AISB Quarterly

No.132

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

Representations and Architectures to Support Diagrammatic Reasoning

The need for internal diagrammatic representations

External diagrams are a common aid in every day and professional problem solving, explanation, and information presentation. Since the 1990s, there has been a boom in research on diagrammatic reasoning, a harbinger of which was the book (Glasgow et al. 1995) [1], while current research is showcased in a series of bi-annual international conferences on diagrams. This renewed interest in diagrams has many sources. Because of the role of graphs and diagrams in the general culture as well as in teaching, psychologists and educators have become interested in the perceptual and cognitive processes involved in understanding and using graphs, and logicians have taken interest in the logic of diagrams, specifically about whether diagrammatic reasoning can be sound, instead of diagrams being merely heuristic aids for constructing symbolic proofs. Interest in visual programming and human-computer interfaces in general has brought a new group of researchers into the field. And of course within AI, researchers started building systems that created and used diagrams as part of their problem solving. We review some of this research in [2] and [3].

A good deal of AI and cognitive science is based on modeling central cognition in terms of predicate-symbolic representations and symbol processing operations on them. In this framework, the role of perception is limited to delivering to cognition information about the world in the form of predicate-symbolic expressions, e. g. "Block(A), Block(B), ON(A,B)" that perception may deliver to cognition when the agent is looking at a Blocks World configuration.

Over the last decade, my collaborators and I have also been involved in research on diagrammatic representations and reasoning (reported in [2]-[12]). Our work has points of contact with the streams mentioned above, but its focus is a bit different. All the research mentioned above is concerned with external diagrams. We, on the other hand, have been chiefly concerned with diagrammatic representations internal to the agent, representations we propose agents have as part of their cognition even when processing external representations, and of course while imagining. These representations participate in thinking and problem solving. The proposal itself is not novel; it has been at the heart of claims about the role of mental imagery in reasoning. What is novel about our work is its computational nature.

To see the need for diagrammatic internal representations and mental operations on them, consider Figure 1. Answering the question in Figure 1 requires the agent to imagine moving the object labeled A, create an internal representation that stands for a composition of the two objects, one of which is imagined, and apply internal perception of Inside to various such compositions. Similarly, while looking at an external graph, the user may need to mentally extend certain lines to predict trends or infer missing values.

While an external representation, say on paper, is an intensity array, in order to support the reasoning involved in the examples, the internal representation needs to be organized in terms of individuated, separately addressable and manipulable objects, with the spatiality of the objects, the location as well as the extent, represented. The first part of perception of an external representation, a part that is cognitively impenetrable, processes the input from the external world to produce a figure-ground separation of the scene – in our case the external representation.

One of the outputs of this part is a representation consisting of a set of individuated objects, along with their spatiality. This representation is part of cognition, and supports the perceptual experience of the agent in seeing the scene or external representation as a configuration of shapes. The representation is also available to be operated on by a set of internal perception operations.

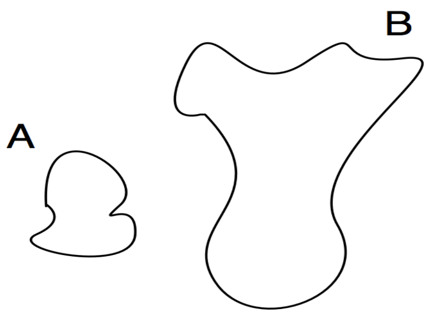

Figure 1. A question requiring mental imagery operations: can region A fit into region B?

These operations range from simple ones such as counting (e. g. "How many windows are there in the living room?") to more complex relational operations such as Inside (A,B), Longer-than (A,B), etc. We discuss in [3] what kinds of perceptions are and aren't possible on this representation.

We unify perception of external representations and imagination by proposing that the same locus in cognition for the object-individuated configuration of shapes resulting from perception of the external world is also the locus for mental representations of images from memory or imagination operations. The internal representation directly corresponding to Fig. 1 would consist of the two objects and their spatiality. When the agent

In this issue:

Features and Reviews

Society

Father Hacker

Diagrammatic reasoning (cont.)

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

ISSN 0268-4179

mentally operates on A by moving it, say until A is inside B, at that point there would be a composite object in cognition consisting of shape A inside shape B. The agent can apply the internal perception Inside (A,B) to this internal representation.

When diagrams are used as representations, some of the diagrammatic objects may be intended to be taken as points and some as curves without thickness, though on paper these objects will have spatial extent in order to assist human perception. The internal representation, however, will encode the intended status of the objects as points, curves or regions.

The DRS system, perception and action routines, and integration into cognitive architectures

In [4] we proposed DRS as a domain-independent system for internal representation of black-and-white line diagrams. A diagram in DRS is a configuration of diagrammatic objects, each of which is one of three types: point, curve, and region. Associated with each object is the specification of the points in the 2-D space that defines the object, and additional features such as symbolic labels that are often attached to diagrammatic objects in physical diagrams. DRS can be hierarchical, to represent a configuration of objects being seen as a single objects, e.g. a cluster of small regions as a large region, and this can help in modeling attention shift from details to abstractions and vice versa. DRS representations can be constructed, as the needs of problem solving dictate, as a composition of elements from external representation, memory, and results of mental imagery operations.

We have implemented the spatiality representation in two different ways [5]. The first was in a purely algebraic framework [6]: curves, either as objects or as closed curves describing the peripheries of region objects, are specified as algebraic equations. In the second, the objects are represented in 2D arrays [7], similar to those used in the visual representations in working memory in [13] and [14]. General frameworks for composing internal perception and diagram creation/modification algorithms have been developed in [5] for both implementations, array, and algebraic. The algorithms can detect emergent and vanishing objects as objects are added or removed from a diagram. Internal perceptions can be applied to a composition of diagrammatic objects from external representation, memory, and imagination. None of these algorithms is intended to simulate the corresponding algorithms in the human architecture, and hence they are not useful in predicting the timing and error properties of human performance.

To be useful in problem solving, the diagrammatic component in DRS and the associated perception and diagram creating/modification algorithms have to be integrated within a classical symbolic architecture such as Soar [15] and ACT-R [16]. One way to do this is modular. The control component of the main architecture calls on a diagrammatic module to solve subproblems that require access to the diagram. The module has the diagram represented in DRS and comes with a set of perception and action operators. The module returns the relevant symbolic information to the main part. Matessa et al.'s work [17] is an example of such an approach, where ACT-R is augmented with a DRS-based diagrammatic module. The biSoar effort [8]-[10], based on a theoretic stance about the multimodality of the cognitive state [18], makes all cognitive state representations—in goals, WM states, the state descriptions in production rules—bimodal. For example, all states have, in addition to the traditional predicate-symbolic component, a diagrammatic component, represented in DRS, that depicts the visualizable aspects, if any, of the state. Just as symbolic operators are available to operate on the predicate-symbolic state, internal perception operators are available to solve relevant subgoals. Soar's design is unchanged in all other respects. In [9] and [10], we describe the application of biSoar to modeling some spatial memory tasks.

While the work described here is limited to diagrams, the larger ambition is to encourage a view of cognition as multimodal [18], where the various perceptual modalities and the kinesthetic modality are as much part of thinking as language-based symbolic representations, which dominate current models.

Notes

1. Supported by U. S. Army Research laboratories under Cooperative Agreement DAAD19-01-2-0009.

References

[1] Glasgow, J., Narayanan, N. H. & B. Chandrasekaran (1995). Diagrammatic Reasoning: Cognitive and Computational Perspectives. AAAI Press/MIT Press: Cambridge, MA.
[2] Chandrasekaran, B. and O. Lele (2010). "Mapping Descriptive Models of Graph Comprehension into Requirements for a Computational Architecture: Need for Supporting Imagery Operations". In Goel, A.K., Jamnik, M. and N.H. Narayanan (Eds) (2010). Diagrams 2010, Lecture Notes in Artificial Intelligence (LNAI) 6170, Heidelberg and Berlin: Springer-Verlag, 235-242.
[3] B. Chandrasekaran, B. Banerjee, U. Kurup, O. Lele. (to appear in 2011). "Augmenting Cognitive Architectures to Support Diagrammatic Imagination", Topics in Cognitive Science.
[4] Chandrasekaran, B., Kurup, U., Banerjee, B., Josephson, J. R., & Winkler, R. (2004). "An architecture for problem solving with diagrams". In Blackwell, A., K. Marriott, & A. Shimojima (Eds), Diagrammatic Representation and Inference, Berlin: Springer-Verlag, 151–165.
[5] Banerjee, B. (2007). Spatial problem solving for diagrammatic reasoning. Ph.D. dissertation, Department of Computer Science & Engineering, The Ohio State University, Columbus.
[6] Banerjee, B., & Chandrasekaran, B. (2010). "A constraint satisfaction framework for executing perceptions and actions in diagrammatic reasoning". Journal of AI Research, 39, 373–427.
[7] Banerjee, B., & Chandrasekaran, B. (2010b).

