 2 shows an example of the kind of improvement that can be achieved by just adding a simple dynamic restarting mechanism to a PSO algorithm.

In the near future, we plan to exploit the information provided by run-length distributions and other analysis tools to propose improvements to the most promising PSO variants. Our research is focused towards the engineering of effective swarm-based optimisation algorithms.

Marco A. Montes de Oca *et al.*
Institut de Recherches Interdisciplinaires et de Développements en Intelligence Artificielle
Université Libre de Bruxelles
mmontes@ulb.ac.be

References:

1. J. Kennedy and R. Eberhart, Particle Swarm Optimization, In **Proceedings of IEEE International Conference on Neural Networks**, pages 1942-1948. 1995. IEEE Press.
2. H. H. Hoos and T. Stützle, **Stochastic Local Search. Foundations and Applications**. 2005. Morgan Kaufmann.
3. M. A. Montes de Oca, T. Stützle, M. Birattari and M. Dorigo, A Comparison of Particle Swarm Optimization Algorithms Based on Run-Length Distributions. In Proceedings of the 5th international Workshop on Ant Colony Optimization and Swarm Intelligence. **ANTS 2006**. **LNCS 4150.** pages 1-12.

Nature-Inspired Approaches to Clustering: *continued from p. 4*

are often poorly understood. An assessment of the suitability of these algorithms for the clustering of a particular type of data can therefore be very challenging.

A second category of nature-inspired methods for clustering comprises algorithms for general-purpose optimization, such as ant colony optimization, particle swarm optimization and evolutionary computation, which are adapted specifically to clustering by the definition of an appropriate representation, objective function(s) and operators.[5] The advantage of these methods compared to more traditional, greedy clustering algorithms, like k-means, is their ability to search globally for optima; another advantage is the explicit and open choice of objective function, which allows one to better understand and predict the performance of the clustering algorithm on particular types of data and, ultimately, helps in the choice and design of these algorithms. Furthermore, the explicit single-objective optimization may be extended to a multiobjective one, that is to say, several explicitly defined clustering criteria may be optimized simultaneously. This is an approach that has been explored only recently, but for which very promising results have been obtained.[6] In particular, it has been shown that good clustering solutions can correspond to a trade-off between two or more clustering criteria, and that, therefore, these solutions are accessible only to a multiobjective optimization approach; thus, multiobjective clustering which can find either extrema of a single objective, or tradeoffs of several, enjoys a general performance advantage.

Julia Handl and Joshua Knowles
University of Manchester
Email: j.handl@postgrad.manchester.ac.uk

References:

1. P. Arabie, L. J. Hubert and G. De Soete, Clustering and Classification, **World Scientific**, 1996.
2. F. Picarougne, H. Azzag, G. Venturini, C. Guinot, A new approach of data clustering using a flock of agents, **Evolutionary Computation**, 2006 (to appear).
3. E. Lumer and B. Faieta, Diversity and adaptation in populations of clustering ants, From Animals to Animats 3, pages 501-508, MIT Press, 1994.
4. J. Handl,J. Knowles and M. Dorigo, Ant based clustering and topographic mapping, **Artificial Life 12**, pp. 35-61, 2006.
5. V. J. Rayward-Smith, Metaheuristics for Clustering in KDD, Proceedings of the IEEE Congress on Evolutionary Computing, pages 2380-2387, **IEEE Press**, 2005.
6. J. Handl and J. Knowles, An evolutionary approach to multiobjective clustering, **IEEE Transactions on Evolutionary Computation**, 2006 (to appear).

Program Analysis in Genetic Programming

In TV programs such as Star Trek, computers are "programmed" by a simple verbal command "Computer! Plot a course to Alpha Centuri." To do this for real, two technologies are obviously needed: accurate speech recognition, and detailed natural language processing. However, a third technology is less obvious — these actions are generated by giving a statement of what the computer should do rather than how to do it. Automatic conversion of such statements of desired behaviour into program code are a core component of a current research project in the University of Kent Computing Laboratory.

