AISB Quarterly

No.115

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

In this issue:

Features

- **Cox and Young** 2
  Exploratory learning
- **Potter/Tate/Dalton** 3
  Collaborative task support
- **Anderson/McOwan** 4
  Towards affective computing
- **Angelo Cangelosi** 5
  Grounding language
- **Luc Berthouze** 6
  Models of motor development
- **Niels Taatgen** 7
  Production compilation
- **Sandy Louchart** 8
  Empathic characters

Reviews

- **Book** 9
  David Young on Computational Vision
- **Conference** 10
  Carlos Bento on Creative Systems

News 11
- *Membership report*
- *AISB Convention '04*
- Expert Systems *offer*

Father Hacker

- *How to be a Research Leader* 12

Danger theory: the missing link between artificial immune systems and intrusion detection

The central challenge in computer security is determining the difference between normal and potentially harmful activity. For half a century, developers have protected their systems by coding rules that identify and block specific events. However, the nature of current and future threats, in conjunction with ever larger IT systems, urgently requires the development of automated and adaptive defensive tools. A promising solution is emerging in the form of artificial immune systems (AISs). The human immune system (HIS) can detect and defend against harmful and previously unseen invaders. The question is, can we build a similar intrusion detection system (IDS) for our computers? Presumably, such systems would then have the same beneficial properties as HIS: error tolerance, adaptation, and self-monitoring, for instance.

Current AISs have been successful on test systems, but the algorithms cannot be scaled up to real-world requirements. This is because they rely on the ability to discriminate between what is and isn't the system's *self*, as stipulated in classical immunology. However, immunologists are increasingly finding fault with this traditional thinking and a new *danger theory* (DT) is emerging.

The new theory suggests that the immune system reacts to threats based on the correlation of various danger signals, and it provides a method of *grounding* the immune response: i.e. linking it directly to the attacker. Little is currently understood of the precise nature and correlation of these signals and the theory is a topic of hot debate. It is the aim of our research to investigate this correlation and to translate the DT into the realms of computer security, thereby creating an AIS that is no longer limited by self/non-self discrimination. It should be noted that we do not intend to defend this controversial theory *per se* although, as a deliverable, this project will add to the body of knowledge in this area. Rather, we are interested in the merits of the DT for scaling up AIS applications.

With our growing understanding of cellular components involved with apoptosis, it will be possible to compare the differential proteomic profile between necrotic ('bad') and apoptotic ('good' or 'planned') cell death, particularly with respect to the activation of antigen presenting cells. A vital necessity will be maintaining the physiological relevance of the system to be used, so that the power of the DT in protecting against false positives is preserved. In essence, it is thought that apoptosis has a suppressive effect and necrosis a stimulatory immunological effect: although they might not actually be as distinct as currently thought. In the IDS context, this can be read in two ways: either the necrotic signals act to say that the previous pattern of apoptotic signals is dangerous, or the apoptotic signals indicate that the necrotic signals are a false alarm.

A variety of contextual clues may be essential for a meaningful 'danger signal', and immunological studies will provide a framework of ideas showing how such danger is assessed in the human immune system. Such ideas can be fruitfully applied to the AIS arena. In the latter context, the danger signals should show up after limited attack to minimise damage, and therefore have to be quickly and automatically measurable. Once the danger signal has been transmitted, the AIS can react to those artificial antigens that are near the emitter of the danger signal. This allows the AIS to pay special attention to dangerous components, and would have the advantages of detecting rapidly spreading viruses and scanning intrusions at an early stage to prevent serious damage. See Figure 1 for a graphical illustration.

This project is funded by EPSRC with further support from Hewlett Packard Labs, Bristol, and ECSC Computer Security, Bradford.

Uwe Aickelin
School of Computer Science
University of Nottingham
uxa@cs.nott.ac.uk
http://www.aickelin.com

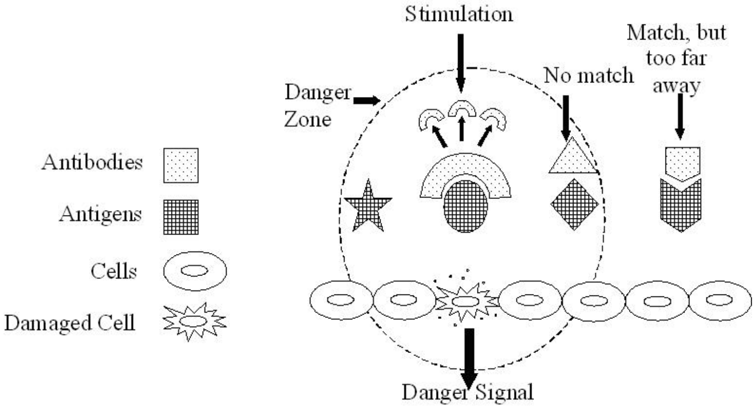

Figure 1. The Danger Theory model.

A rational framework for modelling exploratory learning

In the real world, people rarely read instructions. Instead, they learn a new piece of equipment or software application by actually trying to use it: drawing on a combination of prior knowledge, information from the interface itself, and problem-solving skills. This phenomenon is known as exploratory learning. Though most equipment is not designed to be learned by exploration alone, a study investigating the real-world strategies users employ[1] showed that people prefer to learn by exploration in the context of a task they need to perform. They dislike taking time out to experiment with it, or work through documentation in a task-independent manner.

Our earlier research[2] led to the development of a rational framework for modelling this kind of exploratory learning. We have already used it to explain empirical findings such as why people do and learn different things from free, as against focused, exploration.[3] We are now using it as a framework to build models of single-level menu exploration which we will compare to data gained from empirical studies.

The framework and how it works

At any moment, the model has a number of possible things it can do. These might be pressing one of a number of buttons, considering some feedback that has just been received, reading a label, formulating a hypothesis to test, etc. These are referred to as Exploratory Acts or EAs. EAs differ in the number of stages they consist of and therefore in the amount of time and effort required to carry them out. For example, pressing a button would require a single quantum of effort, whereas interpreting feedback from the device might require a number of quanta: one to read the display; another to realise that the information has changed; and a third to understand what type of change it was.

Each EA has a cost and a value determined by the costs and values of its constituent quanta. Each quantum has a cost reflecting the time it would take to perform it, and a value related to the estimated increase in information that would be gained. The cost of an EA that proposes a single button press, therefore, is likely to be low as it only requires one quantum of processing: but the value may be high or low depending on what is known about the button already. The cost of an EA that proposes interpreting some feedback would be higher than that of the button press, however. This is because it requires more processing quanta (three in this example) and, again, the value would vary depending on the amount of information expected to be gained.

Which of the proposed EAs is chosen in any given situation is determined by rational analysis, which suggests that, when trying to learn about the device, the next move chosen will be the one that is believed to elicit the highest amount of information for the least cost. The framework describes a cycle of three stages. In stage one, the efficiency of all the EAs possible at that moment are calculated. The efficiency of each is equal to the expected amount of information gained by the EA, divided by the cost of executing it. In stage two, the EA with the highest efficiency is chosen. The model will therefore choose whichever proposes the highest information gain per unit of cost. Finally, in stage three, the chosen EA is executed.

