AISB Quarterly

No.114

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

In this issue:

Features

Reviews

Father Hacker

Family resemblance, Bayesian networks and exemplars

Much of our daily reasoning appears to be based on stereotypes, exemplars, and anecdotes. Yet, basic statistics informs us that decisions based only on limited data are, at best, likely to be inaccurate, if not badly wrong. However, exemplars and stereotypes are not arbitrary data points, they are chosen based on experience and represent prototypical situations. The ability to predict the behaviour of a consumer, observe that two people are related, diagnose an illness, and even how an MP might vote on a particular issue, all depend on a person's past experience—that is the exemplars and stereotypes a person learns. If this hypothesis—namely that we can form and reason well with exemplars—is true, we should be able to identify exemplars from data. To achieve this, we need to answer the following questions. What is an exemplar and how can it be represented? How do we learn good exemplars incrementally? How can exemplars be used?

Here we outline a particular approach to these questions that involves the use of the notion of family resemblance to learn exemplars and Bayesian networks to represent and exploit them.

The central problem can be visualized as moving from a situation like that in Figure 1—which has three categories, A, B, and C, with a lot of data—to one like that in Figure 2 where we have exemplars representing the categories. Given that both membership of categories and the extent to which exemplars represent other points are graded, we use Bayesian networks for this task. Figure 3 shows the representation used, where *ei* denotes exemplars associated with the categories, *fi* denotes features, and the arcs denote dependencies. The node *Ve* is a virtual exemplar, introduced to take account of all the data points that have not been seen as the model learns incrementally. It is needed to satisfy the conditions of the particular kind of network that we use, called the noisy-OR model,[1] which enables us to adopt a more efficient propagation algorithm. Given such a model, we assess whether a new point is represented by an exemplar using a propagation algorithm to compute its probability given the point's features. The exemplar with the highest probability can then be used to determine the point's category.

How do we learn such a model incrementally? As an initial experiment, a simple, greedy learning strategy is adopted: if a new training case is not represented by an existing exemplar, then it is added as a new one, otherwise the exemplar that represents the training point and the new training point compete, with the better one being retained. To assess which exemplar is better, we adopt Rosch and Mervis' view[2] that a good exemplar is one that has high family resemblance with those it represents (focality) and low resemblance with those considered outside the family (peripherality). Given our representation, we can interpret family resemblance as the probability of an exemplar representing a point, which in turn can be used to compute focality and peripherality. The difference between the the two can then be used as a measure of the prototypicality of the exemplar.

The model has been implemented and tested on the *animals*, *votes*, and *audiology* data sets available from the University of California at Irvine Machine Learning Repository. The experiments used a 70/30 training/testing split and 20 random trials were performed. The *votes* data set records the voting behaviour of USA congressmen on 16 issues and their party affiliation, classified into Democrats or Republicans. On average, two exemplars were retained to achieve an accuracy of 96% for the Republican category, and four exemplars were retained to achieve an accuracy of 84% for the Democrats. The results were similar for the *animals* data set, with two exemplars at most in any of the seven categories of animals, and an overall accuracy of 92%. The results on the *audiology* data set are much more varied, but close to those achieved by Bareiss (1989) when the exemplars were hand crafted with the aid of experts for his PROTOS system.[3] A general characteristic of the results is that forcing an increase in the number of exemplars retained reduces the accuracy,

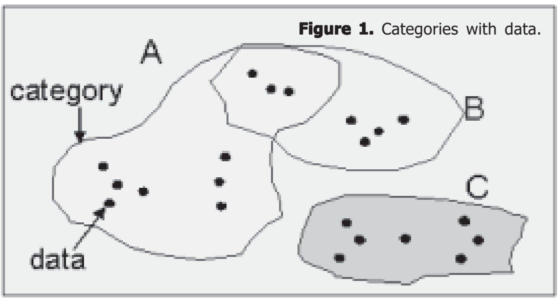

Figure 1. Categories with data.

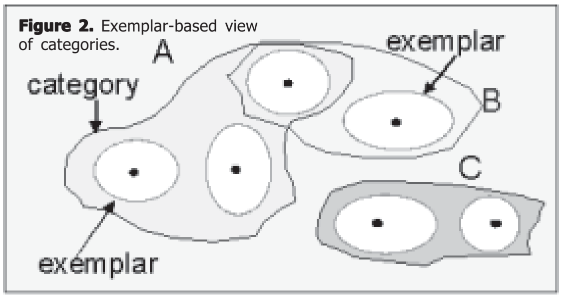

Figure 2. Exemplar-based view of categories.

Sunil Vadera
University of Salford
Continued on p. 11

References
1. J. Pearl, **Probabilistic reasoning in intelligent systems: networks of plausible inference**, Morgan Kaufmann, 1991.
2. E. Rosch and C.B. Mervis, *Family resemblance studies in the internal structure of categories,* **Cognitive Psychology 7,** pp 573-605, 1975.
3. R. Bareiss, **Exemplar-based knowledge acquisition: A unified approach to concept representation, classification, and learning,** Academic Press Inc., 1989.

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

ISSN 0268-4179

Editor/Managing Editor:
Sunny Bains
aisbq@aisb.org
http://www.aisb.org.uk
and click on *AISBQ*.

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

AISB PATRON
John Barnden
University of Birmingham

INTELLIGENT INTERACTION

The IndagoSonus composition system

The IndagoSonus project is about using blocks that are both similar in size and shape to children's wooden building blocks, and in their ability to be combined together to make physical structures. However, here, each block has the ability to play *and compose* music, so building a physical structure results in creating a piece of music.

Imagine you have a single block. You give it a small musical fragment (its 'home' music), put it on the table, and start it off playing. The block then begins repeating its musical phrase, either sporadically or continuously. Make another block, this time with a different piece of 'home' music. Place it next to the first block and they start playing together, and in synch. Now press a button on block 1 labelled 'send music.' This causes block 1 to send its music to block 2. Block 2 then performs a composition activity. Each block contains a genetic algorithm to achieve this, which evolves musical solutions based on the block's 'home' music, the music it is currently playing, and the music it has just been passed. The compositional aim for the block is to produce a *new* musical section that has a thematic relationship with both its home music and the music it has received. Block 2 then starts playing its new music.

Now imagine a chain or group of several blocks in any 3D structure. All blocks have a 'send music' button, so the start of the chain does not have to be block 1. If a block is passed some music, it recomposes itself then passes its new music on to *all* of its neighbours. By sending out music from a starting point, all other blocks within a specified range recompose, and the collective music of the structure is transformed. It is important to clarify that each block holds onto its 'home' music throughout, enabling any music composed by it to remain thematically related, despite the constant process of re-composition each block undertakes. In this way the composer of the home music for all blocks maintains a compositional thumb-print on the evolving musical structure.

The development of the overall piece of music will therefore be determined by: the design of the composition system described here; the nature of the music imported as 'home' music into each block; how the blocks are built into structures; how these structures change over time; and how the user sends music around the structure. In effect, the listener/performer is able to *shape* the overall music by choosing to send musical fragments from blocks they like to *influence* other blocks. One part of the structure may have composed some ideas the user likes. A block from that group could then be placed in another part of the structure to see what effect it has. The whole system acts as a kind of genetic algorithm 'patch bay,' allowing users to direct the evolution down a musical path without having to make judgments after every evolved population.

IndagoSonus is a development from my earlier work in interactive music, and sees the human composer's role as being the process of defining a search space for exploration by the user. This is an extension of the notion of the act of composition itself being an exploration of the potentially very large search space permitted by the compositional practice of the composer, even considering limitations of instruments, styles, genres etc.. I now regard this exploration as a two-stage process. Firstly, the composer limits the larger search space by defining an area within it. Secondly, the performer or listener explores this more limited space and defines his or her own aural realization.

