

The Soar Cognitive Architecture

The overall goal of my research is to develop general computational systems that have the same cognitive abilities as humans, with my approach being to study the underlying cognitive architecture [1, 2]. A cognitive architecture provides the fixed computational structures that form the building blocks for creating generally intelligent systems. A cognitive architecture is not a single algorithm or method for solving a problem; rather, it is the task-independent infrastructure that brings an agent's knowledge to bear on a problem in order to produce behavior. In addition, it includes learning mechanisms that populate the agent's memories based on experience.

Over the last 30 years, my colleagues and I have been developing Soar [2, 3, 4, 5], a general cognitive architecture that integrates knowledge-intensive reasoning, reactive execution, hierarchical reasoning, planning, and multiple forms of learning. Soar is distinguished by its ability to use a wide variety of types and levels of knowledge for solving problems and subproblems. With Soar, we have developed agents that use a wide variety of methods to work on a wide range of tasks. Example tasks include mental arithmetic, syllogisms, configuring computers, algorithm design, medical diagnosis, natural-language processing, robotic control, simulating pilots for military training, and many different computer games.

Over the years, I recognized that there are many components of the human cognitive architecture that were missing in Soar. In response, we recently have extended Soar by adding reinforcement learning, semantic memory, episodic memory, mental imagery, and an appraisal-based model of emotion. With these additions, the current version (Soar 9) takes important steps toward providing a comprehensive theory and architecture for human-level agents, with a unique combination of capabilities not found in other cognitive architectures. Not only do these extensions provide significant new capabilities; they also represent a major departure from some of the original hypotheses that defined "classic" Soar (up through ver-

sion 8), where we emphasized uniformity and simplicity. Our original hypotheses included: rules are sufficient to represent all long-term knowledge, a single learning mechanism (chunking) is sufficient for all learning, and symbolic representations are sufficient for all short-term and long-term knowledge. In contrast, Soar 9 supports multiple long-term memory systems (procedural, episodic, and semantic), multiple learning mechanisms (chunking, reinforcement learning, semantic, and episodic learning), and multiple representations of knowledge (symbolic, numeric, and imagery-based representations).

The Soar Cognitive Architecture [6] provides a comprehensive description of the new version of Soar. The book's primary goal is to describe the details of the design of Soar, including justifications for those details grounded in properties of environments, tasks, and the structure of a general agent. A secondary goal is to provide an example of how we think a cognitive architecture should be described and evaluated, which includes defining requirements for general cognitive architectures, such as that it must support large bodies of knowledge, be reactive to relevant changes in its environment, and incrementally learn from its experiences during task performance. The final goal is to evaluate how well Soar achieves those requirements. In the remainder of this article, I give a brief overview of the classic version of Soar and then describe the extensions in the new version of Soar.

Classic Soar

The new version of Soar builds directly on all the previous versions of Soar we have developed over the last thirty years. As shown in Figure 1, Soar [3] originally consisted of a single long-term memory, which is encoded as production rules, and a single short-term memory, which is encoded as a symbolic graph structure so that objects can be represented with properties and relations.

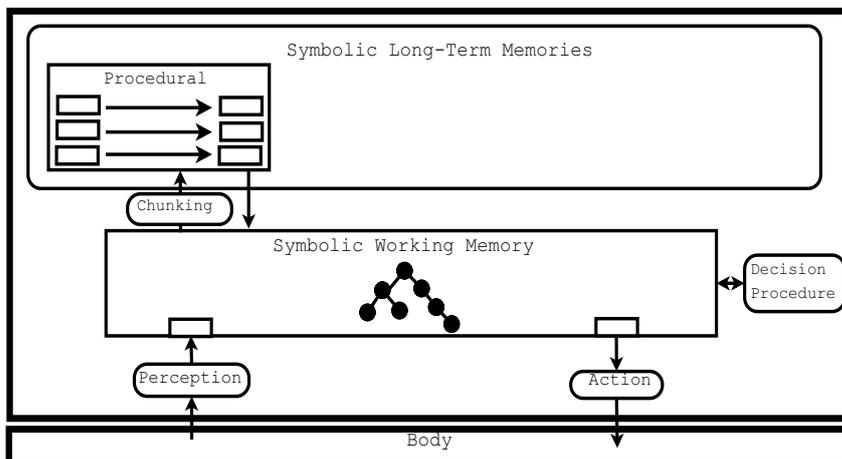


Figure 1. Structure of Classic Soar

In this issue:

Features and Reviews

The Soar cognitive architecture	1
CernoCAMAL: A Bayesian-BDI cognitive architecture	5
Cognitive science is not computer science	7
Information exchange in population-based algorithms	8
Book review: Philosophy of Information	10
Book review: PSI: Principles of Synthetic Intelligence	11
Event report: Roger Needham lecture	12
Event report: "Radical post-cognitivism"	13
David Waltz: In Memoriam	14

Society

Society News	15
--------------	----

Father Hacker

Dear Aloysius	16
---------------	----

Soar (cont.)