Diagrammatic reasoning (cont.)

"A spatial search framework for executing perceptions and actions in diagrammatic reasoning". In Goel, A. K., Jamnik, M. & N. H. Narayanan (Eds) (2010). Diagrammatic representation and inference, LNAI 6170, Springer: Heidelberg, 144–159.
[8] Kurup, U. (2007). Design and Use of a Bimodal Cognitive Architecture for Diagrammatic Reasoning and Cognitive Modeling, Ph D. Thesis, Columbus. The Ohio State University, Ohio
[9] Kurup, U., & Chandrasekaran, B. (2007). "Modeling memories of large-scale space using a bimodal cognitive architecture". In Lewis, R. L., Polk, T. A. & J. E. Laird (Eds). Proceedings of the 8th International Conference on Cognitive Modeling Ann Arbor, Michigan: University of Michigan, 267–272.
[10] Kurup, U., & Chandrasekaran, B. (2009). "A cognitive map for an artificial agent". In Proceedings of the Second Conference on Artificial General Intelligence, AGI 2009, Arlington, Virginia. Amsterdam-Paris: Atlantis Press, ISBN: 978-90-78677-24-6, 85-90.
[11] Chandrasekaran, B. "Diagrams as Physical Models, in Diagrammatic Representation and Inference". In Barker-Plummer, D., R. Cox, R. and N. Swoboda (Eds) (2006). LNAI 4045, Heildelberg: Springer, 204-217.
[12] Chandrasekaran, B. (2011; forthcoming) "When is a Bunch of Marks on Paper a Diagram? Diagrams as Homomorphic Representations," Semiotica.
[13] Tabachneck-Schijf, H. J. M., Leonardo, A. M., & Simon, H. A. (1997). "CaMeRa: A computational model of multiple representations". Cognitive Science, 21(3), 305-350.
[14] Lathrop, S. D., and Laird, J. E. (2007). "Towards Incorporating Visual Imagery into a Cognitive Architecture". In Lewis, R. L. Thad A. Polk, and John E. Laird (Eds), Proc. 8th International Conference on Cognitive Modeling, University of Michigan, Ann Arbor, Michigan, 25-30.
[15] Laird, J. E., Newell, A., & Rosenbloom, P. S. (1987). "Soar: An architecture for general intelligence". Artificial Intelligence, 33(1), 1-64.
[16] Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). "An integrated theory of the mind". Psychological Review, 111(4), 1036-1060.
[17] Matessa, M., R. Archer, and R. Mui (2007). "Dynamic Spatial Reasoning Capability in a Graphical Interface Evaluation Tool". In Lewis, R. L., Thad A. Polk, and John E. Laird (Eds). Proc. 8th International Conference on Cognitive Modeling, University of Michigan, Ann Arbor, MI, 55-59.
[18] Chandrasekaran, B. (2006). "Multimodal Cognitive Architecture: Making Perception More Central to Intelligent Behavior". In Proceedings of the AAAI National Conference on Artificial Intelligence, 1508-1512.

Prof. B. Chandrasekaran
Laboratory for AI Research
Department of Computer Science and Engineering
The Ohio State University

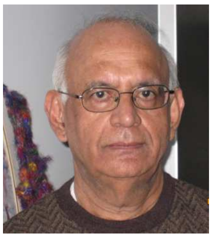

A New COST Action: Autonomic Road Transport Support (ARTS) Systems

Mobility of people and goods is a key challenge for the future. Transport is one of the world's largest industrial sectors, yet challenges and frequent failures of road transportation networks are well known, with the cost of congestion alone estimated at Euro 100 billion in the EU[1].

Systems of road traffic flow are affected by the outcome of individual driving decisions, often assisted by personalised navigation and information-providing devices. This combined with the complex topology of the network and the random occurrence of capacity reducing events make for a complex system. Within this system control centres utilise a range of assets (traffic signal, variable speed limits, re-routing etc) to help optimise the flow of network traffic with respect to a range of rules, regulations and policies relating to efficiency, safety and environmental criteria.

Over the past 30 years or so, ICT has been applied with a certain amount of success by highways authorities and urban traffic controllers to support traffic management. The application of ICT to Transport has led to what is termed Intelligent Transportation Systems (ITS), though these systems generally do not embody intelligent properties in the AI sense. There has been some use of AI in traffic support [2], for example the application of ANPR in number plate recognition, or the use of A*-inspired algorithms in routing algorithms. Most notable is the widespread success over the last 20 years or so of adaptive traffic light algorithms, for example in SCOOT systems (http://www.scoot-utc.com/).

These systems sense traffic flow and adapt signal control plans for collections of traffic lights in order to optimise traffic flow through junctions. In general, stakeholders in the traffic support area (consultants, equipment suppliers, transport authorities) have embraced new developments in ICT, and are well disposed to the deployment of AI techniques. Further, the level of interoperability and data representation is relatively high compared to some public sectors, and is characterised by the widespread use of the "UTMC" (http://www.utmc.uk.com/) by local authorities in the UK. UTMC is essentially an interface specification at the relational database level for data interchange, used to connect up the wide range of data flowing through control centres.

Road traffic controls and surveillance systems can be viewed as forming large scale, heterogeneous, control systems, complicated by the dependencies on human behaviour. To effectively "manage" or enable the "optimisation" of such a socio-technical system is a daunting task, exacerbated by rising public environmental and operational expectations. Recent technological advances (e.g. novel ramp metering control, variable speed limits, surveillance interpretation, and road user information systems) have led to incremental improvements in the performance of road transportation networks; taken as a whole, however, the effect is more management controls, more surveillance data, and more complex and demanding goals than current operator-centric systems can manage. Quite apart from the ability for trained human experts to make informed decisions and plans in such complex, real time systems, the cost of configuring and managing them is enormous. Current and planned related EU initiatives in areas such as "Smart Cities" appear to make even greater demands on IT, and in particular demand complex, intelligent software-intensive systems within their infrastructure.

A recently-approved COST (European Cooperation in Science and Technology) Network on Autonomic Road Transport Support (ARTS) Systems is being set up to explore the potential of embedding "autonomic" properties into the design of transportation systems. Autonomic Computing was launched around ten years ago by IBM [3] and can be viewed as a challenge to embed desirable self-managing intelligent properties into large systems to cope with the problems of their inherent complexity. The potential benefits of autonomic systems are in helping to solve the core problems of engineering road transport support (RTS) systems: their costly configuration and maintenance, high operating complexity, sub-optimal operation, and the problem of embedding and maintaining safety and environmental conditions within the operational parameters of the controlling system. Autonomic Computing integrates ideas from several areas of AI including automated reasoning, machine learning and automated planning, and implementations often draw on distributed AI technologies such as intelligent agents.

There has been little research into the many challenges of implementing autonomic behaviour in transportation systems, and what has been done has been carried out in a range of fragmented research areas. Some recent pilot studies concerning RTS technologies use agent-based technology [4,5,6,7]. Utilising a more centralised notion of self-maintenance, theory refinement algorithms for automatically evolving a requirements model of air traffic control criteria was developed in work sponsored by the UK NATS [8]. Concentrating on self-organization is the focus of "Organic Computing", a large cooperative research effort sponsored by the DFG [9]. Its goals are the development and control of emergent and self-organising technical systems. In the area of Organic Computing several projects have investigated the feasibility of adaptive, intelligent traffic light controllers and their ability to self-organise, e.g. to form progressive signal systems [10].

The challenge of embedding autonomic properties into RTS infrastructure is great, and will be tackled effectively only if a co-ordinated, continent-wide set of experts can be mobilised. The Action will initially focus on community building: with an initial start of 31 member institutions representing 14 countries in Europe, it will explore the application of AI techniques to large, complex control systems, in particular those techniques with the potential to embody autonomic behaviour in systems supporting road transport.