IndagoSonus is therefore a composing system, rather than playback system for pre-composed fragments or loops. It has a simple interface designed to encourage depth of involvement, and to provide the user with a high degree of creative influence on the actual realization, whilst at the same time enabling a composer's thumbprint to be present

Currently the project exists as a working 3D graphical prototype, see Figure 1, with plans to develop the hardware model during 2003/4.

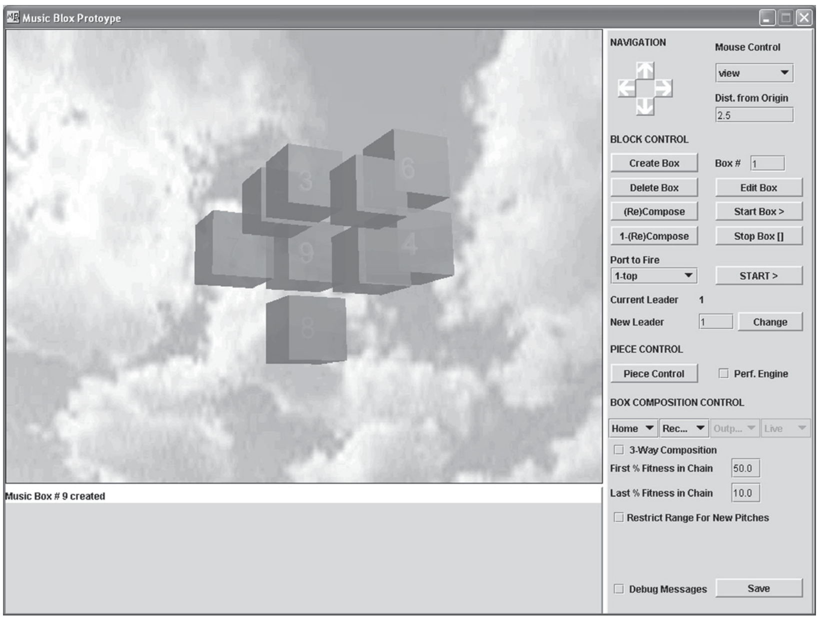

Figure 1. This screenshot shows the current, functional 3Dgraphical prototype.

Andrew Gartland-Jones
Computational Creativity Research Group
The University of Sussex
Brighton, UK
E-mail: drew@atgj.org
http://www.atgj.org/drew

Informal human mathematical reasoning

We encounter mathematics in every aspect of our lives. Some of the deepest and greatest insights into reasoning were made in mathematics, so it is not surprising that emulating such powerful reasoning on machines is one of the important and difficult aims of artificial intelligence. Human mathematicians have many problem-solving strategies. They use diagrams to better convey problems and generate intuitive and easily-understandable solutions. They also learn from examples of solutions to related problems, and exploit analogy and symmetry. My research explores the nature of such informal reasoning.

Though informal human reasoning is very powerful, its potential has largely not been exploited in the design of mechanised reasoning systems: systems that use some logic formalism to (semi-) automatically solve problems. This can perhaps be explained by the fact that we do not have a deep understanding of informal techniques and their use in problem solving. To advance the state of the art in automated reasoning systems, I think it important to integrate some of the informal human-reasoning techniques with the proven, successful, formal techniques, such as different types of logic. This will not only make reasoning systems more powerful, but will allow them to serve as tools with which we can study the nature of human reasoning. My aim is to formalise and emulate two things on machines: human reasoning with diagrams and human learning.

In automated systems, theorems are usually proved with formal, logical—so-called symbolic—proofs. However, there is a subset of problems that humans can prove by the use of geometric operations on diagrams, so-called diagrammatic proofs. Figure 1 presents an example of such a proof[1] concerning the sum of odd naturals $n^2=1+3+5+...+(2n-1)$.

The proof consists of repeatedly applying *Lcut*s to a square (an *Lcut* removes an *ell* shape which is formed from two adjacent sides of a square—see Figure 1). Notice that an *ell* represents an odd natural number since both sides of a square of size *n* are joined (*2n*), but the joining vertex was counted twice (hence *2n-1*). We showed how such diagrammatic reasoning about mathematical theorems can be automated, and demonstrated the approach with the diagrammatic reasoning system called Diamond.[2]

In Diamond, concrete, rather than general, diagrams are used to prove particular instances of a universal statement: e.g., in the example in Figure 1, the instance is $n=6$. The "inference steps" of a diagrammatic proof are formulated in terms of geometric operations on the diagram: e.g., the *Lcut*s in the diagrammatic proof in Figure 1. A general schematic proof of the universal statement is induced from these proof instances by means of the constructive omega-rule. Schematic proofs are represented as recursive programs which, given a particular diagram, return the proof for that diagram. It is necessary to reason about this recursive program to show that it outputs a correct proof. One method of confirming that the abstraction of the schematic proof from the proof instances is sound is proving the correctness of schematic proofs in the meta-theory of diagrams.

Diamond can only tackle theorems that can be expressed as diagrams. However, there are those that may require a combination of symbolic and diagrammatic reasoning steps: so-called heterogeneous proofs. I am currently investigating how a system could automatically reason about these, and learn them in general from examples. An example below demonstrates a heterogeneous proof that consists of a combination of symbolic and diagrammatic inference steps. The theorem states an inequality: $(a+b)/2 \geq \sqrt{(ab)}$ where $a,b \geq 0$. The first few symbolic steps of the proof are:

$$
\begin{array}{lcl}
(a+b)/2 & \geq & \sqrt{(ab)} \\
 & \downarrow & \text{square both sides of } \geq \\
(a+b)^2/2^2 & \geq & ab \\
 & \downarrow & \times 4 \text{ on both sides of } \geq \\
(a+b)^2 & \geq & 4ab \\
 & \downarrow & \\
a^2+2ab+b^2 & \geq & 4ab
\end{array}
$$

The second part of the proof,[1] which is presented in Figure 2, shows diagrammatically the inequality $a^2+2ab+b^2 \geq 4ab$.

Rather than learning low-level proofs, I aim for a system that can learn diagrammatic *proof plans*. *Proof planning*[3] is an approach to theorem proving that uses high-level *proof methods* rather than low-level logical inference rules to find a proof of a conjecture at hand. A proof plan consists of some combination of proof methods, which in turn specify and encode a general-reasoning strategy that can be used in a proof. It typically represents a number of individual inference rules: e.g., mathematical induction can be represented as a proof method. Our heterogeneous proof plans will be formed from geometric operations plus symbolic inference steps. The system will be able to learn such proof methods from examples of the use of lower-level methods. Eventually, it will be able to learn new diagrammatic and heterogeneous proof plans.

The hope is that, ultimately, learning new, general, and complex proof methods and plans may lead to the discovery of new and interesting proofs that use diagrams for inferencing.

Mateja Jamnik
Computer Laboratory, University of Cambridge, UK
http://www.cl.cam.ac.uk/users/mj201

Figure 1. Diagrammatic proof for: $n^2=1+3+5+...+(2n-1)$.

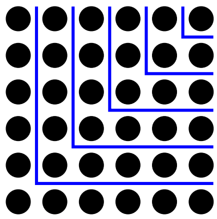

References

1. R.B. Nelsen, *Proofs without Words: Exercises in VisualThinking*, **The Mathematical Association of America,** 1993.
2. M. Jamnik, *Mathematical Reasoning with Diagrams: From Intuition to Automation*, **CSLI Press, Stanford,** 2001.
3. A. Bundy, *The use of explicit plans to guide inductive proofs*, in R. Lusk and R. Overbeek, editors, **Lecture Notes in Computer Science 310,** pp. 111-120, Springer Verlag, 1988 (9th Conference on Automated Deduction).

Figure 2. Diagrammatic proof for $a^2+2ab+b^2 \geq 4ab$.