Comparison of the model to empirical data

We have built a model that explores the menu of the *Cricket Graph* programme on an Apple Macintosh. The model is given a task called *create a graph*. This menu and task have been used in both empirical studies[4] and cognitive models of exploratory behaviour.[5] The model makes use of four different types of EA with ascending levels of associated cost. These are: reading the label; considering the semantic relatedness of the label to the goal; choosing to pull down the menu to view the items under the top level label; and deciding that a particular label is the correct one to choose.

Initial results suggest that the model shows evidence of label following and spatial search strategies, as demonstrated by the empirical and modelling work referred to above. The model appears to ‘zoom in’ fairly quickly on the `Graph’ label and, once it has found it, identifies this label as the one it would choose. It would appear, therefore, that the model does not exhibit iteratively deepening attention as described by Rieman et al..[5] It may be the case that, in this example, the information ‘scent’ from the label is sufficiently strong for us to expect this pattern of behaviour. We are currently investigating how the model behaves on other menus with different patterns of information scent.

Anna L Cox & Richard M Young
Depts of Computer Science & Psychology
University of Hertfordshire, UK
Email: {a.cox, r.m.young}@herts.ac.uk

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

ISSN 0268-4179

References

1. Rieman, J. A. *Field Study of Exploratory Learning Strategies*, **ACM Transactions on Computer-Human Interaction 3** (3), pp. 189-218, 1996.
2. A. L. Cox, ‘**Exploratory Learning of Interactive Devices: What People Do and What People Learn,** unpublished Ph.D. thesis, University of Hertfordshire, UK, 2002.
3. R. M. Young, and A. L. Cox, *A New Rational Framework for Modelling Exploratory Device Learning ... but does it Fit with ACT-R?* **Proc. of the 7th ACT-R Workshop,** Carnegie Mellon University, 2000.
4. M. Franzke, *Turning Research into Practice: Characteristics of Display-Based Interaction,* **Proc. ACM CHI Conf. on Human Factors in Computing Systems,** 1995.
5. J. Rieman, R. M. Young, A. Howes, *A Dual-Space Model of Iteratively Deepening Exploratory Learning*, **Int’l J. of Human-Computer Studies 44,** pp. 743-775, 1996.

Collaborative task support and e-Response

Rather than seeking to automate the tasks performed by humans, the trend in modern AI research and application development is towards providing support to human agents in the workplace. The impetus for this lies in an acknowledgement of the differing capabilities of humans and computers, and its aim is to engineer environments where these capabilities will complement each other to greatest effect. Another influence is the dramatic shift in work practices in recent years, with the rise of the internet and the web (and with the semantic web on the horizon), ensuring that knowledge management has become central to the philosophy of the modern organisation.

The I-X programme is typical of this type of modern AI project. Its overall aim is to create an enabling environment for mixed-initiative (i.e., involving both human and computer agents) synthesis tasks. The definition of such a task, as it is considered here, is general enough to embrace tasks as diverse as designing an aircraft engine, devising a marketing strategy, and writing a joint report. Such tasks occur regularly in organisations and usually require some degree of creativity, something that is difficult to emulate on a computer. (This is not to say that computers do not have a role to play in the task—for instance, in simulating design concepts.)

I-X draws on (and is a natural successor to) several decades of AI experience at Edinburgh in planning, scheduling, and—more recently—process, workflow and activity management. Born of this experience, and lying at the conceptual heart of the programme, is a unifying upper ontology for a shared representation of a synthesis task, whatever the precise nature of the task or its domain may be. This conceptualisation, the *<I-N-C-A>* ontology, is based on the notion of both the processes governing and the products emerging from the task being composed of abstract nodes. These are related by a series of constraints, and issues concerning them are cyclically generated and resolved so as to refine the set of nodes and their relationships.

This model allows flexibility in the extent and nature of the formalisation of the representation. So, an informal approach to representing a constraint might suffice when coordinating joint memorandum-writing activities ('finish by next Friday'). However, a more formal scheme might be required for a design task where precision is required or automated constraint-solver agents are to be invoked ('has-orientation(fin-9102, horizontal)').

As well as encouraging a well-founded encapsulation of the task, the model also provides the basis for a systems architecture and communication framework, allowing the concrete realisation of I-X systems. Figure 1 depicts the manner in which the I-X tools serve to construct a task-solving environment.

This technology is being put to use in the Collaborative Advanced Knowledge Technologies in the Grid (CoAKTinG) project, funded by the UK e-Science Programme. Its aim is to assist remote collaboration by creating a mediated conceptual space in which to work. This is realised by adopting concepts from AI and knowledge management—ontologies, issue-based information systems, presence visualisations and so on—and embodying these in a set of complementary tools. These include the I-X tools, as well as others developed by the project partners (the AIAI, the Knowledge Media Institute at The Open University, and the Intelligence, Agents, Multimedia group at the University of Southampton). One CoAKTinG demonstration scenario, called e-Response, concerns an evolving environmental emergency: an oil spill threatening a sea-bird reserve. The response team (which, collectively, has a wide-ranging scientific background) has to generate a plan for responding to this emergency: the creation of this plan is the synthesis task here.

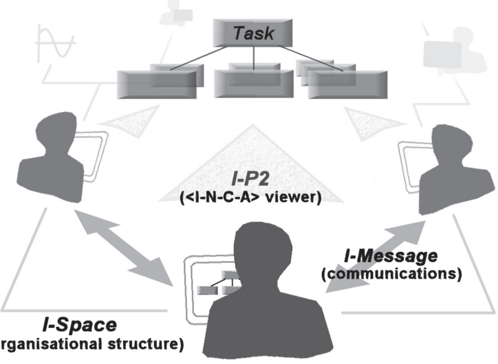

Figure 1. Shown is the I-X environment. The I-P2 (Process Panel) tool visualises the current task for a user; the I-Space tool establishes organisational relationships between agents, dictating the sort of appropriate interactions among them, and the I-Message tool allows messages of variable formality to be sent. In addition, various editors allow human agents to express and formalise <I-N-C-A> entities, and task-specific plug-ins can provide additional visualisations, editors, and solvers.

In constructing their plan, the members of the team follow—both individually and as a group—specific response procedures. While some of these may be extemporised and contingent on circumstances, others may be instances of 'standard operating procedures': generic approaches to archetypal activities that can be downloaded from a central web-store. In addition to the human agents in this environment, automated agents exist to provide tide data and weather forecasts, simulate the progress of the oil slick, poll centralised data stores for details of available human expertise in specific fields, and so on. The interactions are governed by the activities, issues, and constraints that arise, and mediated by the I-X interfaces of the team members. These present the current state of the collaboration from their individual perspectives, and allow them to decompose activities, refine elements of the plan, delegate issues, invoke the automated agents, etc.: all serving to facilitate the team's task.

Stephen Potter, Austin Tate and Jeff Dalton
Artificial Intelligence Applications Institute
Informatics
The University of Edinburgh
Edinburgh, UK
E-mail: stephenp@inf.ed.ac.uk
http://i-x.info
http://www.aktors.org/coakting
http://e-response.org

Towards affective computing: a real-time system for recognising facial expression

Affective computing has been defined as, "computing that relates to, arises from, or deliberately influences emotion." As humans, our evolutionary and developmental histories are fundamentally social, giving rise to a vast repertoire of skills for social interaction that allow us to assess, identify and predict the 'internal states' of others. A vital component of this capacity to adapt appropriately whilst interacting is an ability to display, recognise, and interpret facial expressions. Allowing synthetic systems to make use of such modes of interaction would be a significant step forward, and in this article we describe a computer technology that allows a system to respond in real time to the facial expression of the user. Both sociable robots and synthetic characters have been enabled to express emotional state. However, less progress has been made in providing such agents with the skills to understand the facial expression of others in real time.