This work is based on genetic programming, a technique first explored in the early 1990's which generates program code automatically for a task. Genetic programming takes a population of putative solutions, evaluates their success on the programming task at hand, and then uses processes such as mutation and crossover to generate new programs.

Our project focuses on the evaluation part. In traditional GP, the success of a program is measured by its success on a training set of input data. However, for safety-critical applications, this is a problem — how can we be certain that the program will work outwith the dataset used for training? Furthermore, for complex tasks, simply running the program often enough on enough training data makes applying these evolutionary methods intractable.

The approach that we have taken to this is to concentrate on automated analysis of the program and its properties rather than running it on test data. This has used a number of techniques, such as static analysis of variables in the program, constraint satisfaction, and model checking. We will use this latter technique as an example in the rest of this article.

Model checking is concerned with describing a program in terms of temporal logic statements about program state. That is, we describe how the state of a program will change over time, by a mixture of traditional predicate logic and statements about time-dependencies: for example "until", "while", and "always".

As an example, consider a simple coffee machine that can be in three states: start, coffee (dispensing coffee) and reset (in the process of resetting). In addition there is a variable describing whether there is a coin in the machine or not. A temporal logic description of the desired behaviour of this machine is given in Figure 1.

Variables	coin, taking values {0,1}	
States	coffee	
Statements	EF("coin=1")	1. We can reach a state where the coin has inserted
	EF("coffee") EF("reset")	2. We can reach a state labelled "coffee" 3. We can reach a state labelled "reset"
	AG("userStart" –>AF("coffee"))	4. Once the user process has started, coffee must be served
	EF("coffee")&EF("reset")& AG("coffee"–>EF("reset"))	5. States coffee and reset can be reached, and whenever coffee has been reached, the machine *can* reset
	EF("coffee")&EF("reset")& AG("coffee"–>AF("reset"))	6. States coffee and reset can be reached, and whenever coffee has been reached, the machine *must* reset
	EF("coffee")&AG("coffee"->AX(AG(!"coffee")))	7. Once coffee has been served, we must not serve another coffee
	EF("coin=0")&EF("coffee")& AG("coin=0"–>AF("coffee"))	8. If no coin has been inserted, we cannot get a coffee

Figure 1: Coffee Machine Specification

How do we use this? Model checking algorithms exist which take a program (in the form of a finite state machine) and confirm, by intelligent enumeration and grouping of states, whether each of the statements in a description such as that in Figure 1 are satisfied, regardless of the route taken through the program. More sophisticated algorithms can also give an estimate of the number of runs in which a program will satisfy the statements.

We use such a program in the fitness evaluation stage of our genetic programming system. The number of statements that are satisfied by each program in the population is used as a measure of how good that program is at the task, and therefore how likely that program will be used as a "parent" of programs in the next generation. Some examples of coffee-machine programs that have been successfully evolved by our system are given in Figure 2.

This ongoing project is looking at many different problems — so far we have applied these techniques to dynamic geometrical placement problems,robotics with safety constraints, and the evolution of sorting networks.A particular priority for future work will be evolving programs that combine traditional test-based methods of fitness evaluation with methods based on the analysis of programs. Overall the aim of this project is to move one small step closer to an industrial revolution in programming, where bespoke programs are generated on the fly by automatic programming mechanisms.

Colin Johnson
Department of Computer Science
University of Kent
Email: C.G.Johnson@kent.ac.uk