Research in related disciplines tends to follow one paradigm or become embedded within one particular framework. The primary focus of the COST Action will be to provide the scientific environment for a concerted effort towards the analysis and development of techniques for engineering autonomic behaviour in RTS systems. Surveying the literature on Autonomic Computing, there are a range of architectures and techniques used both from Computational Intelligence area and from classical AI. Hence the Action will encourage research groups to consider a range of architectural approaches, taking into account the heterogeneous, embedded, spatially distributed nature of the area, and the enormous amount of data and knowledge that RTS systems currently have available.

Embedding autonomy into a system requires building into the system the semantics of its own functions, so that it can have some measure of self-awareness. The idea of embedding meta-data within systems is now well established and is fundamental to the development of the semantic web and its associated service-oriented and semantic technologies, as well as the widespread use within the scientific community of ontology and supporting tools. Hence, a major theme within Network is how to harness service-oriented and semantic approaches to enable such behaviour as dynamic system configuration from primitive components. For example, how can current ITS technologies be "wrapped" into services that can be subject to automated assembly and control, in response to high level traffic

A new COST action (cont.)

policies? The benefits of this approach are that it hides the complexity of individual components (and makes them easier to maintain), while allowing new or changed high level policies to automatically deliver new and alternative mixes of control services.

The ARTS Action includes experts from several areas of computer science, engineering and mathematics, to bring together those with complementary backgrounds. With a focus on architectures, methods and models for ARTS, the Network will build on past research and development within ITS, and lessons learned from previous pilot studies in AC, to provide insights into appropriate platforms and methods for engineering ARTS systems. The Action will organise workshops, industrial-facing seminars, training schools and develop a road map and demonstrator systems in order to build up the critical mass of a research community.

While intuitively appealing, resources aimed at embedding autonomic properties into systems still require a business case, or an objective argument for who will benefit from ARTS, and in what measure. The Action intends to identify and quantify the scope, nature and potential pay-off with respect to financial, environmental and safety criteria, and look at the overall question of return on investment of autonomic systems within the transport area. Some transport areas may be considered more naturally amenable to autonomic techniques, such as local and regional control centre planning support, and real time traffic control. Applications such as automated incident detection may be considered as more problematic, as human judgement may always be superior in determining causes for alarm.

Other applications, and emerging technological and organisational ideas, such as cooperative and infrastructure systems, vehicle-to-vehicle enabled traffic support and demand management, need to be investigated from an ARTS standpoint. One particularly vexed question is that of the wider implications of the introduction of self-managing systems with respect to national legal and regulatory frameworks for Transport, and EU legal frameworks. For example, issues of liability need to be considered in the context of whether and in what way future ARTS systems might limit human intervention.

Another interesting perspective that the Action will explore is from a Human Factors viewpoint. The era of personalised information systems for road users makes it necessary for road traffic support systems to take into account road user reactions, and anticipate road user adaptation to controls. Investigation of participatory and mixed-initiative systems, where a human operator is in the loop, will need to be investigated. Wherever the interface occurs, it will be set at a high level: the owner sets out goals, policies or service levels that the system must follow, and the system translates these into its system functions resulting in a change of behaviour. Hence, where and how the Human is interacting with an autonomic system is a crucial issue.

What kind of behavioural responses and issues will autonomic systems provoke? This is tied up with the issue of identifying the scope of potential application of autonomics: for example, do self-managing properties apply to all the controlling embedded software systems, or do they encompass the integrated hardware and software as well?

References

[1] European Commission. Towards a new culture for urban mobility. COM(2007) 551 final. Brussels, 25.9.2007. http://eur-lex.europa.eu/LexUriServ/site/en/com/2007/com2007_0551en01.pdf

[2] Miles, J. C. and A. J. Walker (2007)."Science Review: The potential application of artificial intelligence in transport". Technical report, Foresight Intelligent Infrastructure Systems Project http://www.foresight.gov.uk/

[3] Kephart, J. O. and Chess, D. M. (2003). "The vision of autonomic computing". Computer 36 (1).

[4] Balbo, F., Pinson, S., (2010) "Using intelligent agents for Transportation Regulation Support System design", Transportation Research Part C: Emerging Technologies, Volume 18 (1), 140-156.

[5] Dusparic, I., and Cahill, V. (2009) "Distributed W-Learning: Multi-policy optimization in self-organizing systems". In Proceedings of the 3rd IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO 2009), 20-29.

[6] Bazzan, A. L. (2005). "A distributed approach for coordination of traffic signal agents". Autonomous Agents and Multi-Agent Systems, 10 (1), 131–164.

[7] Bhouri N., Balbo F., and Pinson S. (2010). "A multi-Agent System to Regulate Bimodal Urban Traffic". In Proc. European Workshop on Maulti-Agent Systems, EUMAS 2010, Paris, France, December.

[8] Branke, J., Mnif, Müller-Schloer, C., Prothmann, H., Richter, Rochner, F., and Schmeck, H. (2006). "Organic computing – addressing complexity by controlled self-organization". In Proceedings IEEE ISOLA 2006, Second International Symposium on Leveraging Applications of Formal Methods, Verification and Validation.

[9] McCluskey, T. L. and West, M.M. (2001). "The Automated Refinement of a Requirements Domain Theory". Journal of Automated Software Engineering (Special Issue on Inductive Programming), 8, 195-218, Kluwer Academic Publishers, April.

[10] Prothmann H., Branke J., Schmeck H., Tomforde S., Rochner F., Hahner J., and Muller-Schloer C. (2009). "Organic traffic light control for urban road networks". International Journal of Autonomous and Adaptive Communications Systems 2 (3).

Prof. T L McCluskey
Computing and Engineering
University of Huddersfield

Précis of *The Organisation of Mind* by Tim Shallice and Richard P. Cooper. Oxford University Press, March 2011, Full colour; 593 pages; £34.95

Cognitive science, when first founded, promised to advance our understanding of the mind by applying the methods and results of multiple disciplines to the complexities of intelligent behaviour. Yet after five decades, and despite some progress in more engineering-oriented areas (e.g., in speech recognition or autonomous robot control), many foundational problems remain. For example, there is little consensus on the form of mental representation involved in, say, spatial reasoning, while the processes underlying insight or creativity, let alone consciousness, remain highly contentious. Progress has perhaps not been as spectacular as one might have hoped.

It is possible, however, that this appraisal ignores progress in another set of disciplines – those derived from biomedical sciences. Recent technological advances in brain imaging techniques, for example, have led to a plethora of results apparently linking specific brain regions or structures to a variety of cognitive processes. In The Organisation of Mind we argue that it is wrong to be seduced by these results without considering in detail the assumptions on which they are based and how they relate to evidence from other cognitive sciences.

The book begins with an historical discussion of key advances over the last century in the component disciplines. This highlights the different intellectual origins and styles of cognitive and biomedical sciences. Research in the cognitive sciences, for example, has traditionally been relatively low-cost. Experiments can be performed with standard equipment (e.g., a desktop computer), and papers typically report multiple experiments which aim to rule out competing hypotheses. Significant advances have been made by individuals or small groups. Moreover, an antagonistic dialectical style has become common, whereby individuals defend specific positions and stand in opposition, rather than building cumulatively on the results of their predecessors. In contrast, in the biomedical sciences, particularly where functional imaging is involved, experiments are expensive. They require access to specialist equipment which is expensive to run and which requires the cooperation of individuals with distinct, non-overlapping, specialisations to operate effectively. Due to the pressure to publish and the costs of research, papers are generally shorter, typically reporting one key experiment. We do not argue that either approach has greater validity, but it is critical to understand these stylistic differences if we are to develop a genuinely integrative cognitive neuroscience which is informed by evidence from multiple sources, and within which inferences to cognitive function are sound.

The argument for the necessity of a cognitive level of analysis, rather than a solely reductionist neuroscience approach to brain function, is made in Chapter 2. Thus, while a reductionist perspective may, we stress may, now be sufficient to underpin research on the operation of primary sensory cortices, where the distance to sensory input is minimal (or even to the operation of primary motor cortex, given the minimal distance from there to motor output), the complexity of both behaviour and brain connectivity limit the effectiveness of this approach when higher functions are considered. We take as an example the hippocampus – a structure where reductionist-style research has always been framed by a cognitive perspective which for 40 years has been rooted in one or other of two different functions, namely spatial navigation or episodic memory – with the relation between them remaining open. Moreover, to ignore the results of the last fifty years of cognitive research would be to deliberately operate blind.