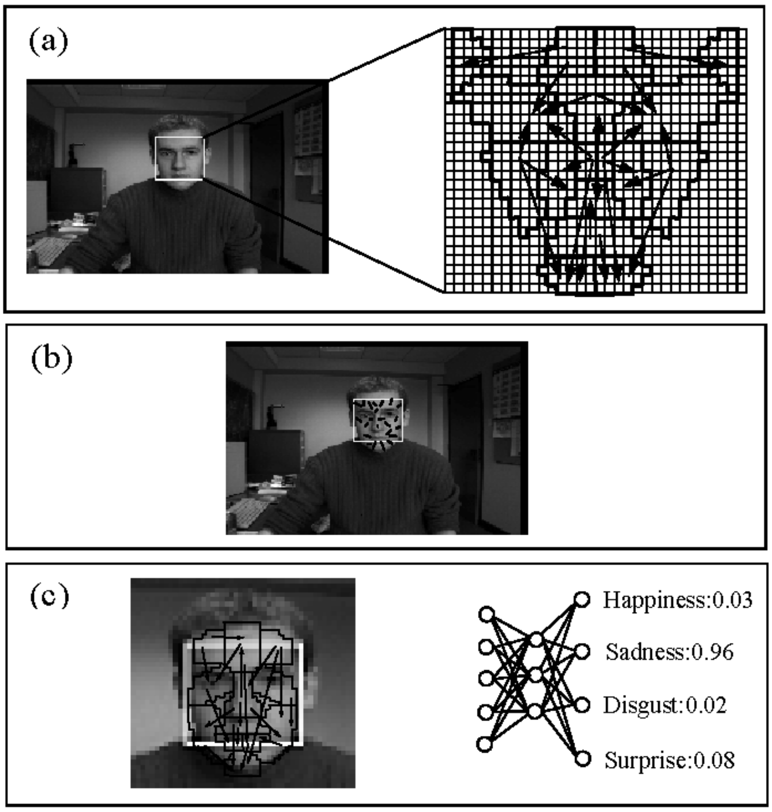

Figure 1: (a) The tracker locates the face using a modified ratio template algorithm. Greyscale values are averaged over the regions indicated, and ratios taken (shown by the arrows). A ratio is satisfied if the average from the first region (arrow tail) to the second (head) exceeds 1.1. (b) Motion detection. (c) For expression classification, motion data is averaged over key parts of the face and ratios of the averaged motion taken. This data is entered into multi-layer perceptrons trained using back propagation.

Our system is a fully-automated architecture that runs in real time at 4fps on a 450MHz Pentium III machine with Matrox Genesis DSP boards. It is able to operate effectively in cluttered and dynamic scenes and recognise four of the six facial expressions universally associated with unique emotional states: happiness, sadness, disgust, and surprise.

There are three main components to the system; face tracking, facial-motion extraction, and expression recognition. To characterise the users facial emotions we first need to identify the face in the scene. The face tracker uses a modified version of the ratio template algorithm.[2] This technique searches the scene for the best match of luminance ratios in a spatial pattern modelling the structure of the human face (see Figure 1a). This ratio-based approach is able to detect frontal views of faces under a range of scene-lighting conditions. Our system improves the algorithm's performance by incorporating proportions into the spatial face model that are more biologically plausible—which improves tolerance to illumination change—and by the examination of higher-order relationships within the initial ratio-template measures.

Our facial template is automatically aligned during face location: we then apply a robust optical flow algorithm[3] at this location (Figure 1b) to determine facial motion (averaged over specified areas such as the chin or brow). By using only motion information, the task of expression recognition is simplified: variations in the texture of different faces are ignored. Overall, head motion is cancelled out using ratios of averaged motion data (see Figure 1c).

The processed motion data is currently entered into multi-layer perceptrons (MLPs), trained using back propagation to classify facial expression. The MLPs were trained and tested using sequences taken from the Cohn-Kanade AU-coded facial expression database.[4] An overall correct-recognition rate of 91% was achieved on the 57 image sequences included in the test set.

To demonstrate the utility of this expression-recognition technology in affective computing applications, it was incorporated into simple, prototype, real-time applications. We have developed, and are testing, a simple chat-room application that automatically inserts emoticons—widely understood symbolic abbreviations, such as :) for happy—into the text for the user. A more complex system currently under development monitors the facial expression of a user and, if one of the learned expressions is displayed, launches a desktop application appropriate to the expression measured. For instance, if the user has a sad expression on his or her face, then either an appropriate web page is automatically opened, or music is played to raise the spirits. The ability to drive computer applications via facial expressions could also have benefits for disabled users.

Keith Anderson* and Peter McOwan**
Department of Computer Science,
Queen Mary, University of London
*E-mail: ka@dcs.qmul.ac.uk
http://www.dcs.qmul.ac.uk/~ka
**E-mail: pmco@dcs.qmul.ac.uk
http://www.dcs.qmul.ac.uk/~pmco

References
1. K. Anderson and P. W. McOwan, *Real-time emotion recognition using biologically inspired models,* **4th Int'l Conf. on Audio & Video Based Biometric Person Authentication,** p. 119, 2003.
2. B. Scassellati, *Eye finding via face detection for a foveated, active vision system,* **Proc. 15th Nat. Conf. on AI,** 1998.
3. A. Johnston, P. W. McOwan, and C.P. Benton, *Robust Velocity Computation from a Biologically Motivated Model of Motion Perception,* **Proc. Royal Society of London 266,** p. 509, 1999.
4. J.F. Cohn, A. Zlochowe, J. Lien, and T. Kanade, *Automated face analysis by feature point tracking has high concurrent validity with manual FACS coding,* **Psychophys. 36,** p. 35, 1999.

Grounding language in sensorimotor and cognitive categories

Language is the preferred means of communication not only amongst humans, but also between humans and machines (be they robots or simulated agents), and even between agents themselves. Scientists in artificial intelligence, robotics, and cognitive science use different approaches to model language depending on their scientific aims and methodologies. For example, in cognitive science, much effort has been put recently on studying the emergence of language using 'synthetic' methodologies such as adaptive behaviour and artificial life.[1] These models use groups of autonomous agents that interact via language games to exchange information about their physical and social environment. Their coordinated communication system is not externally imposed by the researcher, but emerges from the interaction between agents. In such models, the levels of detail of the representation of the agents and of their environment can vary significantly. This constitutes a continuum between abstract models, at one end, and situated and embodied robots at the other. Intermediate approaches also exist, based on grounded simulation models.

In the *abstract models* of the emergence of language, only a few essential communicative properties of the agents and environment are simulated.[2] The environment may consist of a list of abstract meanings (e.g., John, Mary, love) and a set of words (e.g., 'John', 'Mary', 'love', 'Xyz'). The architecture of the agent consists of rules that map meanings to signals (e.g. John for 'John'). This is a useful approach for studying the dynamics of the auto-organization of syntax and its dependence on specific, pre-identified factors. However, this methodology considers language as an independent and autonomous capability of the agent (and, indirectly, of humans). Such models are subject to the symbol-grounding problem. To be psychologically plausible, cognitive models must contain an intrinsic link between the symbols (words) used by the agents and their own representations (meanings) of the external world (referents). That is, communication symbols need to be grounded in the agents' sensorimotor and cognitive representations. The next two approaches to language modelling both deal satisfactorily with the symbol-grounding problem.