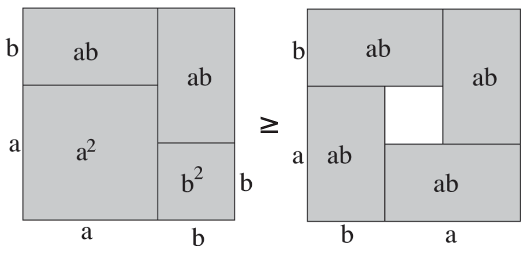

Modelling assertion evaluation

An ultimate goal of AI is the manufacture of an android that can interact meaningfully with human beings. Progress has already been made towards achieving this. Androids exist that can walk, climb stairs, recognize and grasp objects, imitate human behaviour, and so on. In addition to these abilities, however, an android would also need the ability to acquire knowledge about its surroundings: not only its physical and social environment, but also its intellectual environment. Most research on knowledge acquisition focuses on the ability to acquire beliefs through perception.[1] This is important, but equally important is the neglected topic of acquiring knowledge by believing what other people say and what they have written. Usually, when I discuss my research, people cannot believe that there is no general theory of how agents acquire knowledge by accepting other people's assertions.[2] In recent years, philosophers have started to take this topic more seriously,[3] but AI has not yet caught up with them. (In addition to my research, Paul Thagard is also working in this area.[4])

I propose a two-stage model of belief-acquisition, shown in Figure 1. An agent gains information in two main ways: by making judgements about its perceptual environment and by accepting some of the assertions other people make. My research focuses on the second of these. Most of an agent's knowledge has been acquired by accepting other people's assertions, but nobody believes everything he reads or hears. My proposal is that an agent's assessment of the assertions it encounters is governed by the defeasible (capable of being annulled or invalidated) rule to believe them. This sounds simplistic, but working out all the factors that may cause it to be overridden is a difficult task.[5-7]

In the model, faced with an assertion, an agent can either accept it—adding it to his belief-system—or reject it. This is a simplification of what actually happens, since we do not accept everything with the same degree of conviction. This is one of many simplifications I have made. Such an approach is justifiable, however, in order to gain a better understanding of how agents evaluate assertions. I intend to refine the model in the future as computer simulations are evaluated and analyzed.

The heart of the model is the assessment component. This unpacks the defeasible rule to believe other people's assertions. It consists of an ordered set of rules, all of which—except the last—are conditional. The last rule is the non-defeasible rule to believe the assertion in question. There are many reasons why someone may decide not to accept an encountered assertion, and these become the antecedents of the conditional rules. For example: a play is a work of fiction and so we do not normally believe the actors' assertions. This can be captured by adding, to the assessment component, the rule, 'If the assertion *X* is uttered by an actor during the performance of a stage play, then reject *X*'. There is provision in the model for the assessment-component rules to be altered in the light of experience.

So far I have been describing the first stage of belief-acquisition. Clearly, the judgements we make about our perceptual environment, and the evaluations we carry out concerning the assertions we encounter, have to be done in real time. Therefore, assessment cannot be very sophisticated. As a result of this, agents will acquire some false beliefs and reject some assertions that, as a matter of fact, are true. Consequently, the model contains a second stage of belief-acquisition in which a small number of an agent's beliefs are subjected to a thorough investigation of their truth or falsity. Here, an agent can re-evaluate something he believes or re-consider an assertion he previously rejected. This may involve substantial belief-revision.

This research is still in its infancy, but considerable progress has already been made in developing a general theory of how agents learn from others' assertions. Collaborators and doctoral students are welcome: as are competitors to devise rival theories.

Antoni Diller
School of Computer Science
University of Birmingham, UK
E-mail: A.R.Diller@cs.bham.ac.uk
http://www.cs.bham.ac.uk/~ard

Figure 1. A two-stage model of belief acquisition.

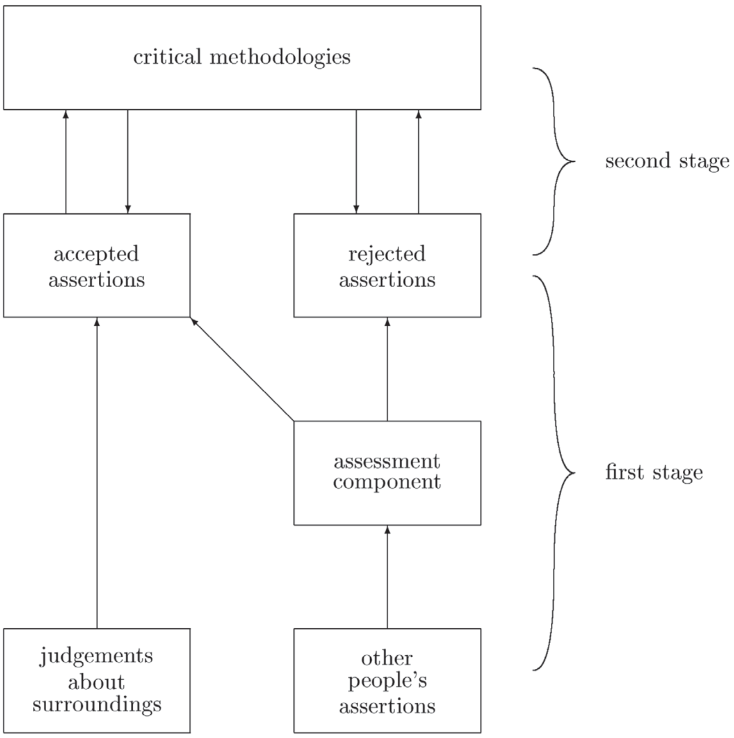

References
1. J. Pollock, **Cognitive Carpentry,** MIT Press, London, 1995.
2. A. Diller, *Designing androids,* **Philosophy Now 42,** pp. 10-13, July/August, 2003.
3. C.A.J. Coady, **Testimony,** OUP, Oxford, 1992.
4. http://cogsci.uwaterloo.ca/Biographies/pault.html
5. A. Diller, *Everyday belief-acquisition,* **Proc. Argentine Symp. on Artificial Intelligence (ASAI 2000),** SADIO, Buenos Aires, 2000.
6. A. Diller, *Evaluating information found in journal articles,* **Proc. Logic, Language and Information (JOLL 2000),** Kronos, Sevilla, 2000.
7. A. Diller, *Acquiring information from books,* **Proc. Res. and Dev. in Intelligent Systems XVII (ES 2000),** Springer, London, 2001.

GENETIC PROGRAMMING

Routine, human-competitive, high-return machine intelligence

Genetic programming is a systematic method for getting computers to automatically solve problems. The technique starts from a high-level statement of what needs to be done, and automatically creates a computer program to do it, by progressively breeding a population of programs using analogues of the naturally occurring operations: reproduction, crossover (sexual recombination), and mutation.[1-4] Genetic programming now delivers routine, high-return, human-competitive machine intelligence.

In fact, there are now 36 instances where this method has produced a human-competitive result. Of these, in 15 cases something was created that either infringed or duplicated the functionality of a previouslypatented 20th-century invention (a further six for 21st-century inventions), and twice a new, patentable invention emerged. One of the latter is a general-purpose controller that outperforms others by employing tuning rules that have been widespread in industry for most of the 20th century (see Figure 1). Such genetic-programming creations exhibit the same kind of creativity, logical discontinuity, and departure from established ways of thinking as are the essence of human invention.

We say that a result is 'human-competitive' if it satisfies at least one of eight criteria: for instance, that the machine-produced result is publishable in its own right as a new scientific result—independent of the fact that the result was mechanically created—or that the machine-produced result was patented in the past or would qualify today as a patentable invention.

A result is 'high-return' if it has a high AI (artificial-to-intelligence) ratio. We define this as the contribution of the automated operation of an artificial method over the intelligence pre-supplied by the human. Manifestly, the aim of the field of machine intelligence is to get computers to automatically generate human-competitive results with a high AI ratio, not to demonstrate that humans are capable of producing human-competitive results themselves. A method is 'routine' if it is general and relatively little human effort is required to get the method to successfully handle new problems or those from a different domain.