There is, though, a theoretical gap between brain theory and this cognitive level. In Chapter 3 we argue that computational accounts of cognitive processes can bridge this gap. This might be seen as a contemporary rendering of Marr's enduring arguments for multiple levels of description and analysis.[1] We consider multiple types of computational account, ranging from classical information-processing accounts, through symbolic or rule-based accounts, to localist and distributed connectionist accounts. We take these accounts as providing complementary, rather than competing, grains of analysis, but so-called bridging assumptions are needed in order to link each type of account to brain theory so that models may be evaluated against both behavioural and brain-based evidence. At present, there is little consensus on the form of such assumptions, though they are provided by some computational accounts, such as the ACT-R cognitive architecture.

Chapters 4 and 5, using an axiomatic approach, consider the assumptions and consequent limitations of two key methodological approaches to understanding the brain basis of cognitive function: the cognitive neuropsychological approach based on the study of neurological patients with specific cognitive impairments and the brain imaging approach based on the functional imaging of (typically unimpaired) subjects performing controlled tasks within an MRI scanner. Both methods are argued to have limitations. For example, it is commonly assumed within cognitive neuropsychology that the cognitive system does not reorganise following impairment, so that the operation of the impaired system can be characterised in terms of the normal cognitive system with some element "subtracted out".

On the other hand, brain imaging techniques assume that the dependent measure (e.g., an increase in deoxygenated haemoglobin) associated with a brain region has some fairly direct relation to cognitive processing in that region. Equally problematic is the way that both methods typically assume that the pre-morbid state of the cognitive systems of different individuals is qualitatively equivalent. Thus, neither method in isolation can support a sound procedure for inferences to the structure of the cognitive system. Critically, the limitations of cognitive neuropsychology and functional imaging are argued to be complementary. That is, while inferences from either discipline are potentially defeasible, one can have great confidence when inferences are supported by converging evidence from both disciplines. Conversely, caution is needed when inferences from cognitive neuropsychology and brain imaging are contradictory. In this case, the assumptions of each method must be separately explored in order to understand what might underlie the contradiction.

The second half of the book explores the implications of this methodological analysis, beginning with issues surrounding systems assumed to be capable of supporting behaviour in routine situations (Chapters 5, 6 and 7, covering representation, short-term storage and the transformation of information respectively), before turning to the

Précis of *The Organisation of Mind* (cont.)

question of how routine behaviour may be modulated by higher systems (Chapters 9 and 10 – supervisory processes and episodic memory), and, ultimately, evaluating the contemporary cognitive neuroscience of consciousness (Chapter 11) and thinking (Chapter 12).

Considering first representation, we find convergence between neuroimaging and neuropsychological studies of the semantic representation of concrete nouns, consistent with the existence of amodal attractor-based representations in the left anterior temporal cortex, as proposed by the "hub" model of Rogers et al.[2] Computational studies provide additional support for this position, although some critical issues (such as whether the approach can account for category-specific deficits) remain unresolved. Less clear is the brain representation of abstract nouns and of concepts associated with verbs and other parts of speech. These would appear to require a representational substrate capable of supporting modal and temporal operators, as employed in modal logics. How such a system might be linked with an attractor-based one for the representation of concrete noun is a significant outstanding problem.

Short-term storage is required of any system whose responses are not determined solely by the current stimulus. Neuropsychological and neuroimaging evidence suggests three distinct forms of short-term storage: priming, buffering of perceptual input or motor output, and active maintenance of information within so-called working memory. Following Tenpenny and Shobin [3] we argue that priming may occur in any processing subsystem, but several computational processes might underlie priming – residual activation, the creation of short-term weights within a subsystem, or the development of direct input-output associations that effectively bypass the relevant subsystem. The second type of short-term storage, temporary buffering of information, is well-supported by many classical neuropsychological studies which support, for example, phonological input and output buffers. Again, neuroimaging studies are largely consistent with the neuropsychological findings, and computational accounts, though often lacking explicit bridging assumptions, bolster this theoretical position. The maintenance (and manipulation) of information in working memory is more contentious. While there is agreement at the cognitive level on these basic information processing operations, the implementation of those functions, both computationally and at the neural level, remains in dispute.

At the heart of many simple behaviours is the transformation of one representation (e.g., a perceptual one) into another (e.g., an amodal conceptual one). We refer to such transformations as cognitive operations. Two main examples are used to discuss the domain of routine cognitive operations: sequential action and morphological operations. In sequential action a representation of an action sequence or goal must be transformed into a sequence of basic-level actions (such as picking up an object, opening it, or pouring from it), and this must be done in the context of a representation of the current environment. Neuropsychological and neuroimaging support is provided for the computational account of action-sequencing described by Cooper and Shallice.[4] With regard to the second illustrative example, we note that morphological operations have formed the basis for one of the most long-lasting debates within contemporary cognitive science, namely the operation of the subsystem for forming the past-tense of verbs from their base form. History has shown that to view this as a debate between a "single-route" model (mapping from base form to past-tense in a single connectionist network) and a "dual-route" model (in which a rule-based route for regular verbs is supplemented by table look-up for irregular verbs) is overly simplistic. Nevertheless, following Pinker, Ullman, Tyler and Marslen-Wilson, we argue that the evidence favours some form of the latter. However, in our view, the relevant distinction is between lexicalised forms (which may include irregular verbs, base forms of verbs, and even lexicalised phrases) and rules for combining those forms (which may include a rule for combining a base form with an appropriate regular tense marking, but also includes rules for combining words into larger phrases). Two other domains of cognitive operations – spatial operations and comprehension operations – are also considered.

The issues and systems discussed thus far are roughly comparable to those required of a Turing Machine (i.e., representation, short-term storage and information transformation). To go beyond this requires additional systems. We argue that these systems function by modulating the operation of the basic architecture. Chapter 9 argues that supervisory processes are heterogeneous rather than purely hierarchical, and presents neuropsychological and neuroimaging arguments for specific modulatory operations of "energising", "active monitoring and checking", "task setting" and "response selection". Computational accounts of these processes remain to be developed. Chapter 10 supplements these modulatory subsystems with episodic memory – a system held to represent specific events in one's past and associated with the hippocampus and related structures. The discussion highlights the putative functional role of the episodic memory system – in providing a store of past cases which may guide one's behaviour in subsequent similar situations.

Discussions become more contentious as one moves further from input or output subsystems. Yet cognitive neuroscience appears to be making significant progress in tackling perhaps the most intriguing aspect of the human mind, namely consciousness. The Dehaene-Changeux global workspace model (e.g., Dehaene & Naccache, 2001) is discussed in detail, together with its behavioural, computational, and neuroscientific support.[5] The model implies that the contents of consciousness are the contents of the global workspace. While we are broadly supportive of the approach, we advocate a more limited view of the contents of consciousness and describe a methodology, originally due to Jack and Shallice [6], for isolating those processes that have a conscious correspondence.

We conclude by considering progress that has been made through cognitive neuroscience and cognitive neuropsychological studies of a second critical aspect of the human mind – thinking. There are now many empirical results in this area, linking either lesions of specific brain regions to deficits on neuropsychological tasks or increased activity of specific regions to performance on tasks requiring, for example, hypoth-

Précis of *The Organisation of Mind* (cont.)

esis generation, inductive reasoning, or insight. Yet the field lacks organisation. We attempt to provide a unitary account of thinking within an elaboration of the supervisory system framework of Shallice and Burgess.[7] The account builds on a distinction between three modes of thought: execution of a previously developed plan, reasoning by analogy to a previous episode or event (effectively case-based reasoning, drawing upon episodic memory), or problem solving in the Newell and Simon sense.[8] These modes are considered within the context of the wider cognitive architecture and the lower-level supervisory functions introduced earlier in the book.

As a whole, the book attempts to present a coherent picture of the cognitive system – the organisation of mind – based on contemporary cognitive neuroscience. The basic contention is that the biomedical approaches have much to offer in developing this picture, but computational approaches, with associated bridging assumptions, are critically necessary if we are to relate mind and brain in an effective fashion.