In the studies based on *grounded simulation models,* the agents' environment is modelled with a high degree of detail upon which emergent meanings (i.e. sensorimotor and cognitive categories) can be directly constructed and which constitute the grounding for language.[3] This type of model permits the investigation of the direct interactions amongst various sensorimotor, cognitive, linguistic and neural abilities of the agents.

Some grounded models use neural networks to control all these abilities, and the analysis of these networks can highlight the neural mechanisms responsible for the integration of language and cognition. For example, categorical perception measurements are used to compare the internal representations of words. In models where agents use verbs and nouns, it has been shown that categorical perception helps the emergence of syntax. In the neural networks of these agents, the similarity space of the representations of verbs is enhanced and optimised with respect to that of nouns (Figure 1). In addition, synthetic brain-imaging techniques permit the investigation of the direct relationship between language processing and sensorimotor abilities. Analyses have shown that the neural representations of syntactic word classes are sensitive to the level of integration of linguistic information and sensorimotor knowledge. The neural networks show functional organisations that reflect those observed in experiments on human language processing using brain imaging.[3]

Finally, a better way to study the grounding of language into behavioural and cognitive categories is through *embodied and situated robots*. Here, communication results from the dynamical interaction between the robot's physical body, its nervous and cognitive system (embodiment), and the external physical and social environment (situatedness). Amongst the robotic approaches, evolutionary robotics[4] has been proposed to study the grounding of language. In a recent model,[4] we evolved the control system of artificial agents that were asked to categorise objects (e.g. cubes and spheres, see Figure 2) and communicate about them. Simulation experiments showed that the ability to form categories from direct interaction with the environment constitutes the ground for subsequent evolution of names of objects (e.g. nouns). This was also extended to analyse the emergence of action names (e.g. verbs). Comparisons between experiments supported the language origin hypothesis that nouns precede verbs both in phylogenesis and ontogenesis.

The use of grounded models, with simulated agents or embodied robots, has significantly contributed to a better understanding of the strict interdependence between language and other sensorimotor and cognitive capabilities. This will be essential in

Angelo Cangelosi
University of Plymouth

Continued on page 8.

References

1. A. Cangelosi and D. Parisi (Eds.), **Simulating the Evolution of Language,** Springer, London, 2002.
2. S. Kirby, *Spontaneous evolution of linguistic structure: An iterated learning model of the emergence of regularity and irregularity*, **IEEE Trans. on Evol. Comp. 5** (2), pp. 102-110, 2001.
3. A. Cangelosi and D. Parisi, *The processing of verbs and nouns in neural networks: Insights from Synthetic Brain Imaging*, **Brain and Language,** *in press.*
4. D. Marocco, A. Cangelosi and S. Nolfi, *The emergence of communication in evolutionary robots*, **Phil. Trans. Royal Soc. A 361,** pp. 2397-2421, 2003.

Figure 1. Categorical perception effects in agents using verbs and nouns.

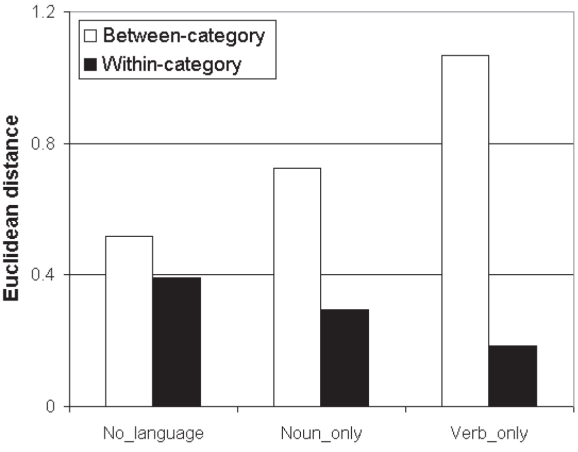

Towards computational models of motor development

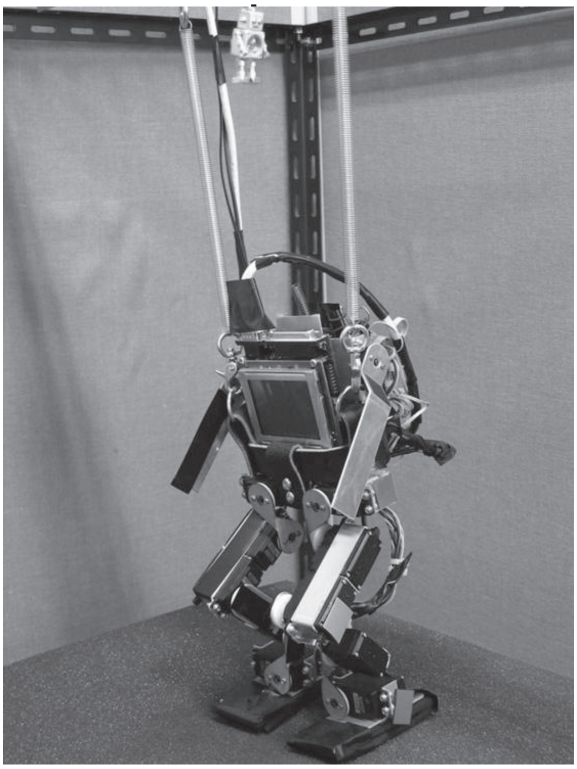

Figure 1. Our 12-DOF humanoid robot.

Since infants’ early cognitive development is largely driven by their sensorimotor interaction with the world, it is not surprising that their motor activity displays developmental characteristics similar to those of their more cognitive skills. For example, U-shaped development—a period of progressively worse performance followed by a gradual recovery to original levels—is observed in both infant stepping and learning of past tenses. Motor development has been defined as, “a progressive change in motor control and motor behaviour brought about by the interaction of both maturation and experience.”[1] While maturation refers to the development of the central nervous system (CNS), as well as to measurable changes in physical growth, experience refers to environmental factors that modify developmental characteristics through learning.

By emphasizing the interactive nature of motor activity, this definition shifts the focus away from the traditional neural pattern generation issue (i.e., the neurophysiological perspective) and closer to a *dynamical systems* understanding of motor development as a dynamic pattern generation process.[2] From this perspective, the temporal patterning of muscle activation is not the mechanism that produces the motor output: it is only representative of a motor pattern attained as a result of the system’s natural tendency to move towards steady states.

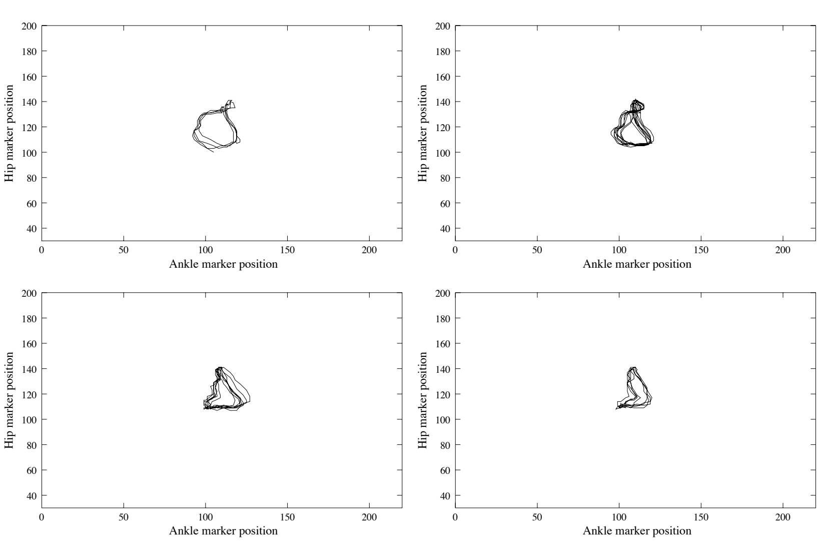

Figure 2. Phase plots corresponding to four preferred oscillatory patterns within one single parameter configuration.