Many of the results produced by genetic programming are the result of the successful reuse of substructures. Complex structures are almost always replete with modularities, symmetries, and regularities. Reuse avoids reinventing the wheel on each occasion, and requires a particular sequence of already-learned steps. Genetic programming can reuse code by means of automatically-defined functions, iterations, loops, and recursions: it can also reuse the result of executing code by means of automatically-defined stores. Thus, it can dynamically determine—during the run—the number of reused structures, their type, and the nature of the hierarchical references among the substructures.

Results produced using genetic programming come mainly from fields such as circuits, controllers, antennas, networks of chemical reactions, metabolic pathways, genetic networks, and game-playing.

Summarizing work over the 15-year period between 1987 and 2002, genetic programming has delivered a progression of qualitatively-more-substantial results in synchrony with five approximately order-of-magnitude increases in the expenditure of computer time. The five now-identifiable groups of ever-better results include:

- solving toy problems of the 1980s and early 1990s from artificial intelligence and machine learning
- producing human-competitive results not involving patents
- duplicating or infringing 20th-century patents
- duplicating or infringing 21st-century patents
- creating patentable new inventions.

As far as we know, genetic programming is, at the present time, unique among methods of artificial intelligence and machine learning. This is because of its duplication of numerous previously patented results, its generation of patentable new results, the breadth and depth of the problems it can solve, its demonstrated ability to produce parameterized topologies, and its delivery of routine, high-return, human-competitive machine intelligence.

Looking forward, we believe that genetic programming will be increasingly used to automatically generate ever-more-complex, human-competitive results.

For further information about genetic programming, please consult the references listed.

John R. Koza
Consulting Professor
Biomedical Informatics, Dept. of Medicine
and Dept. of Electrical Engineering
Stanford University, CA, USA
E-Mail: koza@stanford.edu
http://www.smi.stanford.edu/people/koza

References
1. John R. Koza, **Genetic Programming: On the Programming of Computers by Means of Natural Selection,** MIT Press, 1992.
2. John R. Koza, **Genetic Programming II Videotape: The Next Generation,** MIT Press, 1994.
3. John R. Koza, Forrest H. Bennett III, David Andre, and Martin A. Keane, **Genetic Programming III: Darwinian Invention and Problem Solving,** Morgan Kaufmann,1999.
4. John R. Koza, Martin A. Keane, Matthew J. Streeter, William Mydlowec, Jessen Yu, and Guido Lanza, **Genetic Programming IV: Routine Human-Competitive Machine Intelligence,** Kluwer Academic Publishers, 2003.
5. http://www.genetic-programming.org
6. http://www.genetic-programming.com

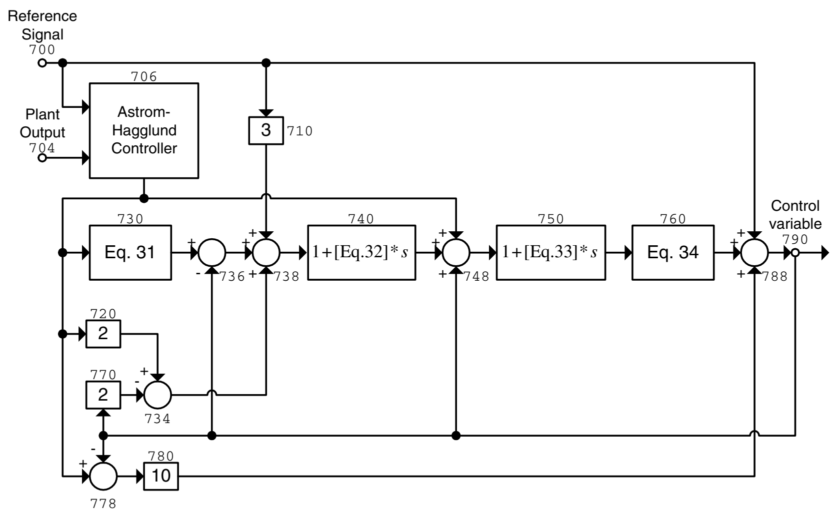

Figure 1. A new, patentable invention generated by genetic programming: a general-purpose controller.

Investigation of robot-environment interaction using chaos theory

Figure 1. The fundamental triangle of robot-environment interaction: a robot's behaviour always has to be seen in the context of robot, task and environment.

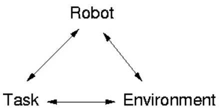

References

1. Ulrich Nehmzow and Keith Walker, *The Behaviour of a Mobile Robot Is Chaotic,* **AISB Journal 1** (4), 2003.
2. Ulrich Nehmzow, *Quantitative analysis of robot environment interaction—towards 'scientific mobile robotics',* **Int'l J. of Robotics and Autonomous Systems 44** (1), pp. 56-58, July 2003.
3. Ulrich Nehmzow and Keith Walker, *Quantitative Description of Robot-Environment Interaction Using Chaos Theory,* **Proc. European Conf. on Mobile Robotics (ECMR),** Warsaw, 2003.
4. Ulrich Nehmzow and Keith Walker, *Is the Behaviour of a Mobile Robot Chaotic?,* **Proc. AISB Convention,** Aberystwyth, April 2003.

In collaboration with Point Loma Nazarene University in San Diego, current research at the University of Essex investigates how dynamical systems and chaos theories can be used to quantify robot-environment interaction. Mobile robotics research to date is still largely reliant on trial-and-error procedures rather than exploiting established theories describing robot-environment interaction in a formal manner, making falsifiable predictions, and allowing quantitative descriptions of a robot's behaviour. Quantitative performance measures are the first step towards a theory of robot-environment interaction, as well as having practical applications to mobile robotics research. We are therefore interested in establishing such measures as standard within the field.

Quantitative measures of interaction

The behaviour of a mobile robot cannot be discussed in isolation: it is the result of properties of the robot itself (physical aspects, the 'embodiment'), the environment ('situatedness'), and the control program (the 'task') the robot is executing (see Figure 1). This triangle of robot, task and environment constitutes a nonlinear system, whose analysis is the purpose of any theory of robot-environment interaction.

If we assume, for argument's sake, that the behaviour resulting from the interaction of agent, task, and environment can be described quantitatively, a number of possibilities arise:

• Two of the three elements in Figure 1 are kept unchanged, and the third is modified in a systematic way. The quantitative performance measure then characterises that third component.

- For instance, to characterize two environments quantitatively, the same robot and control program can be used in either environment, and the quantitative description of behaviour used to identify (in the system-identification sense) the environments.

- Likewise, by systematic modification of just one of the three components shown in Figure 1, optimal parameter settings (with respect to some desired behaviour) can be determined in a systematic way.

• Experimental results can be stated quantitatively. This allows replication and verification of experimental results, which is currently hardly possible in mobile robotics research.

• Predictions made by the theory of robot-environment interaction can be made quantitatively, and tested against the actual experimental results.

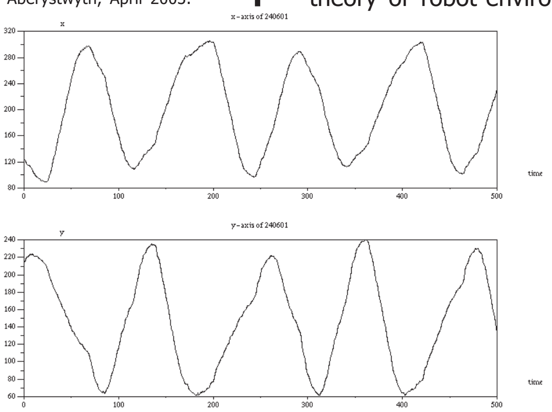

Figure 2. The *x* and *y* coordinates of part of the trajectory shown in Figure 3.

Figure 3. 'Billiard Ball' behaviour in a square arena—the entire trajectory (left) and 150 data points (right).