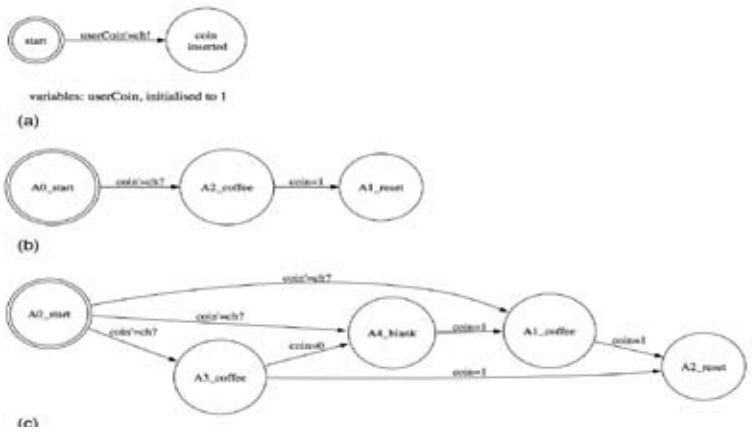

Figure 2: Successfully evolved coffee machine automata

Conference Report from the Genetic and Evolutionary Computation Conference (GECCO 2006)

From July 8th to 12th, the Genetic and Evolutionary Computation Conference (GECCO 2006) was held in Seattle, Washington, USA. GECCO 2006 continues the tradition of "One Conference, many mini-conferences". Over 500 researchers from the entire spectrum of research in genetic and evolutionary computation gathered together to report and hear about the latest developments in the field. The conference included two days of workshops and tutorials, fifteen program tracks, human competitions, etc. In this report, I will describe several tutorials and tracks of the conference, through which I hope to give you an idea of what's been going on in those areas of evolutionary computation.

Neuroevolution is to evolve neural networks using evolutionary algorithms. In GECCO this year, Risto Miikkulainen has given a tutorial about neuroevolution. After introducing the important concepts, challenges and developed solutions in this field, he pointed out some promising directions and showed a lot of impressive demos, either of benchmark problems or novel real-world applications. Compared to traditional methods of training neural networks, neuroevolution is especially strong in solving problems where the targets or correct actions are not always available for every situation. Thus, a feasible solution is trying with evolution algorithms to explore different actions and observe their outcomes. The conventional neuroevolution approach studies about evolving the connection weights of fixed-topology networks. In the tutorial, Prof Risto Miikkulainen paid particular attention to introducing advanced neuroevolution methods, including (1) evolving specified neurons that accomplish different subtasks of a complex problem; (2) evolving neural network topologies through complexification, starting with very simple architecture; (3) indirect encoding, for instance, to evolve the instructions of constructing a network instead of network itself; (4) combining learning and evolution. Above approaches have given promising results in the fields of robot control and rocket control, automated driving and collision warning, coordination of multi-agent systems, game playing, and resource optimization, as the demos reveal. So they may be quite inspiring for researchers of evolutionary algorithms or neural networks to work with.

As an extension of Evolutionary Algorithms, **Coevolutionary algorithms** (CEAs) offer the potential to address problems for which no accurate evaluation function is known, and the evaluation of an individual depends on other evolving individuals. Coevolution's first appearing in GECCO was a workshop given by Richard K. Belew and Hugues Juillé in 2001. The past a few years have been an exciting time for coevolution research. In GECCO 2006, tutorials of introductory and advanced coevolution were organized, and coevolution has one program track.

The introductory coevolution tutorial was presented by Sevan Ficici and Anthony Bucci. They started from reviewing the early work and notable results through the 1990s. In so doing, they made clear the range of applications for which coevolution has been applied and the common algorithmic paradigms. The attendees were impressed that Coevolutionary algorithms advance over traditional evolutionary algorithms in terms of their adaptability and potential open-endedness. Nevertheless, the presenter then clarified a variety of evident pathologies that coevolutionary algorithms have frequently exhibited followed by an outline of several attempts to remedy them.

The advanced coevolution tutorial was divided into three parts, being respectively given by Edwin de Jong, Paul Wiegand and Kenneth Stanley. In first part, Edwin de Jong specified the problems that coeovlutionary algorithms have been trying on and formalized the commonly employed solutions, for instance, cooperative coevolution, pareto-optimal coevolution, etc. Then Paul Wiegand introduced and analysed many developed tools of coevolution, such as evolutionary game theory, different coevolution models distinguished according to features of individuals' interaction or the population, and performance measurement. The final part was presented by Kenneth Stanley where he addressed the problem of "Representation in coevolution" and looked at how this affects the coevolutionary algorithms' performance. The strategy proposed was to design complexification methods which enable continual elaboration of evolved solutions.