Our research is concerned with the emergence of new behavioural forms, and the transitions between them. To study the interaction between the different sub-systems (organism, task, and environment) involved in producing a given motor pattern, we use a *synthetic modelling methodology,* an approach defined as, “understanding by building.” If the constructed system respects the embodiment of the system under study—i.e., if it preserves how the system’s neural structures interact with the environment through their physical embedding—then that system can be used to investigate the behaviour of the real system.

We used a 12-DOF (degrees of freedom) small-sized humanoid robot (see Figure 1) to investigate one particular instance of early motor activity: swinging. The robot was controlled by neural oscillators, elements of the CNS that can produce patterned output even in the absence of patterned sensory inputs. With their ability to *entrain* to afferent signals, oscillators are very suitable for studying the interaction of neural structures with proprioceptive and environmental feedback. The system’s exploration of its parameter space (timing adjustment) was simulated by a Boltzmann-like search process—a probabilistic exploration scheme—regulated by a value system and habituation.

When all degrees of freedom were involved, the system displayed several preferred states (see Figure 2) with abrupt phase transitions between them. While consistent with some observations of infants’ early exploratory activity, this result also illustrated what Bernstein[3] called the *degrees of freedom problem:* i.e. the fact that too many degrees of freedom are involved in producing a single motor trajectory. Bernstein suggested that, in children, the peripheral DOFs would initially be reduced to a minimum *(freezing)* and subsequently released *(freeing)* as a result of experiment and exercise. We tested this hypothesis by initially freezing the system’s lower degrees of freedom, and releasing them once the system settled in a stationary regime. This process resulted in the system showing high-amplitude and stable oscillatory patterns, (mostly) irrespective of the choice of initial parameters. We showed that the corresponding behavioural characteristics could not be attained by simply manipulating the neural parameters, but resulted from the *global*

Luc Berthouze
AIST Japan

Continued on page 9.

References
1. D. L. Gallahue, **Understanding Motor Development in Children**, John Wiley and Sons, New York, 1982.

References continued on page 9.

Production compilation: a versatile cognitive mechanism

The goal of cognitive modelling with cognitive architectures is to explain as much of cognition as possible using only a small set of explanatory mechanisms. In this article, I will focus on a particular mechanism, production compilation,[1] that learns new production rules on the basis of both existing ones and on declarative knowledge. The basic mechanism is very simple but, in combination with other ACT-R mechanisms and some general cognitive strategies, it allows for a broad spectrum of models that all entail learning a particular skill on the basis of discovery, experience, and instruction.

A fundamental assumption of the ACT-R theory is the distinction between a declarative (fact-based) and a procedural (rule-based) memory. Declarative memory contains pieces of knowledge that are open to conscious inspection, consisting of episodic knowledge, self-derived facts, and knowledge gained through perception. Each fact in declarative memory has an activation value that reflects its past relevance and connection to the current context. Procedural memory contains rules that act upon the current goal and state of the perceptual and motor systems. With each rule, a utility value is maintained that represents an estimate of its success rate and costs. Production rules are not open to conscious inspection.

The production compilation mechanism generates a new rule whenever two rules fire in sequence. The two rules are combined into one new rule that performs the same function. If the two rules interact with declarative memory, i.e., when the first rule requests a certain fact from declarative memory and the second rule uses the knowledge in some fashion, this fact from memory is substituted directly into the new rule. Although production compilation always specializes rules, it can nevertheless achieve some sort of generalisation when the rules comprise some general cognitive strategy that is applied to a specific example.

For instance, when children learn the past tense, they at some point discover the regularity in some verbs, which is that the past tense is the stem plus *-ed*. This rule can be learned by a model that tries to use an analogy strategy to generate past tenses.[1] Given a verb that has to be inflected in this way, the model attempts to retrieve an example from declarative memory (see Figure 1). If this example is not the word we are trying to find (e.g., we try to retrieve the past tense of *walk* but instead find the past tense of *work*) then the model will try to find a pattern in the retrieved example and apply this to the current verb. In the case of *work* and *work-ed,* the pattern is to add *-ed* after the stem.

Production compilation will specialise the analogy strategy on the basis of the regular example (in this case), producing the regular rule. The model produces the U-shaped learning effect associated with the past tense, and does so without external feedback on its own behaviour: this is in contrast to several existing models.

Let's consider another example. When people have to learn a new complicated task, they are initially very slow and make many errors. As experience increases, progress becomes faster and the number of errors decreases. This behaviour can be explained by a model that starts with a declarative representation of the task instructions and production rules that interpret these instructions. After some initial practice on a new task, experience will build up, and those experiences can be retrieved later when a similar situation occurs.

We[3] modelled this process in the case of the Kanfer-Ackerman Air Traffic Controller task,[4] one in which different planes have to be landed on different runways in different weather conditions. Figure 1 shows two examples of how production compilation can produce efficient rules that produce a significant speed-up in performance. In the first example, a declarative instruction is compiled into a production rule. After the rule is learned, the model no longer has to retrieve the instruction but can act at once at the appropriate moment. The second example shows -how previous experience is compiled into a rule: in this particular case, that a DC10 can be landed on a wet runway. The model

Niels Taatgen
Carnegie Mellon & U. Groningen

Continued on p. 10

References
1. N. A. Taatgen and J. R. Anderson, *Why do children learn to say "broke", a model of learning the past tense without feedback,* **Cognition 86** (2), pp. 123-155, 2002.
2. J. R. Anderson, D. Bothell, M. D. Byrne, S. Douglass, C. Lebiere and Y. Qin, *An integrated theory of the mind,* **Psychological Review**, *in press.*
3. N. A. Taatgen and F. J. Lee, *Production Compilation: A simple mechanism to model Complex Skill Acquisition,* **Human Factors 45** (1), pp. 61-76, 2003.
4. P. L. Ackerman, *Determinants of individual differences during skill acquisition: Cognitive abilities and information processing,* **J. Experimental Psychology: General 117**, pp. 288-318, 1988.