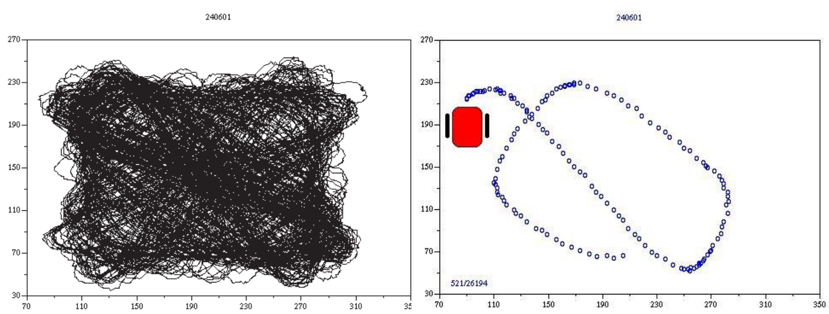

Is robot behaviour chaotic?

A mobile robot interacting with its environment is a nonlinear system. Any prediction of the robot's trajectory is therefore limited to short time horizons, and the actual trajectory will diverge. One way to investigate the behaviour of nonlinear systems is chaos theory, in particular quantitative descriptions of the system's phase space (dimensionality of attractor and sensitivity to initial conditions).

In our experiments, we used a Pioneer II mobile robot executing various control programs in various environments. The robot's trajectories were logged using an overhead camera system, resulting in paths such as those shown in Figures 2 and 3. We then analyzed the trajectories for the four defining characteristics of systems exhibiting deterministic chaos: stationarity, determinism, aperiodicity, and sensitivity to initial conditions. Please see the references for details.

Results

We conducted a wide range of experiments, modifying either the environment, or the task the robot was executing. As a result, both the dimension of the attractor, and the degree of sensitivity to initial conditions, were quantitative descriptions of the parameter changed.

Our findings were as follows. First, robot-environment interaction does exhibit deterministic chaos. Second, chaos theory can indeed be used to quantify robot-environment interaction. Finally, different aspects of robot-environment interaction have different influences on the deterministic chaos seen: for instance, changes to the environment are less influential than changes to the robot's control program. A full presentation and discussion of experimental results can be found in the references listed.

Ulrich Nehmzow
Department of Computer Science
University of Essex
Wivenhoe Park
Colchester CO4 3SQ, UK
http://sh656.essex.ac.uk/staff/udfn/

Towards a model of expectation-driven perception

Human perception is an active process by which meaningful information is gathered from the external environment. Application areas such as human-computer interaction (HCI), or the role of human experts in image analysis, highlight the need to understand how humans, especially experts, use prior information when interpreting what they see. Here, we describe how a model of expert perception is currently being extended to support expectation-driven perception of bitmap-level image data, focusing particularly on its ability to learn semantic interpretations.

The chrest model

CHREST (Chunk Hierarchy and REtrieval STructures[1]) is a computational model of perception and learning, designed to capture the perceptual knowledge acquired by an expert[2] (see also Gobet's individual article on page 8). Figure 1 illustrates the model's three main components: mechanisms for interacting with the external world; multiple short-term memories (STMs) to hold information from different input modalities; and a long-term memory (LTM) where information is held within a discrimination structure known as a 'chunking network'.

Recent work with CHREST is attempting to integrate three key processes for using expectations in perception: the use of bitmap data (whereas previous work has relied upon symbolic information), the creation of links between visual and verbal information, and the role of heuristics to guide the simulated eye. We describe the latter two in more detail here.

Combining visual and verbal chunks

CHREST's LTM holds information in the form of 'chunks', each of which is a familiar pattern in the environment. CHREST stores a chunk in a dual fashion. Firstly, the chunk itself is stored in a format representative of the data within it: in a visual domain, the chunk may be in the form of a bitmap; for a verbal pattern, it may be a sequence of phonemes. Secondly, the chunk's location in the model's LTM may be addressed directly with a link. Links are formed between nodes in the multiple STMs when they share an important relationship: such as being present in the environment simultaneously.

Figure 2 illustrates how a chunk acquired visually may be named by forming an association with a chunk acquired verbally. There are three steps. First, the visual pattern is sorted through LTM, and a pointer to the node retrieved is placed into visual STM. Second, the verbal pattern is sorted through LTM, and a pointer to the node retrieved is placed into verbal STM. Finally, a 'naming link' is formed between the two nodes at the top of the STMs.

Simulations with the CHREST model using semantic associations, such as those illustrated in Figure 2, demonstrate that CHREST captures several important phenomena illustrating the role of expectations in perception. These include: improved classification accuracy, faster classification, and the use of reconstructive memory to identify very noisy objects.[3]

Heuristics to guide eye fixations

An extended bitmap image cannot be perceived in its entirety. Instead, CHREST uses a simulated eye directed at a focus of attention—the fixation point—and has a limited field of view. The position of the eye is controlled with a set of heuristics that interact with each other. Prior work[2] has used various groups of heuristics that combine both bottom-up and top-down sources of information to guide the eye. In the top-down category, CHREST attempts to complete information held at a node referenced in the STM, to follow a test link, or to deepen the search within the LTM. Additional sources of information/heuristics include salient objects, novel objects, or default scanning of the scene.

CHREST is uniquely placed as a cognitive model of human learning in perceptual domains, with each area of Figure 1 interacting closely to gather and use meaningful information from a complex environment. With its recent extensions and use in domains with bitmap-level data, CHREST is currently being applied to domains involving the semantic analysis of complex images.

Peter Lane* and Fernand Gobet**
*Dept. of Computer Science
University of Hertfordshire, UK
E-mail: peter.lane@bcs.org.uk
http://homepages.feis.herts.ac.uk/~comqpcl
**Dept. of Human Sciences, Brunel University, UK
Details on page 8.

References
1. F. Gobet, P.C.R. Lane, S.J. Croker, P.C-H. Cheng, G. Jones, I. Oliver, and J.M. Pine, *Chunking Mechanisms in Human Learning,* **Trends in Cognitive Science 5** (6), pp. 236-43, 2001.
2. A.D. de Groot, F. Gobet, **Perception and Memory in Chess: Heuristics of the Professional Eye,** Van Gorcum, Assen, 1996.
3. P.C.R. Lane, A.K. Sykes, and F. Gobet, *Combining Low-Level Perception with Expectations in CHREST,* **Proc. of the European Cognitive Science Conf.,** 2003.

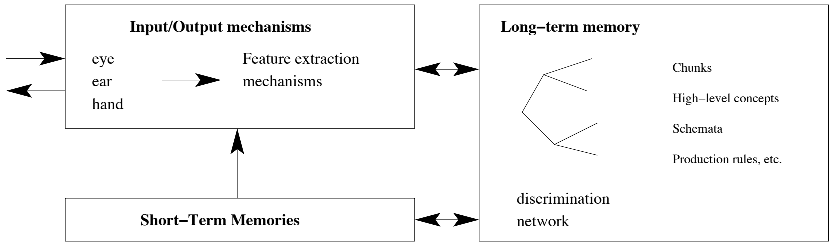

Figure 1. The CHREST Model.

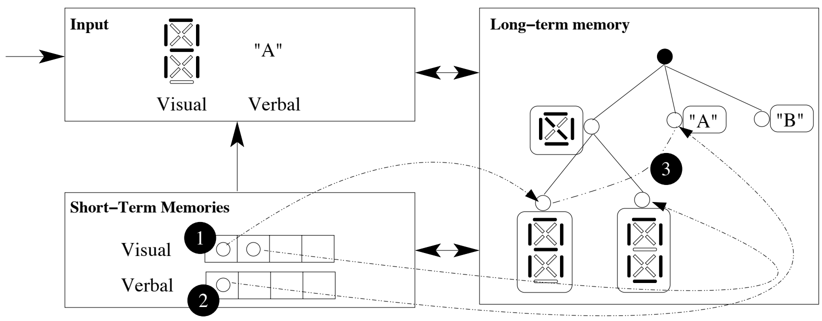

Figure 2. Learning a 'naming link' across two modalities.