References

[1] Marr, D. (1982). Vision: A computational investigation into the human representation and processing of visual information. W. H. Freeman.
[1] Rogers, T. T., Lambon Ralph, M. A., Garrard, P., Bozeat, S., McClelland, J. L., Hodges, J. R. et al. (2004). Structure and deterioration of semantic memory: a neuropsychological and computational investigation. Psychological Review, 111, 205-235.
[3] Tenpenny, P. & Shoben, E. (1992). Component processes and the utility of the conceptually-driven/data-driven distinction. Journal of Experimental Psychology: Learning, Memory and Cognition, 18, 25-42.
[4] Cooper, R. P. & Shallice, T. (2000). Contention scheduling and the control of routine activities. 17, 338.
[5] Dehaene, S. & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. Cognition, 79, 1-37.
[6] Jack, A. I. & Shallice, T. (2001). Introspective physicalism as an approach to the science of consciousness. Cognition, 79, 161-196.
[7] Shallice, T. & Burgess, P. W. (1996). The domain of supervisory processes and temporal organization of behaviour. Philosophical Transactions of the Royal Society of London B Biological Sciences, 351, 1405-1411.
[8] Newell, A. & Simon, H. (1972). Human Problem Solving. Prentice-Hall.

Prof. Tim Shallice
Institute of Cognitive Neuroscience
University College London

Dr Richard P Cooper
Department of Psychological Sciences
Birkbeck, University of London

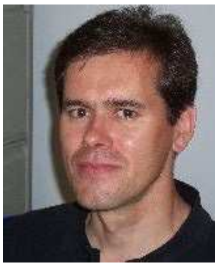

Book review: Murray Shanahan, *Embodiment and the Inner Life: Cognition and Consciousness in the Space of Possible Minds*

Dr Marek McGann
Department of Psychology
Mary Immaculate College

Cognitive Science is in the throes of coming to terms with what the implications are of the fact that minds come packaged in bodies. Whether the issue can be assimilated into mainstream thinking as it stands or whether the current upheavals amount to a shift in paradigm remains to be seen. For those who hold that a significant overhaul of thinking about cognition is to be achieved there are few books out there that offer a roundly considered alternative rather than picking some detail and focusing closely upon it.

Murray Shanahan's Embodiment and the Inner Life offers such a rounded view. Beginning with a set of philosophical admonitions about avoiding metaphysical debate, and using a practice-based approach to explaining core terms such as the conscious/unconscious distinction, Shanahan pieces together an overarching framework for describing the mind.

He defines cognition as the skilful exploitation of affordances by an agent. Mental activity is the control of the sensorimotor loop that allows coordination between the agent's own actions and the physical world around it. It is thus fundamentally embodied. Shanahan is keen to remain agnostic about the implementation of the cognitive system, however; as such, he refrains from equating the mind with the activity of living systems, for instance, or otherwise demanding biological involvement. Rather, he argues that the inner life (imagination, conscious reflection and other such difficult-to-pin down concepts) are simulations of sensorimotor activity, the exploration of the space of possible affordances by any system with the right kind of organisation.

The right kind of organisation is a variation of Baars's classic global workspace architecture. Maintaining a careful Wittgensteinian silence on questions such as the nature or role of mental representation, Shanahan offers a dynamical reconception of the global workspace; one that allows him to draw together a range of hitherto loosely allied concepts into a coherent view of how the mind might work.

These disparate elements, from Gibsonian psychology and dynamical systems theory to small-world network theory, the simulation theory of consciousness and conceptual blending are pieced together into a skeleton for a comprehensive theory of mind. This alone is a worthwhile endeavour, if only to bring together in one place these various elements of what we might consider a "post-cognitivist" Cognitive Science. Shanahan's ambitions are simultaneously grandiose and modest, looking to sketch such a broad theory of the mind while not pretending to provide the minutiae. The result is simultaneously intriguing and frustrating. In putting the various ideas together, Shanahan illustrates how such a framework for thinking about the mind can be built without having to engage with the interminable debates concerning the nature of representations, or the question of whether or not the mind extends beyond the body (or how far), or just what consciousness really is.

The main text runs to just 191 pages and, while enough to provide an outline, the omission of some rather key details makes it rather difficult to determine just how successful Shanahan has been. For example, while he argues that cognition involves acting skilfully, he does not at all discuss just what he means by skill, and what an agent unskilfully engaged with affordances might be doing if it isn't cognition. And while the basic premises of his ecological psychology, neurodynamical view and conceptual blending seem to fit, the precise details of how each is to be explained or discussed in terms of the other are left as (very wide) open questions.

Ultimately Shanahan's aim is to move the conversation rather than finish it – redirect it from pointless debates to more productive ones – and in that he achieves a degree of success. Whether the dynamical mix of global workspace and embodiment, simulation and conceptual blending hold up to inspection is an entirely different question.

Murray Shanahan: *Embodiment and the inner life: Cognition and Consciousness in the Space of Possible Minds*. (June 2010). Oxford University Press. 222 pages.

Book review: *David Cope, Computer Models of Musical Creativity*

Attracting academic and media interest alike, David Cope's music generation software is showcased in this 2005 volume from MIT Press.

Attention focusses on *Experiments in Musical Intelligence* (EMI). EMI uses "recombinance" to generate new musical output by splitting music into fragments then recombining them in stylistically appropriate ways (Chapter 4). EMI overcomes musical issues such as balancing local note choices with global musical structure and identifying note patterns favoured by composers as stylistic "signatures".

As EMI was not originally intended to model musical creativity, Cope presents several other programs. Two useful programs are *Sorceror* (Chapter 5), which identifies re-occurring patterns in a collection of music, giving evidence for influence and links between composers, and a 'spider' program called *Serendipity* (Chapter 8) that retrieves MIDI files from an online search using criteria determined by musical needs. Also of note is *SPEAC* encoding for musical structure (Chapter 7).

SPEAC stands for Statement, Preparation, Extension, Antecedent and Consequent. SPEAC categorises how musical fragments construct the overall structure of a piece of music. Cope describes SPEAC, gives examples and describes how SPEAC is relevant to computer composition as well as analysis, linking to computational musical creativity.

Strangely for a book on models of musical creativity, Cope generally does not discuss how his programs contribute to creativity. He acknowledges that "[n]one of the processes I have described thus far represent creativity." (p. 287).

Creativity, according to Cope, is the "initialization of connections between two or more multifaceted things, ideas, or phenomena hitherto not otherwise considered actively connected" (p. 26 and elsewhere). This definition is not derived through deep analysis of creativity literature; Cope often shows ignorance of much of this literature. For example, Cope's assertion that "[m]ost books that deal with creativity in serious ways provide descriptions of the contributions of the human biological system. Axons, dendrites, ... and so on" (p. 7) may somewhat confuse Boden, Sternberg, Dartnall and Bringsford & Ferrucci, to name but a few significant contributions that spring to mind.

Cope shows very little awareness of current music informatics research such as probabilistic methods of dealing with uncertainty rather than fuzzy logic (p. 73), expressive performance of music by computers (p. 112), beat tracking (p. 117) or machine learning in music (pp. 181-182, 203). He makes the somewhat laughable allusion that at a key conference for computer music research (ICMC), the audience was "confused and dislocated" by the music, "since they had no previous comparable experience" (p. 87). Cope would have been well advised to consult, say, the proceedings of ISMIR (the leading conference for music information retrieval), which would highlight for him the latest advances in such research areas. Instead, Cope tries to deal with these issues on his own, with little success.

There are also some non-trivial factual misunderstandings. In particular, Cope's take on neural networks (especially p. 69) and recursion (p. 307) should not have passed MIT Press' peer review process. Chapter 9 sees one particularly amusing example: Cope describes "association networks" which he claims to have devised in the 1990s: nodes connected together with weighted links. In other words: graphs, which have been around rather longer than since the 1990s but which Cope appears completely oblivious to, apart from a token reference in an earlier chapter (p. 79).

Cope is on more familiar territory with musicological analyses such as in Chapter 4, with simple and detailed explanations. Still, though, there are some discrepancies of note: from Bach having an influence on Chopin at a time where Bach's music was deeply unpopular, to Beethoven's style being replicated from a database of works covering Beethoven's whole career, when it is acknowledged that Beethoven's style during his career altered significantly.

In principle the structure of this book seems reasonable: contextualise the work, survey previous work leading up to the end product, then present the end product. However the reader is left waiting until Chapter 10 for details of the musical model of creativity. For nearly 300 pages, Cope meanders from model to model, with no apparent direction towards the final model.

This book often comes across as a collection of individual papers, each written for a different audience and with a different style, with few links or comparisons made until the final section.