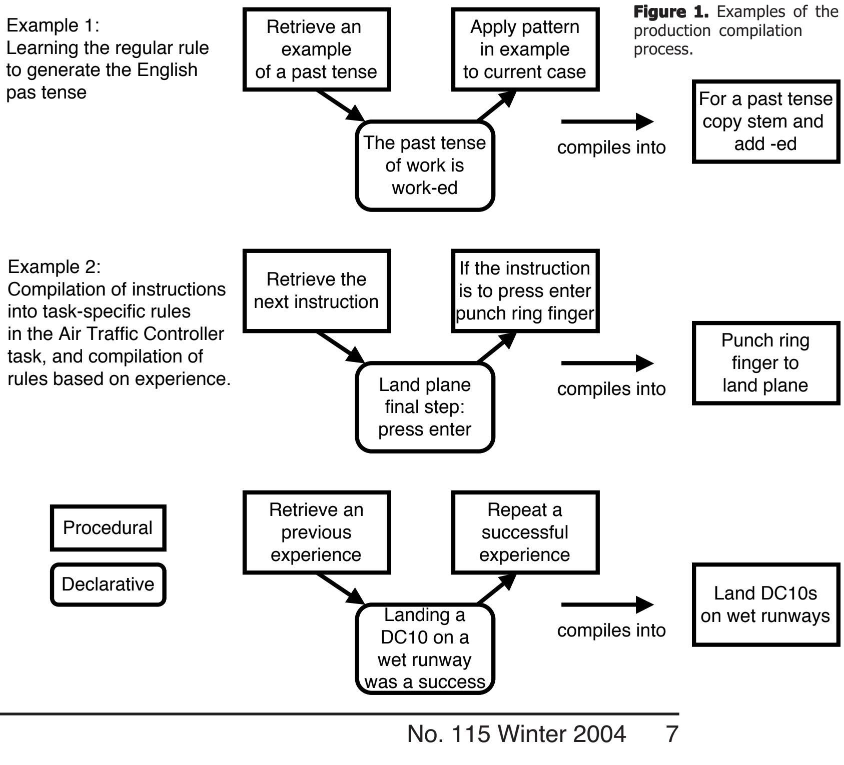

Figure 1. Examples of the production compilation process.

Virtual ICT with empathic characters

VICTEC (Virtual Information and Communication Technologies with Empathic Characters)[1] is a project the goal of which is to investigate the technologies required to create empathy between a human user and a synthetic character running in a 3D, interactive, graphical world. The intention is to produce a system to help with anti-bullying education for children aged 8-12—and by extension, enhance other areas of personal and social education—by building empathy between a child user and a synthetic character in a virtual drama. Scientifically, the project is expected to contribute to three technical areas: to theories of interactive narrative in virtual environments, known as emergent narrative; to an understanding of the role of empathy in creating social immersion; and to the evaluation of virtual-environment ICT for child users. On the psychological side, it hopes to contribute to a deeper understanding of empathy and its role in bullying, and to the relationship between theory of mind and bullying behaviour. The building of empathy between child and character is seen as a way of creating a novel educational experience.

The emergence of research on synthetic characters has seen a growth in possible applications and implications: such characters are now becoming widespread as a way to establish communication between users and computers. One of the main challenges for the VICTEC project is to design and build empathic synthetic characters: i.e., characters that, by their appearance, situation, and behaviour, are able to trigger empathic relations with the user and also able to show empathy (or not) for other characters. In order to build these, the VICTEC team has conducted and published research on agent requirements such as personality, believability, and empathy properties. Technically, the VICTEC empathic agent is one that is able to perceive and internally represents other agents' emotions and/or experience an appropriate emotion as a consequence. Further, by its behaviour and features, it can allow the users to build an empathic relationship with it.

A premise of this project is that the creation of empathy requires the user to really feel and to be emotionally affected in a way quite different from the indestructible or infinitely regenerating protagonists of most computer games. Because it is essential that the character with which the user is to feel empathy shows coherence and consistency, the VICTEC project suggests a need for unique narratives: that is, narratives with different characters and events for different users, rather than scripted stories that repeat identically. The mechanism being investigated for continuing, but different, stories is *emergent narrative*—generated by interaction between characters in the style of improvisational drama—rather than the authored tales in more widespread use. Narrative is generally seen as an artefact that can be studied, not as the dynamic process resulting from the interaction between characters and its impact on the user: the 'storification' process. This view of narrative-as-artefact is difficult to apply to the VICTEC project in which many similar-but-unique narratives are required rather than one pre-scripted one.

We have produced and published research and technical material on the project's web site. Amongst this material are VICTEC trailers and videos that present and illustrate the VICTEC demonstrator and give an overview of the final product. In addition, the VICTEC technical products and psychological theoretical models will be put into practice, formally tested, and evaluated. This will happen in June 2004 during a 10-day evaluation period involving more than 400 children at the University of Hertfordshire.

Sandy Louchart
Centre for Virtual Environment
University of Salford
E-mail: S.Louchart@Salford.ac.uk
http://www.nicve.salford.ac.uk/sandy

Reference

1. VICTEC is a Framework V project, part of the EU e-learning actions programme. It involves five partners in three different countries: Germany, Portugal and the UK. More information can be found at: http://www.VICTEC.org

---

Grounding language in sensorimotor and cognitive categories

Continued from page 5.

future research that looks at the phylogenetic and ontogenetic origins of language and cognition.

Angelo Cangelosi
Adaptive Behaviour and Cognition Research Group,
University of Plymouth
E-mail: angelo.cangelosi@plymouth.ac.uk
http://www.tech.plym.ac.uk/soc/staff/angelo/

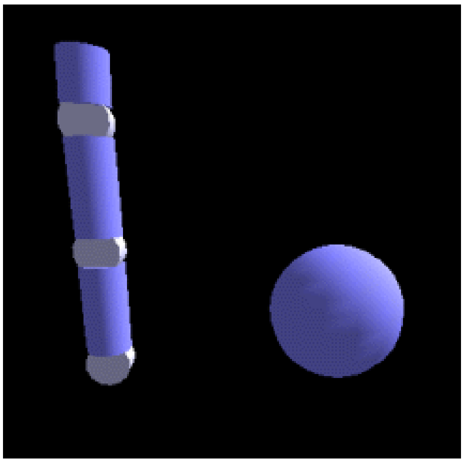

Figure 2. Experimental set-up for studying the emergence of grounded language in embodied evolutionary robots.

Book Review

Computational Vision: Information Processing in Perception and Visual Behavior

Hanspeter A. Mallot, John S. Allen (translator)

Publisher: MIT Press, http://mitpress.mit.edu
Hardback: published January 2001, 296 pp, £34.50.

Let me put the conclusion first: your library should certainly have this book, and if you are interested in vision you should probably have it on your shelf. I write that at the outset because what follows will sound critical, but you should not be put off. This is an ambitious and bold attempt to set aside the distinction between computer vision and biological vision, encompassing them both under the heading 'computational'. However, in moving towards this worthwhile goal, the book made me constantly aware of another, much better (and much longer) book which is still waiting to be written, perhaps by the same author. Nevertheless, in this difficult field, partial success is better than none: Mallot's synopsis has much to offer both postgraduate students and established researchers.

The introductory section starts strongly with a discussion of the perception-action cycle, and the information in images. How disappointing, then, that at the end of Chapter 1 vision is defined as 'inverse optics': the convenient but outdated view that the goal of any visual system is to build an internal representation of three-dimensional scene layout. It is as if the author, having written with conviction of a behaviour-oriented approach, was not himself convinced.

Having thus adopted a relatively conventional stance, the book takes us through a substantial selection of core ideas in visual information processing: imaging and projective geometry, convolution, edge detection, colour representation, stereo, shape from shading and texture, motion detection, and optic flow analysis. In general these are clearly presented with good diagrams and examples, and I found the colour chapter particularly informative. It is a pity, though, that important generic computational concepts are relegated to small examples: Bayesian methods tucked away in the estimation of slant from edge orientations; differential geometry surfacing briefly in shape from shading. There is a sense of truncation to many sections and chapters: abrupt endings that leave many questions unanswered.