In GECCO 2006's Coevolution Track, 8 papers and 4 posters were presented. These work mainly fell into the range of theoretical or empirical study of coevolutionary algorithms. It is hoped that more real-world applications using CEAs will be developing and shown next year.

Evolutionary Computation in practice (ECP). This track differed from GECCO's another normal program Track "The Real-World Applications": presenters of ECP track are generally researchers or managers from industry, governmental agencies and other public sectors. This provides a special forum for Evolutionary Computation practitioners to exchange ideas and to promote a wider usage of the technology in solving real-world problems.

A very interesting presentation of this year was given by Gregory Hornby, who reported the statistical results from the "Survey on members of the Evolutionary Computation community" that was held on and after GECCO 2005. It showed that Evolutionary Computation has become more accepted by the computing community, over the past decade across the world, although there is still a distance from the ideal. This talk was informative for students who would face with job hunting in the fields of Evolutionary Computation.

Finally, in this report, it's only possible to illustrate to you a small fraction of the exciting work presented in GECCO 2006. More information can be found at GECCO 2006's homepage http://www.sigevo.org/gecco-2006/, for the readers who are interested. Excitingly, next year GECCO will be leaving the USA for the first time, and will be located in London, UK, in July 2007.

Qijun Zhang
Interdisciplinary Centre for Computer Music Research
University of Plymouth
Email: qijun.zhang@plymouth.ac.uk

AISB Annual Convention Call for Papers

The AISB convention, 2nd-5th 2007 (Newcastle University, Newcastle upon Tyne, UK) will host nine symposia. The deadline for submitting papers is 8th January 2007 (some symposia 22nd January). Further information can be found at www.aisb.org.uk/convention/aisb07/index

Affective smart environments

Ambient Intelligence (AmI) is an emerging and popular research field, with the goal to create "smart" environments that react in an attentive, adaptive and proactive way to the presence and activities of humans.

Topics of interest include but are not limited to:
Non-invasive methods for sensing, recognizing and modelling the emotional state of users in 'natural', everyday situations
Methods and models for profiling emotion information
- Methods for building the inhabitants' group profiles from their individual models
- Methods for learning long-term features, from tracing of interaction histories.
- Methods for inferring how to adapt the environment to the recognized situation
- Methods for enforcing the sense of trust in the environment
- Affective Conversational Agents
- Social Robots
- Natural language and speech-based dialogue simulators

AI in mobile systems

Today's information technology is rapidly moving towards small computerised consumer devices and hi-tech personal appliances from the desks of research labs onto sales shelves and into our daily life. These include PDAs, embedded computers in cameras, cars, and mobile phones as well as high performance wearable computers and tablet PCs. Many of these devices are becoming essential tools that we increasingly rely on both in private and in professional settings.

The scope of interest includes but is not limited to the following topics (in no particular order):
- mechanisms for location and context awareness (e.g. Knowledge-based acquisition of contextual information, inference of location)
- spatio-temporal issues and methods in mobile and ubiquitous computing (e.g. correlation between spatial abstractions and different interface modalities)
- multi-modal interfaces for mobile and ubiquitous systems
- user interfaces that adapt to the current situation as well as to resource availability (e.g. modelling the trade-offs between reasoning capabilities, resource consumption and real-time constraints)
- plan-based approaches for interaction and adaptation
- user modelling for mobile and ubiquitous computing
- scalable ontologies

Artificial societies for ambient intelligence

In this one-day Symposium we would like to explore the application and development of agent societies for AmI, establish a body of knowledge and a theoretical framework for this, and use this framework to relate existing work on areas such as the semantic web, cognitive and social agents, and ambient and ubiquitous technologies.