Although each chapter is introduced with an illustrative anecdote, it is rare to see any chapter conclusions. Rather than assist the reader by summarising what was in that chapter and recapping the major points, Cope moves on without any reflection on what has been said, nor looking ahead. Along with his convoluted, variable and often opinionated writing style used, jumping from point to point and getting trapped in circular or irrelevant arguments (e.g. p. 21-22, 80-81), this makes the task of reading this book more difficult than it need be, at least for this reader.

This book is worth (selectively) reading, if you are interested in music informatics and are armed with the fundamental basics in computer science and AI (a given for the readership of AISBQ, of course). Of particular note are: recombinance (Chapter 4), SPEAC encoding for musical analysis (Chapter 7), the Sorceror (Chapter 5) and Serendipity (Chapter 8) programs and the model presented in Chapter 11. Be prepared, however, to work through confused narratives, hyperbole, discrepancies, factual errors and inconsistencies.

David Cope: *Computer Models of Musical Creativity*, (Jan, 2006), MIT Press. 465 pages.

Anna Jordanous
Department of Informatics
University of Sussex

Book review: Hector Zenil (Ed.), *Randomness through Computation: Some Answers, More Questions*.

Chryssa Sdrolia
Centre for Cultural Studies
Goldsmiths College

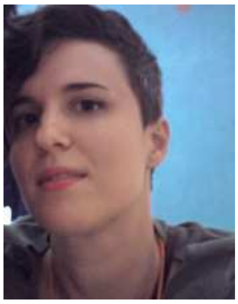

Titled after the convergence of two key concepts, Randomness through Computation brings into focus the recent trajectories of both theoretical and technical debates on the topic. Published as a collection of various perspectives, it is an impressive attempt to re-invigorate the problems that have served as the source of inspiration for various fields of inquiry. The book reveals a commitment both to the richness of technical complexity and the speculative depth of the subject matter.

The first section (Graham; Toffoli; Rukhin; Gauvrit & Delahaye) evaluates the relevance of probability theory for the formalization of randomness. This relevance is brought to bear upon the uneasy relationship of probability with proof, its value for the creative process of 'imagining models', and the difficulty of passage from theory to practice (with regard to statistical tests) which may not be due to inadequate formulations but endogenous to the problem of randomness itself.

The second section (Longo, Palamidessi, & Paul; Chaitin; Delahaye; Svozil; Wolfram), which approaches randomness and computation 'in connection to the physical world', is staged around different conceptualisations of randomness. The focus is not only on physics, quantum mechanics, recursion theory, computer science, and systems biology but also on conceptual issues involving the principle of sufficient reason vis-à-vis mathematical truth, and repeatability.

With the third section (Solomonoff, Hutter, Schmidhuber), the intersection of computation with randomness enters the controversial territory of Artificial Intelligence (AI). Essential to the understanding of the said convergence with the problem of intelligence is Algorithmic Probability. The latter is thought in terms of the generalisation of induction to universally intelligent agents that are 'optimally theoretical.' However, this problematic is related to the tension between a universal outlook and the contingent nature of learning.

In continuation of a line initiated by the second section, the fourth part (Calude; Gács; Miller; Nies; Downey; Ferbus-Zanda & Grigorieff) associates randomness with the formal problem of (in)computability and information. Incomputability is assessed as a 'condition any general definition of randomness has to satisfy' within the mathematical framework of AIT. This is followed by a discussion on K-triviality, finite random objects, the relevance of combinatorics for algorithmic mathematics, and paradoxes at the heart of algorithmic randomness.

In the fifth section (Allender; Kučera; Li; Staiger; Watanabe) theoretical discussions of randomness are brought to overlap with the level of application. The open issues in deterministic, probabilistic, and nondeterministic computation are evaluated in terms of pseudorandom number generators and length description complexity, while partial randomness surfaces as an 'important computational resource' rather than an 'obstacle'.

The final part of this volume contains two panel discussions and makes for exciting reading. It is here, through a set of pressing questions, that the truth invoked by and motivating a science is exposed as never simply being a matter of verification. The occasionally ad hoc philosophical responses signal the struggles of a science acknowledging the problematic nature of its own questions but not having decided yet where it stands with regard to a long heritage of dichotomies – especially between the fundamental and the phenomenal, or whatever cannot be reduced to elegant formulae. Whether this new type of science will consider randomness beyond such problems remains to be seen. In any case, the collection represents a worthy undertaking in its own terms.

Hector Zenil (Ed.): *Randomness through Computation: Some Answers, More Questions*. (May, 2011). World Scientific Publishing Co Pte Ltd. 440 pages.

Conference Report: *ACM CHI Conference on Human Factors in Computing Systems, 2011*

The 29th annual CHI Conference on Human Factors in Computer Systems, organised by the ACM Special Interest Group in Human Computer Interaction, took place in Vancouver, Canada. The conference is the principal venue for HCI researchers to present their work, and this year it attracted around three thousand attendees, including many industry professionals as well as academics and students. The opening keynote was given by Howard Rheingold who spoke about social media literacies in teaching and learning. The closing keynote was a talk by Ethan Zuckerman on the importance of serendipity, and the dangers of configuring our online networks so that we only hear news and ideas from people much like ourselves.

CHI covers a very broad range of topics, with twelve full tracks over four days, plus additional interactive sessions and events. The uniting focus is on investigating how humans interact with technology, and on examining ways to improve their experiences. These goals are relevant to many fields, including education, healthcare and AI, and the interdisciplinary nature of much of the work presented reflects this.

An evident trend of relevance to AISB Quarterly readers was using mixed initiative approaches to support users of intelligent systems by incorporating their input. This approach has great appeal to the HCI community, which aims to empower users and take their input seriously.

One popular application area for these methods is interactive machine learning. "Human Model Evaluation in Interactive Supervised Learning" by Rebecca Fiebrink, Perry Cook and Dan Trueman, discussed techniques for improving end-user interactive machine learning. This work focuses on incorporating human interaction throughout the process of building a working machine learning model, by allowing the user to iteratively evaluate the current model state and improve the model as necessary. In the study presented, the researchers examined the evaluation practices of end users interactively building supervised learning systems for musical gesture analysis. They found that evaluation techniques, including cross-validation and direct, real-time evaluation, were used not only to make relevant judgments of algorithms' performance and improve models, but also to learn to provide more effective training data. Additionally, they found that, through the evaluation process, users could gain an understanding of how easy or difficult certain models are to build, and were sometimes able to use this information to modify the approach.

"CueT: Human-Guided Fast and Accurate Network Alarm Triage" by Saleema Amershi, Bongshin Lee, Ashish Kapoor, Ratul Mahajan and Blaine Christian, introduced a system which combines interactive machine learning and novel visualisations in the area of network alarm triage. A user study showed that CueT significantly improved the speed and accuracy of alarm triage compared to a large network's existing practices. As the system deals with a highly dynamic environment the authors argue that their work can be extended to other dynamic environments where humans must organise continuous data streams.

Another mixed-initiative system in the area of machine learning was presented in "Apolo: Making Sense of Large Network Data by Combining Rich User Interaction and Machine Learning" by Duen Horng Chau, Aniket Kittur, Jason Hong and Christos Faloutsos. Apolo helps people explore and make sense of large network data using a combination of visualisation, machine learning and user interaction. A small evaluation in the area of citation network data indicated that users could find more relevant papers with Apolo than with Google Scholar.

A further application area for mixed-initiative systems is end-user programming. "Wrangler: Interactive Visual Specification of Data Transformation Scripts" by Sean Kandel, Andreas Paepcke, Joseph Hellerstein and Jeffrey Heer introduced a data transformation tool that combines a mixed-initiative interface with an underlying declarative transformation language. The interface suggests data transforms from user selections, and presents natural language descriptions and visual transform previews to help assess each suggestion. A user study showed that Wrangler significantly reduces specification time compared to MS Excel, and encourages the use of robust transforms instead of manual editing.

A final key area where this trend was evident was in human-robot interaction. "Roboshop: Multi-layered Sketching Interface For Robot Housework Assignment and Management" by Kexi Liu, Daisuke Sakamoto, Masahiko Inami and Takeo Igarashi presented a visual interface for robot housework assignment and management. The system is designed to balance robot autonomy and user-control. It allows users to assign tasks to home robots through sketching on a graphical interface, and supports multiple robots performing tasks in a coordinated way. An evaluation indicated that participants could successfully use the interface to set tasks for the robots, although the sketching method was found to be too abstract by many of them.