The mathematics is mostly clearly presented and at an appropriate level, although I am not sure that it was worth giving space to mathematics that can be assumed, such as proofs of the properties of polynomials. A small but telling point: I was pleased to see the fact that sinusoids are the eigenfunctions of convolution given prominence; this is a key insight that is not often made explicit.

The final chapter, on visual navigation, is lively and has a sense of purpose that is not always present in the rest of the book, though it too ends as if guillotined. The bibliography, with about 300 references, is excellent, and reflects what is perhaps the book's main strength: the linking of ideas from psychophysics, neurobiology, ecological psychology and machine vision. I learned from this book and recommend it, but the sense of a missed opportunity remains.

David Young
Senior Lecturer in Artificial Intelligence
University of Sussex, UK

David Young previously studied physics at Cambridge and worked on ecological approaches to visual perception in Edinburgh University's Psychology Department.

Towards computational models of motor development

Continued from page 6.

entrainment between system, neural structures, and environmental feedback.[4]

In a follow-up study,[5] we examined the robustness of the freezing/freeing process by increasing the task complexity through the addition of a strong nonlinear coupling between system and environment. It was found that a single phase of freezing followed by a freeing of all DOFs, as suggested by Bernstein, was not sufficient. Instead, alternate phases of freezing and freeing of degrees of freedom were necessary to achieve robust behaviour. This result suggests that freezing and freeing may not necessarily be an endogenous property of the self-organizing neural system, but rather a process controlled by a *collective variable:* i.e., a measure of the emerging behaviour's dimension. To identify that variable, we are now studying a more complex infant task: bouncing.

Luc Berthouze
Neuroscience Research Institute (AIST), Japan
E-mail: Luc.Berthouze@aist.go.jp
http://staff.aist.go.jp/luc.berthouze

References
Continued from page 6.

2. G. Schöner and J. A. S. Kelso, *Dynamic pattern generation in behavioural and neural systems,* **Science 239,** pp. 1513-1519, 1988.
3. N. Bernstein, **The Coordination and Regulation of Movements,** Pergamon Press, Oxford, 1967.
4. M. Lungarella and L. Berthouze, *On the interplay between morphological, neural and environmental dynamics: A robotic case-study,* **Adaptive Behaviour 10** (3/4), pp. 223-241, 2002.
5. L. Berthouze and M. Lungarella, *Motor skill acquisition under environmental perturbations: On the necessity of alternate freezing and freeing of degrees of freedom,* **Adaptive Behaviour,** *in press.*

Join our committee!

Nominations are currently being sought for committee members. If you would like to help us further the aims of the Society, or know someone who would be good, you can find nomination details at:

http://www.aisb.org.uk/

Conference Review

Creative Systems: Approaches to Creativity in AI and Cognitive Science

Part of the IJCAI'03, Acapulco, Mexico, 2003

One of the most common criticisms of artificial intelligence applied to problem solving deals with its inability to deal with situations not predicted in the specification. Search mechanisms and solution spaces are normally very much constrained and predefined: however flexible, complex, and adaptable the system appears to be. Traditionally—when faced with a problem with no satisfactory solution in its search space—an AI system, at best, returns the least unsuccessful result. This is true even when the solution is incredibly close and only requires a minor change of perspective, relaxation of a constraint, or addition of a new symbol. In other words, such systems cannot perform what we normally call *creative behaviour*, a fundamental aspect of intelligence.

Thus, the question arises as to what can be done to make computers more creative: or even if, with the current computational architectures, it is possible at all. To some extent, current state-of-the-art paradigms (such as evolutionary computation, multi-agent systems, or case-based reasoning) have been responsible for much of the drive and developments regarding the first part of the question. The second half primarily concerns what the essential components of a creativity model could be, and whether these can exist in a formal machine. Both issues were hotly discussed during the Third Workshop on Creative Systems, part of the International Joint Conferences on Artificial Intelligence (IJCAI-03) this year.

The streams were organised as follows: *theoretical foundations*, *models of creativity*, *metaphor and analogy,* and *creative software*. Interestingly, the paper presentations and their specific discussions were more interesting than the plenary discussions held after each session. I see two distinct reasons for this. First, since the work is now starting to mature, experimental results and conclusions could be presented, giving us more to learn from. Second, some of the topics for plenary discussions have not evolved much since the previous events, perhaps because the area is reaching a point of leaving open the question, "What is creativity?" and progressed to, "What should computational creativity be?" A similar thing happened, in the early days, with the concept of intelligence within AI. In this workshop, the only relevant contribution for theoretical foundations of the area came from Geraint Wiggins[1] and there was general agreement that more development needs to take place to structure this area.

Among the researchers that have already reached some level of success in providing a more creative behaviour to the computer is Ricardo Sosa,[2] who applies multi-agent systems to modelling change phenomena in design. Another is Paulo Gomes, who applies case-based reasoning and analogy techniques for software reuse[3]. However, in my opinion, the most important thing is to synthesize the set of methods followed in modelling computational creativity by looking at the broad sweep of papers and discussions.

For instance, there are the internal aspects of the system: it should have a heterogeneous, multi-domain knowledge base;[4] it should be able to get knowledge from outside stimuli;[1,5] it should have methods for establishing analogies;[2] it should seek change[1] and surprise;[8] and it should have ways of generating metaphors and metonymies,[6] i.e. be able to integrate semantically-distant knowledge into useful concepts,[4] as well as being able to perform other sorts of classical reasoning such as deduction, induction, abduction, etc..[7] Externally, social interactions must be taken into account such that the creativity emerges within the situation the system finds itself in, not only as determined by an agent.[1,5] Clearly, this is just one take: a creativity model must obviously consider many other issues as well.

Extended versions of these papers will soon be published in a special issue of the Journal of Knowledge-Based Systems.

Carlos Bento
Co-Chair, 3rd Workshop on Creative Systems
Center for Informatics and Systems (CISUC)
University of Coimbra, Portugal
E-mail: bento@dei.uc.pt

References
All of the following are in **Proc. 3rd. Workshop on Creative Systems: Approaches to Creativity in Artificial Intelligence and Cognitive Science**, part of IJCAI 2003.
1. Geraint Wiggins, *Categorising Creative Systems.*
2. Ricardo Sosa and John Gero, *Design and Change: A Model of Situated Creativity.*
3. Paulo Gomes et al., *The Importance of Retrieval in Creative Design Analogies.*
4. Francisco C. Pereira. *Experiments with Free Concept Generation in Divago.*
5. Fernando Machado et al, *Power to the Critics—A Framework for the Development of Artificial Critics.*
6. Tony Veale, *Qualia Extraction and Creative Metaphor in WordNet*
7. Simon Colton,n*Creative Logic Programming.*
8. LuÌs Macedo, *A Model for Generating Expectations: the Bridge Between Memory and Surprise.*

Production compilation: a versatile cognitive mechanism

Continued from page 7.

is able to fit human data across learning sessions at a general level (number of planes landed), at the unit task level (time to complete certain sub-tasks), and at the level of individual keystrokes.