Implicit learning of expert chess knowledge

Much of what we know about expertise comes from research into chess by de Groot in the forties and Chase and Simon in the seventies.[1,2] Two classic de Groot results demonstrated the importance of perception in expert behaviour. First, even though grandmasters found better moves than strong amateurs in a problem-solving task, there were few differences in their search behaviour. In particular, all players were selective and visited only about one hundred positions. Second, chess masters performed almost perfectly in the recall of game positions (see Figure 1) presented for a few seconds. To explain these results, Chase and Simon developed the 'chunking theory' that proposed mechanisms specifying how knowledge is implicitly acquired during practice. Expertise is seen as the acquisition of a large number of perceptual chunks (groups of features that can be used as units), that give access to relevant information (e.g., what move to play).

Over the last decade, my research has aimed to flesh out these mechanisms computationally and to test them empirically. The computational work has led to the development of CHREST (Chunk Hierarchy and REtrieval STructures), which models expertise as the growth of a discrimination net. Each node (chunk) in the net contains information about the location of pieces, as well as pointers to possible (sequences of) moves. Provision for eye-movement mechanisms enables a close interaction between perception and memory. Finally, high-level schemas are created automatically. The empirical work has investigated expert perception and problem solving using verbal protocols, eye movements, and—more recently—brain imaging. I have also manipulated several variables in recall experiments, such as time of presentation, level of position distortion, and level of position randomisation. In general, CHREST, serving as a subject 'in silico', models the memory experiments well. Here, I focus on the recall of random positions.

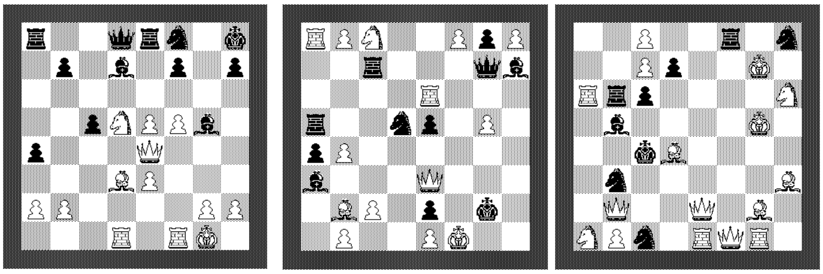

Figure 1. A game position (left), random position (middle), and 'truly random' position (right).

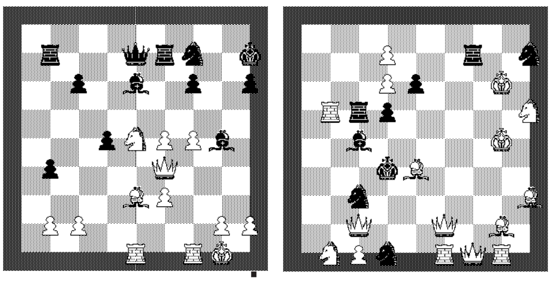

Figure 2. Positions where the pieces have been placed at the intersection of squares: a game position (left); and a truly random position (right).

As documented in psychology textbooks, Chase and Simon found no skill difference in the recall of random positions (see Figure 1). However, CHREST predicts a small difference, as chunks are more likely to be recognized serendipitously in random positions with large nets than with small ones. Re-analysis of the literature, as well as the collection of new data, supported this prediction.[3]

Random positions are typically created by shuffling the piece locations of a game position. Vicente and Wang[4] noted that these positions are not really random, as they still contain information about the distribution of pieces (e.g., only one white King is allowed). They raised the question as to whether skill differences would remain if 'truly-random' positions were used, where both the location and the distribution of pieces are randomised (see Figure 1). CHREST predicts that this would be the case. An experiment with 36 players ranging from weak amateurs to grandmasters confirmed CHREST's prediction: with truly-random positions, there was a statistically reliable correlation between skill and recall performance.[5] This difference remained when variables such as age and visual memory were partialled out.

Current work with Andrew Waters further explores the role of perception in expert memory. We created positions where the pieces lie at the intersection, rather than the middle, of squares (see Figure 2). Results indicate that overall performance drops drastically. While masters still maintain some superiority with game positions, they do not perform better with random and truly-random positions. CHREST simulates these results by assuming that players need to 're-centre' the pieces in their mind's eye in order to facilitate the recognition of chunks. This takes time and thus lowers performance.

Beyond chess, the chunking mechanisms embodied in CHREST have explained empirical data in other domains.[6] Within expertise research, they have accounted for computer programmers' memory and the learning of multiple representations in physics. Beyond this, they have helped model how children acquire the syntactic categories of their native language, and how humans combine information from different input modalities (see Peter Lane's contribution on page 7). Overall, CHREST shows that simple mechanisms leading to the implicit learning of a large number of chunks may underpin (expert) behaviour in a number of domains.

Fernand Gobet
Department of Human Sciences
Brunel University
Uxbridge, Middlesex, UK
E-mail: fernand.gobet@brunel.ac.uk
http://www.psychology.nottingham.ac.uk/staff/frg/

References
1. A.D. de Groot and F. Gobet, **Perception and memory in chess,** Van Gorcum, 1996.
2. F. Gobet, A.J. de Voogt and J. Retschitzki, **Moves in mind: The psychology of board games,** Psychology Press, *in press.*
3. F. Gobet and H.A. Simon, *Five seconds or sixty? Presentation time in expert memory,* **Cognitive Science, 24,** pp. 651-682, 2000.
4. K.J. Vicente and J.H. Wang, *An ecological theory of expertise effects in memory recall,* **Psychological Review 105,** pp. 33-57, 1998.
5. F. Gobet and A.J. Waters, *The role of constraints in expert memory.* **Journal of Experimental Psychology: Learning, Memory & Cognition**, in press.
6. F. Gobet, P.C.R. Lane, S. Croker, P.C.H. Cheng, G. Jones, I. Oliver and J. M. Pine, *Chunking mechanisms in human learning,* **Trends in Cognitive Sciences 5,** pp. 236-243, 2001.

Development and evolvability in hardware

The field of evolutionary computation (EC)—using principles of nature's survival of the fittest to perform robust learning—has been used in many interesting applications over the last two decades. Combined with emerging electronic technology, the field of evolvable hardware (EHW) aims at broadening the scope of EC applications to designing electronic circuits using artificial evolution. Since the field first emerged, its scope has been broadened to generally incorporating biological concepts in electronics or related areas: such concepts include self repair, self reproduction, and adaptation. Today, the field tends to be focussed on areas where sustained operation despite harsh conditions is desirable: space and military applications in particular.

Now EHW has moved beyond the proof-of-concept phase, the move towards larger-scale applications has been limited by the sheer complexity of electronic circuits and the computational resources required for their development. This has pushed a number of researchers towards re-evaluating basic concepts of EC techniques in hardware, inspired by the apparent lack of scaling difficulties of natural evolution. Our work focuses on investigating genotype-phenotype development as one such possible solution within EC in general, and within EHW in particular.

A popular way of overcoming resource limitations is to use the computational equivalent of a biological developmental mapping from genotype to phenotype. EC techniques typically directly encode and optimize traits of possible solutions. Incorporating development in EC implies optimising a generative 'building plan' that develops into a candidate solution before being evaluated (very much like the way human DNA encodes the building plan for an embryo). A schematic overview of a developmental mapping within a typical EC setup is shown in Figure 1.

Typically, such a mapping results in a high degree of gene interaction, a mechanism that is at the heart of why it scales well. However, while there is a good chance that scaling issues can indeed be (partially) overcome using a developmental mapping, it is far-too-often overlooked that uncontrolled degrees of this gene interaction can result in high degrees of gene interdependence: a subclass of so-called gene 'epistasis'. The EC community is well aware that this effect is the key mechanism that makes problems difficult to solve using EC.