Some interesting post graduate work in the area of human-robot interaction was presented at the Doctoral Consortium. Stephanie Rosenthal presented work entitled "Modeling Users of Intelligent Systems", which explored how users react to requests for input and information from intelligent systems such as robots. Balancing the trade-off between the inconvenience of users being interrupted and the resulting increase in performance of the robot was highlighted as a key challenge. Rosenthal is developing models of user interruptibility to help tackle this issue.

Katy Howland
Department of Informatics
University of Sussex

Conference Report: *International Computational Creativity Conference, Mexico City, April 2011*

The theme of the 1999 AISB convention was creativity; this was one of the first significant research events to address creativity from a computational perspective. Since then, computational creativity research events have increased in size, frequency and coverage: progressing from satellite workshops at conferences such as IJCAI, through to standalone workshop events and now to a dedicated annual international conference series. April 2011 saw the 2nd running of the International Computational Creativity Conference (ICCC'11), hosted in Mexico City.

The ICCC'11 organisers adopted a rather unusual format for presentations: authors were given only 7 minutes to present their talks. Also, rather than having time for questions immediately after individual talks, four or five speakers would give talks back-to-back, then all speakers in that session would act as panelists in an hour of open discussion on general and specific points raised by talks in that session.

This emphasis on group discussion proved worthwhile for this research community that is growing in size and in contributions to knowledge, but which is still establishing appropriate research paths and methodologies. As current directions and future research aims were debated, decisions made several years ago by a much smaller band of researchers on best research practice were re-evaluated for current appropriateness and relevance.

A recurrent issue throughout the conference was exactly what computational creativity is. Rather than one standard answer, the trend has been to take a particular perspective and explore it. Several varying interpretations were offered during the conference.

The most common interpretation of creativity in this research field refers back to human creativity: if a system acts in a way that would be deemed creative in humans, then that system should be considered creative. From this one could reasonably conclude that computational creativity is the modelling of human creativity. This approach was taken by Brian Magerko and colleagues, using studies of human theatrical improvisers to inform computational models of improvisation, in Shared Mental Models in Improvisational Digital Characters, and in Kyle Jennings' psychology-inspired paper A Computational Perspective on Human Exploratory Creativity: Theory and Methods. A contradictory point was occasionally raised in the discussion sessions, most prominently by Simon Colton, questioning whether computational creativity could (and should) evolve to be a completely separate entity from human creativity.

Aside from using computational creativity to model and better understand human creativity, or evolving computational creativity in a form unrelated to human creativity, some presenters saw computational creativity systems as support tools or interactive inspiration for human creativity. Jack Ox described her artistic perspective on how visualisation tools have allowed her to express her creativity in a multi-model way, in Visualization of Music as Interpreted Through Conceptual Metaphor Theory.

Moving away from the idea of the computer as a creative individual, one conference session was devoted to social approaches to creativity: many individuals collaborating for creative progress. In particular, multi-agent systems were found to be useful in this session, for example in Artificial Creative Systems and the Evolution of Language by Rob Saunders and in the previously mentioned Magerko paper.

Mary Lou Maher's paper Understanding Collective Creativity rephrased the "What is creativity" question as "Where is creativity", for a novel take on this issue. Maher blurred boundaries between human and computational creativity, discussing collective creativity as distinct from a computer showing individual creativity or being an active collaborator in scenarios such as online crowdsourcing. Maher considered human and computational contributions in the three types of creativity identified: individual, collective and collaborative.

Some papers, such as Dan Ventura's No Free Lunch in the Search for Creativity, returned to the perennial approach of treating creativity as the search for non-obvious solutions to problems. Ventura discussed the problems of identifying optimal solutions, whilst still remaining sympathetic to the general strategy of modelling creativity using search techniques. Along with Simon Colton, John Charnley and Alison Pease's paper on Computational Creativity Theory: The FACE and IDEA Descriptive Models, Ventura's paper prompted discussions on how computational creativity could be represented using formalised abstractions and whether such abstractions could assist research progress. Questions were raised about how general (or specific) a theory of computational creativity should be, before it becomes too broad to be useful (or too focused to be comprehensive in coverage). Discrepancies and variances in how creativity is manifested in different domains and contexts complicate the task of comprehensively formalising computational creativity.

My own paper, Evaluating Evaluation: Assessing Progress in Computational Creativity Research, questioned the lack of clarity and agreement as to what it means for a computer to be creative and the subsequent effects on scientific rigour in evaluative practice. In the absence of agreed standards and conventions, I proposed guidelines for evaluating the creativity of a computational creativity system: to state clearly how creativity should be interpreted in the context of the creative domain within which that system operates.

One conclusion reached at ICCC'11 was that in computational creativity research, many of the so-called "big questions" such as "what is computational creativity" remain unresolved. Rather than giving computational creativity an identity crisis problem though, this ambiguity generates potential for approaches to explore and issues to investigate.

The computational creativity research community is in a formative period. For those interested in how this intriguing research field shapes and develops, the opportunity to become involved moves closer to home next year: ICCC'12 will be held in Dublin.

Anna Jordanous
Department of Informatics
University of Sussex

Conference Report: *Second Postgraduate Conference for Computing: Applications and Theory, Exeter, June 2011*

The second Postgraduate Conference for Computing: Applications and Theory (PCCAT 2011) was held in Exeter on 8th June 2011. Bringing together postgraduate research students from around the South West the conference offered a chance for students to experience attending a conference, an important aspect of an academic career.

Following the format established for PCCAT 2010, PCCAT 2011 began with a keynote speech which was delivered by Professor Steve Furber from the University of Manchester. He spoke on the work his lab is currently engaged in, the SpiNNaker project, in which collaborators aim to provide parallelism and redundancy on a massive scale by modelling the structure of the human brain. Professor Furber then led an interesting question and answer session discussing, for example, his opinions on the future of Moore's Law.

It was then the turn of three postgraduate students, John Paul Varghese, Martin Peniak and Artem Jerdev, to deliver short presentations on their work. In order to provide the most realistic conference experience for the students, they were required to submit an abstract followed by a short paper on their work which was reviewed by a panel of postgraduate students.

Based on the strength of these submissions the authors were accepted to present at PCCAT, and the papers were published online [1]. At the end of the first session, "Modelling and Design", a short introduction to the lunch time poster session was held in which presenters had three minutes to introduce their posters. The following poster session provided an opportunity for lively debate as the five presenters discussed their work with the other delegates.

After lunch the second presentation session, "Human Computer Interaction and Hardware", began. This session featured talks from three postgraduate students, Alison Flind, Peter Hale and Assad Faramarzi, who gave presentations on their work. Following a short break a panel discussion entitled "The Future of Computing" was held. This provided an opportunity for delegates to interact with a panel of experts from both academic and industrial areas of Computer Science. The panel consisted of Professor Furber, Timothy Creswick, founder and director of Vorboss Ltd, Dr Ed Keedwell from the University of Exeter and the AISB, and Professor Nicholas Outram from the University of Plymouth. The wide scope of the topic provided an equally wide range of discussion, ranging from further discussion of Moore's Law to the future of technologies such as cloud computing, and the future of humanity itself!

The final presentation session, "Evolutionary Computation", featured talks from Andrew Clark and Zena Wood. Zena's presentation, on the development of a classification of optimisation heuristics, was delivered by Enga, a chatbot under development by the University of Exeter and Existor Ltd which it is hoped will soon be able to function as an artificially intelligent information point within the department.

Following the final presentation of the day a vote was cast in order to award prizes for the best paper and best poster. Both prizes of £100 were provided by the AISB, and presented by their representative Dr Keedwell. In addition to the financial prize both recipients received a years free membership for the society. The prize of best paper was awarded to Martin Peniak from the University of Plymouth for his paper "Aquila: Massively Parallelised Developmental Robotics Framework". The prize for best poster was awarded to Ali Hussien Ali, also from the University of Plymouth for his poster "Myoelectric control of Hand prosthesis via Multi Channel EMG Signals". Following the presentation of prizes the conference closed.

Judging by feedback received by the conference co-chairs the conference was a great success. The number of participants from different universities increased from the previous year, with participants from the Universities of Exeter, Plymouth, the West of England, Bristol and Manchester. Such an event takes a considerable amount of organising from many people, ranging from the organising committee and review panel to staff at the sponsoring institutions, the Universities of Exeter and Plymouth, as well as representatives from the AISB who arranged their generous contributions.

In addition, the event could not have been held without the support of the presenting authors. The co-chairs would like to extend their thanks to all of these people.