Niels Taatgen
Carnegie Mellon University,
Pittsburgh, USA and
University of Groningen, Netherlands
E-mail: taatgen@cmu.edu
http://www.ai.rug.nl/~niels

Membership report

AISB is one of the three most influential UK Artificial Intelligence societies. We provide an information-packed quarterly newsletter covering AI news and research, plus a twice-yearly journal that covers more in-depth contemporary articles.

To pay for this, our membership fees will rise to £35 for full members from 1st January 2004: a price we feel still represents excellent value for money. Note that fees have only been increased once in the last seven years! To encourage up-and-coming AI researchers, students can still join for £15: particularly reasonable since we give two £300 travel bursaries to students each year.

If you know anyone you think may like to join, why not direct them to our web site below where they can see what the Society offers: among other things, there are back issues of the Quarterly available for download.

Gary Jones
SSAISB Membership Secretary
http://www.aisb.org.uk/

Editor's note: I had to force Gary to submit a truncated report this issue because we had so much other material to squeeze in. However, he will be submitting another installment in the next issue.

AISB'04 Convention: *Motion, Emotion and Cognition*

We are pleased to announce that the AISB'04 Convention will be hosted by the University of Leeds, United Kingdom, from 29 March to 1 April 2004. There will be six symposia at the convention, covering the following topics: *Adaptive Agents and Multi-Agent Systems* (AAMAS-4), chaired by Daniel Kudenko, University of York; *AI in Virtual Reality,* chaired by David Corne, University of Reading; *Emotion, Cognition, and Affective Computing,* chaired by Colin G. Johnson, University of Kent at Canterbury; *Gesture Interfaces for Performing Arts*, chaired by Kia Ng, Nicola Bernardini, and Antonio Camurri; *Immune System and Cognition,* chaired by Simon Garrett, University of Aberystwyth; and *Language, Speech and Gesture for Expressive Characters,* chaired by Ruth Aylett, Marc Cavazza, and Patrick Olivier. The *11th Automated Reasoning Workshop,* chaired by Brandon Bennett, University of Leeds, will be co-located.

The success of the Convention depends on your support. Please consider submitting a paper or coming along: full details are available via the website, details below. The AISB Convention has always been a dynamic event for the reporting of the latest developments, results and concepts, and to foster a multi-disciplinary exchange of ideas. We look forward to welcoming you to an exciting and fruitful AISB convention in 2004.

Kia Ng
University of Leeds
Convention Chair, AISB'04
E-mail: aisb04@aisb.org.uk
http://www.leeds.ac.uk/aisb

Submit your paper to the AISBJ!

Submitted papers should present original and substantial research work in areas of interest concerning artificial intelligence, cognitive science, simulation of behaviour, and any related fields. Interdisciplinary submissions are particularly welcome. Submissions will be anonymously reviewed by at least two researchers working in the relevant field or fields. For full instructions, please see our website:

http://www.aisb.org.uk/aisbj

Expert Systems journal discount for AISB members

Blackwell Publishing has set up a new discount rate of £20/$32 for AISB members interested in their journal *Expert Systems*. The journal is ISI ranked and devoted to all aspects of artificial intelligence and advanced computing. The journal's readers include knowledge engineers, artificial intelligence researchers, computer scientists, and project managers. It is written for those looking for an international perspective on expert systems and neural networks, whether as developer, supplier, or potential user. There is also an annual special issue: in 2003 it was entitled *Rough Sets: Current & Future Developments*.

Abstracts are available to browse online.[1] To subscribe using the discount, please visit the publication website below, or contact the publishers at the e-mail address listed below.

http://www.blackwellpublishing.com/exsy
E-mail: customerservices@oxon.blackwellpublish

Reference

[1] http://www.blackwell-synergy.com/links/toc/exsying.com.

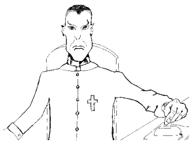

Father Hacker's Guide

for the Young AI Researcher

Cognitive Divinity Programme
Institute of Applied Epistemology

AI research is labour intensive. To make an impact, you need a devoted work force to implement your groundbreaking ideas. To orchestrate raw postgraduate students and postdoctoral researchers into the harmonious, world-class team that you need, it is essential to know....

9. How to be a Research Leader

1. To ensure your group gets its due recognition, it is vital that you maintain high visibility, with regular and frequent media appearances, keynote addresses, international research trips, and expert testament to decision-making bodies. Modern telecommunications mean that is no longer necessary to be physically present in order to stay abreast of the latest developments in your lab. Discuss recent advances with your deputies on the phone: have any papers or slides prepared by your staff and emailed to your 5* hotel just before your presentation, then fly directly to your next important engagement.

2. Naturally, work as important as yours requires a large team. But your vital work will keep you away from your lab for long periods and, when you are present, you will usually be busy in important meetings. You will not have the time to devote to long, face-to-face supervisory meetings with large numbers of junior researchers. Appoint deputies to handle these minor details.

Give your students 24/7, personal, expert supervision...

...with Hacker's *NURSEMAID™* (Necessary Understanding of Research Supervisory Expertise for Making Artificial Intelligence Discoveries).

---

3. Always remember that it is *your* research team. Even work on which you made no explicit contribution was originally inspired by your vision. So you deserve the major credit for all work produced by the group. For instance, your name should appear as principal author of all papers and software, you should deliver all prestigious presentations, and you should collect all awards or prizes. Your team members will benefit from the prestige that adding your name will bring to their work, and so will be grateful for your endorsement.

4. It is also impossible for you to find the time needed to keep up with the latest research results. It is, in any case, improbable that significant advances will occur outside your group. But to err on the safe side, your junior researchers should be instructed to scan the literature, incorporating anything of note into the output of your group, so that you get the credit you deserve for identifying its potential, describing it in your terms, and deploying it effectively.

Take the Royal Road to expert status...

...with Hacker's *BLUFF™* (Bypass Literature Understanding by Faking Fluency). Sprinkle your talks with automatically generated and personalised phrases and slogans—which bear witness to your familiarity with important advances—without the drudgery of unnecessary study.

---

5. The more trailblazing your work, the longer it takes for the international community to reward it with esteem indicators. If journal editorships, programme chairships, keynote talks, distinguished fellowships, and academic society presidencies are not yet flooding in, then speed things along. Set up your own journal, conference series and academic society. Appoint your close colleagues and collaborators to the key posts and place yourself at their head. Reward your research team with the inevitable administrative and managerial tasks that such selfless service to the community inevitably entails. You will soon get the recognition you deserve.

About the Society
The Society for the Study of Artificial Intelligence and Simulation of Behaviour (AISB) is the UK's largest and foremost Artificial Intelligence society. It is also one of the oldest establised such organisations in the world.

The Society has an international membership of hundreds drawn from academia and industry. Membership of AISB is open to anyone with interests in artificial intelligence and cognitive and computing sciences.

AISB membership includes the following benefits:

- Quarterly newsletter
- Biannual Journal
- Travel grants to attend conferences
- Discounted rates at AISB events and conventions
- Discounted rates on various publications
- A weekly e-mail bulletin and web search engine for AI-related events and opportunities

You can join the AISB online via:
http://www.aisb.org.uk

Next AISBQ deadline: 6 February 2004

AISBQ reaches hundreds of Artificial Intelligence and Cognitive Science researchers in the UK, Europe, and beyond. Advertising and other information is available via:

http://www.aisb.org.uk/qhome.html

This page includes full guidelines for the submission of book/conference reviews and technical articles. Books available for AISB members to review are also listed.