The project presented here aims to combine knowledge from EC, biological development, and EHW. It is also intended to offer a first step towards designing and understanding developmental mappings for practical applications like EHW that are both able to encode the typical phenotypes of the application area, and allow favourable evolutionary properties. Practically, this comes down to exploring the balance between efficient phenotype expression on one hand, and evolvable genetic encodings on the other. In our opinion this area is the key to applying development to EC (and EHW), and is far too easily neglected, especially in 'applied' EC.

The project

Our project consists of two main parts: an introductory study of the properties of developmental fitness landscapes, and the design and analysis of a conceptual development model aimed at evolving and developing cellular automata. In particular, the first stage looks at the family of NKd fitness landscapes—an adapted version of Kauffman's NK landscapes[1]—and models the higher impact of early developmental operations versus those later on. Different flavours of NKd landscapes have been investigated for a range of typical properties and compared to the NK variety. Furthermore, actual evolutionary runs on the different landscapes have been compared.[2]

The second stage involves a more elaborate investigation of an actual evolution and development system that is aimed at evolving cellular automata (1D or 2D). This toolbox is designed with the basic properties of biological development in mind, and allows tuning with respect to a broad range of parameters of the developmental system. Generally speaking, the system allows both mosaic and regulative development: both of which are archetypical mechanisms in biological development.

Piet van Remortel, CoMo, VUB, Belgium

Continued on page 10.

References
1. Stuart A. Kauffman, **The Origins of Order**, Oxford University Press, Oxford, 1993.
2. Piet van Remortel and Tom Lenaerts and Bernard Manderick, *Developmental Effects on Tuneable Fitness Landscapes*, **Lecture Notes in Computer Science 2606,** pp. 117-128, 2003.
3. S.F. Gilbert, **Developmental Biology**, Sinauer Associates Inc., 2000.
4. G. Harik, **Learning Gene Linkage to Efficiently Solve Problems of Bounded Difficulty Using Genetic Algorithms**, PhD. Dissertation, University of Michigan, 1997.
5. Rudolf A. Raff, **The Shape of Life**, The University of Chicago Press, Chicago, 1996.
6. Piet van Remortel, **Evolvability and Development in EHW** (working title), PhD. Dissertation, Vrije Universiteit Brussel, Belgium.

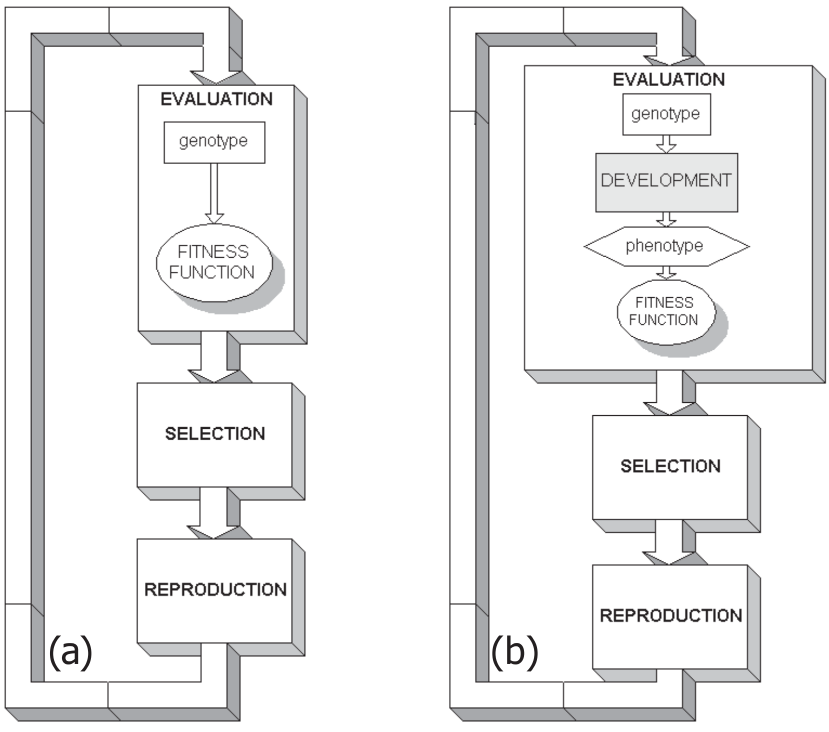

Figure 1. A comparison of a classical EC setup (a), and a setup with development (b). While in (a) the genotype is evaluated immediately and without (much) prior processing, (b) contains a process of development from genotype to phenotype which then is evaluated.

Review

Cognitive Modeling

Edited by Thad Polk and Colleen Seifert

Publisher: MIT Press
http://mitpress.mit.edu
Hardback: ISBN 0262161982
£72.95
Paperback: ISBN 0262661160
£36.50
Pages: 1270

At well over a thousand pages, this book is a bargain. It's a collection of tried-and-true papers in cognitive science, with a special emphasis on human cognition. (So topics such as naive physics, adaptive systems, complex social interactions, and machine learning and tutoring, are deliberately omitted.) As such, it's not exciting: but it is very useful. On the one hand, it could form the basis of graduate and undergraduate courses in various areas, while on the other hand it would be a convenient 'compendium' for the researcher's bookshelf.

The 38 chapters are organized in three parts: Architectures and Approaches (symbolic models and neural networks); Case Studies in Cognitive Modeling (a wide range of specific examples); and Issues in Cognitive Modeling (the general philosophy/methodology of these techniques).

The many classics here include Anderson on ACT and Carpenter/Grossberg on ART; Elman on recurrent networks and Hinton on connectionist learning; seminal papers on Hopfield nets, back-propagation, supervized learning, optimality, and analogical reasoning; Plunkett and Marchman on the past tense; and, last but not least, Newell (and colleagues) on early-Soar and on the general philosophy behind it.

If that sounds worthy-but-boring, you should know that there are also a number of papers which I, at least, hadn't previously come across. Especially in the 'Case Studies' section, there are contributions that may be fresh even to researchers in the field. So one could actually learn a lot, as well as having one's old favourites readily to hand.

In trying to do this, the index will be helpful. Far too often, edited collections include no index: no-one could be bothered to compile one. I hope I'll be forgiven if I don't trawl through the 1200+ pages to check the index out. Possibly, every entry could have had a dozen additional references; and possibly, there could have been many extra entries. (There are no entries for names, as such: "Stroop task" and "Smolensky architectures" are there, but Stroop and Smolensky aren't. Someone who wants to get a sense of Jo Bloggs' overall contribution will gave to read his chapter and consult the bibliographies of the others.) But anyway, at four small-print pages, it's a useful start.

Also useful is the list of postal and e-mail addresses of the 76 contributors. Yes, a mere detail: but details make the difference. All in all, the editors have done a good job.

Margaret A. Boden
Research Professor of Cognitive Science
University of Sussex, UK
E-mail: maggieb@cogs.susx.ac.uk

Boden is author of Purposive Explanation in Psychology *(1972),* Artificial Intelligence and Natural Man *(1977/87), and* Computer Models of Mind *(1988).*

Development and evolvability in hardware

Continued from page 9.

Furthermore, it supports tunable gene regulation, cell communication, a time variable that can be referenced, etc. Without getting into details here, we claim that this model, although conceptual, resembles natural development to a reasonable extent and serves the purposes of our investigation.

Towards results

The general approach in this investigation was to start from the very basics: incorporating development, gradually adding complexity to the model, trying to understand the consequence of every step. In view of this, a first step was subjecting evolution to NKd problems, which turns out not to be very difficult for evolution. Dependent on NKd parameters however, convergence towards a solution can vary from sequential to parallel.[2]

Experimentation with the development system is in its final stage. The cornerstone of this investigation is the modularity of the developmental mapping, since this allows reasonable 'evolvability'. Experimentation is set up in three phases. In the first, basic patterns are encoded by hand in a genetic code, exploring different expression mechanisms and related modularity of the genome (measured based on gene clustering). Preliminary experiments indicate local cell communication (the basis of regulative development) and the parallel signalling channels offered by multiple proteins as important sources of exponential gene interaction.