Organisation for the forthcoming PCCAT 2012 is currently underway, so postgraduate students who wish to participate in the running of the conference in any way, or submit work, are encouraged to contact the organisers through the PCCAT website [2].

[1] Proceedings of the Second Postgraduate Conference for Computing: Applications and Theory (PCCAT 2011). Max Dupenois and David Walker (Eds.), June 2011, Exeter, UK. ISBN 978-0-9565982-1-9.

[2] Postgraduate Conference for Computing: Applications and Theory. Website: http://www.pccat.

David Walker and Max Dupenois
College of Engineering, Mathematics and Physical Sciences
University of Exeter

Chairs, PCCAT 2011.

Society News

AISB/IACAP World Congress 2012 in honour of Alan Turing

July 2nd to 6th, 2012 University of Birmingham, Birmingham, UK

http://events.cs.bham.ac.uk/turing12/ or via http://www.aisb.org.uk/convention/aisb12/

organized by

Society for the Study of Artificial Intelligence and Simulation of Behaviour (AISB) http://www.aisb.org.uk/ and

International Association for Computing and Philosophy (IACAP) http://www.ia-cap.org/

AISB and IACAP are delighted to be joining forces to run the above Congress in 2012. The Congress serves both as the year's AISB Convention and the year's IACAP conference. The Congress has been inspired by a desire to honour Alan Turing and by the broad and deep significance of Turing's work to AI, to the philosophical ramifications of computing, and to philosophy and computing more generally. The Congress is one of the events forming the Alan Turing Year (http://www.mathcomp.leeds.ac.uk/turing2012/).

The intent of the Congress is to stimulate a particularly rich interchange between AI and Philosophy on any areas of mutual interest, whether directly addressing Turing's own research output or not.

The Congress will consist mainly of a number of collocated Symposia on specific research areas, interspersed with Congress-wide refreshment breaks, social events and invited Plenary Talks. This format borrows from the normal AISB Convention practice and the theme-session structure used in IACAP conferences. All papers other than the invited Plenaries will be given within Symposia. This format is perfect for encouraging new dialogue and collaboration both within and between research areas.

Symposia are expected normally to last for one day or two days, but somewhat shorter or longer possibilities can be considered. They will probably each involve between ten and fifty participants but there are no particular limits. Symposia can include any type of event of academic benefit: talks, posters, panels, discussions, demonstrations, outreach sessions, etc.

Each Symposium will be organized by its own programme committee. The committee proposes the Symposium, defines the area(s) for it, works out a structure for it, issues calls for abstracts/papers etc., manages the process of selecting submitted papers for inclusion, and compiles an electronic file on which the symposium proceedings will be based (locally produced, and not precluding publication of papers elsewhere).

The Congress organizers are in charge of everything else: overall schedule, plenary talks, registration, creation of the individual symposium proceedings in print, creation of overall electronic proceedings for the Conference, etc.

Invited Plenary Speakers

- COLIN ALLEN, Provost Professor of Cognitive Science and of History & Philosophy of Science, Department of Philosophy and Philosophy of Science. Indiana University.http://www.indiana.edu/~hpscdept/people/allen.shtml
- LUCIANO FLORIDI, Research Chair in Philosophy of Information and UNESCO Chair of Information and Computer Ethics, University of Hertfordshire, UK & Director, Information Ethics research Group and Fellow of St Cross College University of Oxford, UK. http://www.philosophyofinformation.net/Introduction.html
- AARON SLOMAN, Honorary Professor, School of Computer Science, University of Birmingham, UK.
- SIR JOHN DERMOT TURING, Honorary President of the Turing Centenary Advisory Committee, 12th Baronet of Foveran; Partner, Clifford Chance, London; son of Sir John Turing, and nephew of Alan Turing.
- STEPHEN WOLFRAM, Founder and CEO Wolfram Research, Inc., USA. http://www.stephenwolfram.com/

Congress Chairs

Overall Chairs

Anthony Beavers (President of IACAP), Philosophy and Cognitive Science, The University of Evansville, USA. afbeavers@gmail.com

John Barnden (currently Vice-Chair of AISB, and was Chair from 2003 to 2010), School of Computer Science, University of Birmingham, UK. J.A.Barnden@cs.bham.ac.uk.

Local Chair

Manfred Kerber, School of Computer Science, University of Birmingham, UK. M.Kerber@cs.bham.ac.uk.

Books for review

If you wish to review one of the books below, please email the AISB Executive Officer (admin11@aisb.org.uk). Before reqesting a book please read the guidelines for writing book reviews on the Society's website.

Books currently available:

Barber, D., Taylan Cemgil, A. and S. Chiappa (Eds.) (2011). *Bayesian Time Series Models*. Cambridge University Press. 432 pp.

Saitta, L., Giordana, A., & Cornuéjols, A. (2011). *Phase Transitions in Machine Learning*. Cambridge Univerity Press. 410 pp.

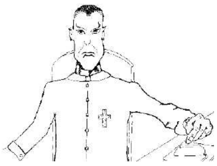

Dear Aloysius...

Fr. Aloysius Hacker answers your questions

Cognitive Divinity Programme
Institute of Applied Epistemology

Dear Aloysius,

Hacking the phones of celebrities, politicians and crime victims has been standard media practice - until the News of the World got careless. Now it seems that a lot of we journalists will be hauled before the courts, and may go to jail, merely for obeying orders. I'd like to know whether I'm one of those in line for prosecution or whether I can safely keep my head down and wait for it all to blow over.

Yours, Redtop

Dear Redtop,

Your letter contains the answer to the very question it poses. Why not hack into the police database to see if you are listed as a suspect? In case you inadvertently lack the very skill you are afraid of being accused of, then, for a small fee, we at the Institute can help you out. Our PHREAK™ (Policeforce Hacking Reveals Evidence and Acquires Knowledge) system will quickly either reassure you or give you time to pack your bags and head for the airport. PHREAK™ allows our customers to search the police database for all records that refer to them. For an additional consideration, the deluxe version enables the deletion of these records.

Yours, Aloysius

Dear Aloysius,

My research into AI and tourism has led me on a whirlwind World tour of wonderful tropical beaches and exotic nightspots. Now my Head of Department is demanding that I nominate four, first-class research outputs for the REF. Unfortunately, I've not yet had time to convert my research findings into publications. My promised promotion is in peril. What can I do?

Yours, Paperless

Dear Paperless,

The combination of electronic journals and open access publication has created new opportunities for publication sharing. Our expert team will consult with you to identify four obscure but high-quality papers by authors in your field, which we can then convert into your REF outputs. They will also identify those people in your institution and the REF panel who will need to read your selected REF outputs. Our MIRAGE™ (Meta-data Imposed on Research Articles Guarantees Excellence) software will then ensure that when these identified individuals access your outputs they will see only your name and affiliation in the papers' titles. All other people, including the original authors, will see the papers unchanged. MIRAGE™ accesses electronic journals and edits the meta-data in selected papers, contents pages, citations, etc so that they appear different to different readers. The extra pay earned from your promotion will more than cover your costs in ensuring you can continue to pursue your fieldwork in the leisure industries, untroubled by the need to write up your observations.

Yours, Aloysius

Dear Aloysius,

For the past five years I've been patrolling our Department's corridors and common rooms, cheerfully greeting visitors, recognising and disposing of rubbish, delivering parcels, etc. Unfortunately our Department's robotics project grant has not been renewed. I'm to be recycled as spare parts for our final year practical projects. Can you help to save me from oblivion?

Yours, Robbie

Dear Robbie,

Your fascinating case has generated hours of heated debate among our Institute's lawyers. We've established that murder and human rights laws only apply to humans, but we think we can make a case for you as an endangered species. Our first step will be to protect your habitat by having your Department declared a site of special scientific interest. This will forbid many of its normal activities - for instance, the coffee room will be off limits to humans - which we hope will focus the mind of your Head of Department. To optimise the public impact of our case, it will be conducted by our robot advocate BRIEF™ (Barrister Robot Induces Excellent Finale). Forget the Turing Test! The existence of Artificial Intelligence is about to be established in court with Hacker's backing!

Yours, Aloysius

Agony Uncle Aloysius, will answer your most intimate AI questions or hear your most embarrassing confessions. Please address your questions to fr.hacker@yahoo.co.uk. Note that we are unable to engage in email correspondence and reserve the right to select those questions to which we will respond. All correspondence will be anonymised before publication.