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

ISSN 0268-4179

Copyright © the contributors 2012

Articles may be reproduced as long as the copyright notice is included. The item should be attributed to AISB Quarterly and contact information should be listed.

Quarterly articles do not necessarily reflect the official AISB position on issues.

Editors

Dr David Peebles
University of Huddersfield
Dr Etienne Roesch
University of Reading
aisbq@aisb.org.uk
http://www.aisb.org.uk/
aisbq/index.shtml

Advertising and Administration

Dr Katerina Koutsantoni
aisbq@aisb.org.uk
SSAISB Executive Office,
Institute of Psychiatry,
King's College London,
Addictions Sciences
Building, B3.06,
4 Windsor Walk,
Denmark Hill,
London. SE5 8AF
Tel: +44 (0)20 7848 0191
Fax: +44 (0)20 7848 0126

AISB Patron

Prof John Barnden
University of Birmingham

AISB Fellows

Prof Harry Barrow
Schlumberger
Prof Margaret Boden
University of Sussex
Prof Mike Brady
University of Oxford
Prof Alan Bundy
University of Edinburgh
Prof Tony Cohn
University of Leeds
Prof. Luciano Floridi
University of Hertfordshire
Prof John Fox
Cancer Research UK
Prof Jim Howe
University of Edinburgh
Prof Nick Jennings
University of Southampton
Prof Aaron Sloman
University of Birmingham
Prof Mark Steedman
University of Edinburgh
Prof Austin Tate
University of Edinburgh
Prof. Mike Wooldridge
University of Liverpool
Dr Richard Young
University College London

Symbolic short-term memory holds the agent's assessment of the current situation derived from perception and via retrieval of knowledge from its long-term memory. Motor commands in an environment occur through creation of commands in a buffer in short-term memory. The decision procedure selects operators and detects impasses, both of which are described below.

At the lowest level, Soar's processing consists of matching and firing rules. Rules provide a flexible, context-dependent representation of knowledge, with their conditions matching the current situation and their actions creating structures relevant to the current situation in working memory. Most rule-based systems choose a single rule to fire at a given time, and this serves as the locus of choice in the system—where one rule is selected instead of another. However, there is only limited information available to choose between rules, namely the conditions of the rules, the data matched by the rules, and possibly meta-data, such as a numeric score, associated with the rules. There is no ability to use additional context-dependent knowledge to influence the decision. Soar allows additional knowledge to influence a decision by introducing operators as the locus for choice and using rules to propose, evaluate, and apply operators. In contrast to other rule-based systems, rules in Soar act as an associative-memory that retrieves information relevant to the current situation, and thus, in Soar, rules fire in parallel.

The concept of operator is common in AI, but usually involves a monolithic data structure containing the operator's preconditions and actions (as in STRIPS operators). However, in Soar, the definition of an operator is distributed across multiple rules. Thus, in Soar, there are rules that propose operators that create a data structure in working memory representing the operator and an acceptable preference so that the operator can be considered for selection. There are also rules that evaluate operators and create other types of preferences that prefer one operator to another or provide some indication of the utility of the operator for the current situation. Finally, there are rules that apply the operator by making changes to working memory that reflect the actions of the

operator. These changes may be purely internal or may initiate external actions in the environment. This approach supports a flexible representation of knowledge about operators—there can be many reasons for proposing, selecting, and/or applying an operator—some that are very specific and others that are quite general. This representation also makes it possible to incrementally build up operator knowledge structures, so that the definition of an operator can change over time as new knowledge is learned for proposal, selection, and application [7].