In a second phase, evolutionary experiments investigate the advent of modularity through mechanisms such as gene duplication and divergence[3] and ideas based on linkage learning.[4] Once more modular, the effect of biological concepts such as heterochronic mutations[5] are investigated in the phenotype. Finally, a conceptual investigation of non-deterministic development is performed. This, given the right circumstances, can improve the robustness of evolved phenotypes. Further information on these experiments will be forthcoming.[6]

Piet van Remortel
Computational Modelling Lab (CoMo)
Vrije Universiteit Brussel (VUB), Belgium
http://como.vub.ac.be/Members/Piet.htm

References on page 9.

Review

The Analogical Mind
Edited by Dedre Gentner, Keith Holyoak, and Boicho Kokinov

Publisher: MIT Press
http://mitpress.mit.edu
Paperback: ISBN 0262571390
£26.50
Pages: 520

This book is well described in its introduction as, "a set of papers that collectively lay out the 'state of the art' in our current scientific understanding of the mental processes involved in the use of analogy in cognition." Analogy itself is defined, first loosely, as the ability to think about relational patterns. But most of the book's contributors take on Dedre Gentner's characterisation of analogy as a process of 'structural alignment' between two domains. The importance of this (suprisingly recent) idea is that it goes beyond simple listing of common properties to explain some of the richness of analogical reasoning.

The initial chapters present theoretical and computational models of this process of structural alignment. The architectures vary from propositional, to agent-based, connectionist and hybrid systems. All have been shown capable of replicating human data on analogical thinking and, as presented, it is difficult to compare them. It is good, however, that all address the theme of how to go beyond the various successful matching algorithms to embed these systems in wider cognitive contexts: e.g., reasoning (Forbus), memory (Kokinov and Petrov), development (Wilson et al), and cognitive neuropsychology (Holyoak and Hummel).

There follows a series of chapters that summarise observational and experimental study of the use of analogy in various domains. These vary from rather focussed synopses of the author's own research using a specific paradigm (e.g. Bassok's study of 'word problems' in mathematics, or Dunbar's investigations of analogy use in science and politics) to more general overviews of areas of application (e.g. effects of analogy on political decision-making and consumer choice by Markman and Moreau, or in expression of emotion by Thagard and Shelley). In my view the weakest chapters were the final two, on analogy use in infants and primates. The authors seemed to over-interpret the rather scant evidence that analogy use is not unique to mature humans: tasks that seemed to indicate little more than simple learning.

This raises the interesting issue of what is *not* analogy use. The majority of the authors in this book are keen to claim structural alignment as a fundamental underlying process in nearly all activities that are considered cognitive: e.g. language use, reasoning, perception, memory. It is certainly hard to find sharp distinctions between 'analogy', 'metaphor', and 'similarity': Gentner et al. explicitly present, "an approach that unifies metaphor with processes of analogy and similarity." Fauconnier argues that, "structure mapping is inherent in all our thought processes," including the construction of meaning, and discusses 'conceptual blending' as a process akin to analogy. However, Keane and Costello cast a critical eye on this tendency and present arguments and evidence for an alternative explanation of concept combination. The 'epilogue,' by Hofstadter, presents a thought-provoking essay on why he feels, in cognition, "analogy is everything, or very nearly so." Whether readers are sympathetic or antagonisitic to this conclusion, this book will certainly allow them to become better informed about the current theory and evidence in this area.

Barbara Webb
Reader
Institute for Perception, Action and Behaviour
University of Edinburgh, UK
http://homepages.inf.ed.ac.uk/bwebb/

Family resemblance, Bayesian networks...
Continued from page 1

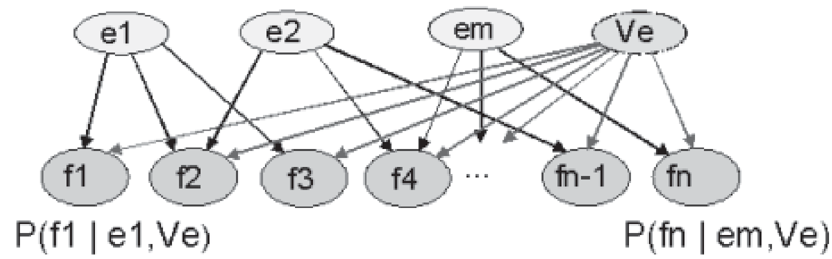

Figure 3. Bayesian network representation of an exemplar-based model.

and a good indication of the accuracy within a category is the proportion of training cases the exemplars represent.

Sunil Vadera,* Andres Rodriguez, and Enrique Succar†**
*University of Salford
http://www.ibi.salford.ac.uk/isr/vadera/svadera.html
**Instituto de Investigaciones Electricas, Mexico
†ITESM, Campus Morelos, Mexico.

References on page 1.

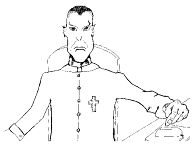

Father Hacker's Guide

for the Young AI Researcher

Cognitive Divinity Programme
Institute of Applied Epistemology

Congratulations! You have succeeded in getting a research paper accepted for a conference or being invited to give a seminar. Now, it is vital that your presentation makes the best possible impression and further enhances your international reputation. You need to know...

8. How to give a presentation

1. As we discussed in Guide #6, the ideal presentation is erudite, profound, insightful and entertaining. It's a rare presenter who can meet this specification, so in this guide I will adopt the second best option: convincing your audience that you have reached a level of enlightenment that is beyond ordinary mortals. If you can't lose your audience in the first five minutes, then your research can't be really challenging. Your presentation should consist of baffling technical detail, peppered with ambitious claims for its significance.

2. The convergence of tools for preparing papers and presentations means that it is a trivial matter to copy whole paragraphs and proofs from the paper and paste them into a slide. Now your audience have the full benefit of the complex and detailed arguments from the paper to complement your verbal presentation. Don't worry if their eyes glaze over with the effort of absorbing so much information, that will only demonstrate the depth and difficulty of your research.

Announcing a New, Improved version of Hacker's OBSCURANT™ (OBSCURANT Baffles with Science, Clouds Understanding, Reasoning, and Arguments with Nonsensical Transformation)

Now with unintelligible jargon, confusing typos, and miniscule fonts.

---

3. The combination of laptop and data-projector enables your presentation to be right up to date. Not only can you leave the preparation of your slideshow to the plane or train journey, but you can be polishing it as you await your turn to speak: or even during the talk itself. Don't worry about the many bugs that such last minute preparation inevitably introduces; your audience will be entranced by the feeling of immediacy and engagement that this engenders.

Construct your slides as you talk with Hacker's SCRAWL™ (Speech Captured, Recognised, Abstracted, Written, and Lit-up)

State-of-the-art speech recognition software will process, précis, and present your talk as you deliver it.

---

4. It is not necessary to master this new technology. There is bound to be a whiz-kid in the audience who can get your laptop working in harmony with the data-projector in only 5 or 10 minutes. It is not your problem if these teething troubles disrupt the published conference schedule; good programme chairs always build-in plenty of slack to mop up timetable overspills. If not, then why shouldn't less important talks be foreshortened to make space for yours?

5. Similarly, don't feel constrained to fit within the narrow-minded time constraints set by your intellectual inferiors. Running overtime will demonstrate the impossibility of fitting your wealth of achievement into artificial limits. Better still, if you are forcibly prevented from completing your presentation then any aggressive questioners can be referred to the undelivered part of your talk.

Dealing with questions will be the subject of a future guide. For those who are unable or unwilling to wait for this, note that if you have, as recommended, run overtime, then there will be no time for questions. In case a short question period is, nevertheless, imposed, feel free to ignore the question and use the opportunity to continue your interrupted presentation.

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

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

Next AISBQ deadline: 14 November 2003

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.