We also hope to present current research in the area of agent societies for AmI, where human activities are supported by social organisations of agents, computing devices or both, and assess the outcomes of such research. The Symposium will identify issues for future investigation, establish links between researchers and encourage international collaborations.

Topics of relevance include, but are not limited to, the following.
- Social architectures
- Agent interaction
- Reasoning and knowledge representation
- Reactivity and pro-activity
- Learning
- Decision making
- Co-operation and co-ordination
- Social emergence and evolution
- Normative reasoning and regulations
- Security, trust and privacy
- Service-oriented approaches
- Interaction design and interfaces
- Mobility
- Applications

Games for education

The mechanisms of imitation and social learning are not well-understood, and the connections to social interaction, communication, development, and learning are deep, as recent research from various disciplines has started to uncover. Comparison of imitation in animals and artifacts reveals that easy tasks for machines can be hard tasks for animals and vice-versa. However, computational complexity issues do not explain, by themselves, the existence or not of imitation behaviours in animals, and the integration of higher level cognitive capabilities like agent's goals, intentions and emotions, may play a fundamental role in explaining these differences. Areas of interest include bute are not limited to:

- Cognitive Development and Imitation
- Neurobiological Foundations of Imitation
- Social interaction and Imitation
- Language acquisition
- Imitation, Intentionality and Communication
- Imitation in Animals
- Learning by Imitation to bootstrap the acquisition of skills & knowledge
- The Role of Imitation in the Development of Social Cognition
- Robot Imitation
- Computational mechanisms of imitation
- Joint-attention and perspective taking
- Cultural transmission of skills
- Teaching and scaffolding of behaviours
- Imitation and motivation

AISB Annual Convention Call for Papers

continued from p. 10

Mindful environments

Research into expressive characters is a growing field, and simultaneously new work in human-robot interaction (HRI) has also focused on issues of expressive behaviour. However, progress is hampered by the need to integrate work in various sub-fields of psychology, in natural language processing, speech and in computer graphics, carried out by many different groups in communities that do not always intersect. Other areas, such as integrating gesture and facial expression and affective state with language and speech, are less developed but vital to progress. Contributions are invited in one or more of the following topics:

empirical studies of gesture and facial expression
- frameworks for the specification and analysis of gesture and facial
- expression for expressive characters
- gesture and facial expression modelling and animation
- evaluation of expressive characters
- appropriate natural language processing architectures
- natural language generation
- dialogue systems and question answering
- language and gesture coordination
- language and facial expression coordination
- language and action integration
- emotional language
- personality modelling, language and speech
- lip synchronisation and combination with facial expression
- affect in speech synthesis and recognition

Spatial reasoning and communication

Spatial cognition has a significant role in our everyday lives. When commuting from our home to our work place, we need a spatial map that enables us to find a reasonable route through the city's road network. When looking for a folder or a textbook in our office, it helps if we know the spatial location at which the item is to be found. When constructing a building, it is essential to understand the spatial-functional relations between the parts of the building: ceilings have to be supported by walls, windows should be inside walls, etc. We welcomes contributions to all aspects of spatial cognition concerning communication and computation, including (but not limited to) new results about:

- Formal analyses of spatial calculi and models
- Integration of spatial calculi with other reasoning formalisms
  (e.g., temporal calculi)
- Spatial database queries
- Context-sensitive interpretation and formalization of spatial language, and its mediation towards system-relevant aspects, for example via spatial ontologies
- Spatial human-machine communication via language and/or other modalities
- Computational treatment of functional-spatial relationships in natural environments
- Handling of different spatial granularities
- Dealing with uncertainty in spatial cognition

The reign of Catz and Dogz?

A major concern for human computer interaction researchers is how to construct interfaces to future ambient and pervasive technologies which are naturalistic, unobtrusive and implicit. Sustained consumer interest in off-the-shelf robotic animals provides evidence of the widespread appeal of interacting with artificial, representations of creatures.. As the designers of such toys and applications are no doubt aware, an accepted consensus within anthrozoologic research is the quantifiable positive effects of human-animal relationships. Topics of interest include, but are not limited to:

- virtual creature/character interaction
- embodied versus screen based interactions
- social and ethological robotics
- virtual pets and companions
- comparisons of interactions between real and artificial creatures
- virtual creatures as interfaces to larger pervasive systems
- virtual creatures inhabiting mobile devices
- the uncanny valley and its effect on virtual creatures/characters
- the 'illusion of life' – making believable artificial creatures
- concurrent interaction with a number of virtual characters
- presence of persona in virtual characters and embodied devices
- companionship and comfort from virtual characters/creature

The Life of A. Hacker

by Fr. Aloysius Hacker

Cognitive Divinity Programme
Institute of Applied Epistemology

Episode 3: Accelerating intellect

Funded by lucrative income from my company **BOOTLEG™** (**B**lack-boxes, **O**ther **O**utfits and **T**hingamajigs, **L**ikewise **E**lectronic **G**izmos), in 1956 I set out for Dartmouth, USA for the first ever AI conference. I had already made contact with two of the pioneers of AI: Alan Chewring and John McCarthief. This was an unparalleled opportunity to meet the other key players — people such as The Marvellous Maxsky and the Herb Knowalls. The presentation of my **CHESS™** (**CHESS™** is a **H**euristic **E**xpert-**S**ystem **S**olver) program attracted a lot of interest, including some job offers. With my eye for recognising a rising star, I wisely chose the offer from The Maxsky to help co-found the new AI Group at **META** (**M**echanics, **E**lectronics and **T**echnology **A**cademy) in Boston. As the **META** AI Group grew into the hotbed of AI research, I hoped my fortune would grow with it.

The batch-processing computers of the late 50s were painful and slow to use. Frustrated by the delays and wasted time this caused, I dreamt of an online, interactive system that everyone could use to edit, run and analyse their programs in real time. My ideas to realise this dream were very well received at **META**, and went on to become the world-famous Project **HACK** (**H**eterogeneous **A**ccess to **C**omputational **K**nowledge), which gave birth to modern, networked computing. My influence was also felt in many of the other major AI breakthroughs at **META**. For instance, my **HOTLIPS™** (**H**igher-**O**rder **T**heories for **I**nference and **P**roblem **S**olving) program language anticipated both **LISP** and **PROLOG** by several years, but was never acknowledged by either the functional or the logic programming communities. My **QWERTY™** (**Q**ueries in, in **W**ritten **E**nglish; **R**esponses out, **T**o **Y**ou) natural language understanding system was a major unrecognised influence on Lous O'Grad's **SCHRDLU**. Other people got the PhDs, the prizes and the promotions, whereas I could not even get published. Disheartened by this lack of recognition of my genius, I resolved to strike out on my own.

Religion and computing had both played formative roles in my early life, so in defining my new initiative, it was natural to combine them; they made natural bedfellows. Computational thinking held the promise of illuminating many of the central, but illusive enigmas of theology. With my help, several philosophers had already used the software/hardware distinction to address the mind/body problem. Could miracles, the holy trinity and free will be far behind? Thus was born the field of Computational Theology. Back in the UK, I founded **CATHOLIC™** (**C**hurch of **A**loysius **T**heobald **H**acker for **O**rdinations, **L**iturgy, **I**nquisitions and **C**hristenings). **CATHOLIC™**'s First Act was to spin out a new pedagogic enterprise: The Institute of Applied Epistemology, which was aimed at satisfying the thirst for qualifications within the newly meritocratic, British, swinging sixties. It's a little known fact, for instance, that it was my Institute that furnished the Rev. Dr. Ian Paisley with his PhD. Its Second Act was to inaugurate the Cognitive Divinity Programme: whose investigations into the deepest religious mysteries would bring the academic respectability to underpin the more profitable parts of my new adventure. A New Age had begun.