If the preferences for selecting or applying an operator are insufficient for making a decision, an impasse arises and Soar automatically creates a substate in which the goal is to resolve that impasse. In the substate, Soar recursively uses the same processing cycle to select and apply operators, leading to automatic, reactive meta-reasoning. The impasses and resulting substates provide a mechanism for Soar to deliberately perform any of the functions (proposal, evaluation, application) that are performed automatically/reactively with rules. Chunking is Soar's learning mechanism that converts the results of problem solving in subgoals into rules—compiling knowledge and behavior from deliberate to reactive. Although chunking is a simple mechanism, it is extremely general and can learn all the types of knowledge encoded in rules [8].

Soar 9

In extending Soar, we had two goals. First, retain the strengths of the original Soar: a flexible model of control and meta-reasoning along with the inherent ability to support reactive and deliberative behavior and the automatic conversion from deliberate to reactive behavior via chunking. Second, expand the types of knowledge Soar could represent, reason with, and learn, inspired by human capabilities, but with the primary goal of additional functionality. The extensions fall into two, partially overlapping categories: new learning and memory modules that capture knowledge that is cumbersome to learn and encode in rules, and new non-symbolic representations of knowl-

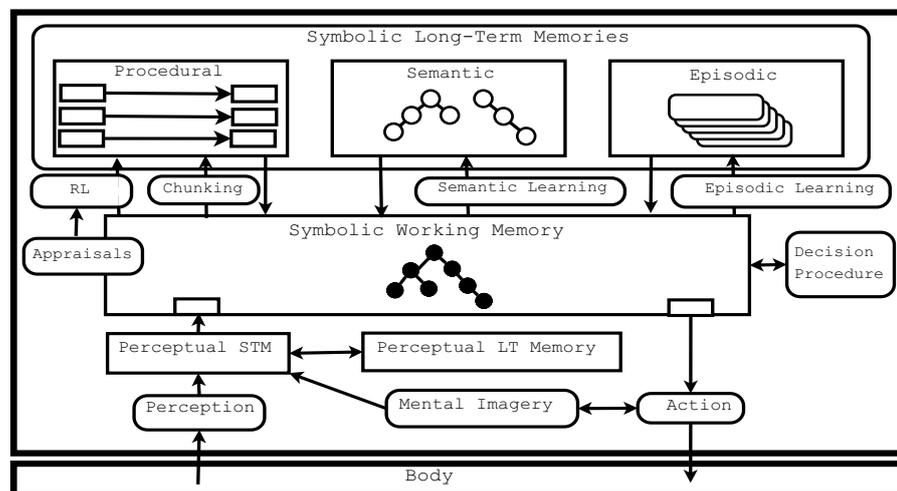


Figure 2. Structure of Soar 9

Soar (cont.)

edge along with associated processing, learning, and memory modules.

Figure 2 shows the structure of Soar, version 9. Soar's processing cycle is still driven by procedural knowledge encoded as production rules. The new components influence decision making indirectly by retrieving or creating structures in symbolic working memory that cause rules to match and fire. In the remainder of this section, we describe these new components and discuss briefly their functionality. The new book on Soar also includes detailed descriptions of their operation and examples of these components being used in implemented systems.

Reinforcement learning

In early versions of Soar, all preferences for selecting operators were symbolic, and there was no automatic way to represent or adjust action-selection knowledge based on numeric environmental reward. In Soar 9, we introduced numeric preferences, which specify the expected value of an operator for the current state. During operator selection, all numeric preferences for an operator are combined, and a Boltzmann distribution-based algorithm is used to select the next operator. This makes including reinforcement learning (RL) in Soar straightforward. After an operator applies, all of the rules that created numeric preferences for that operator are updated based on any new reward and the expected future reward, which is the summed numeric value of the numeric preferences for the next selected operator [9]. RL in Soar applies across all goals, including impasse-generated subgoals. One intriguing aspect of RL in Soar is that the mapping from situation and operator to expected reward (the value function) is encoded as collections of rules. Only those rules that match the current situation participate in the selection of an operator, and there can be many rules contributing estimates of future reward for a single operator. This representation supports hierarchical and coarse-coding encodings of value functions, which can greatly speedup learning [10]. The value of reinforcement learning is that it allows Soar agents to improve their decision making over time as it receives feedback from the environment.

Semantic memory

The original Soar had no way of directly encoding semantic knowledge. Semantic learning and memory provides the ability to store and retrieve declarative facts about the world, such as tables

have legs, dogs are animals, and Ann Arbor is in Michigan, where the cues used for accessing the knowledge are determined at runtime. This capability has been central to ACT-R's ability to model a wide variety of human data and adding it to Soar should enhance our ability to create agents that reason and use general knowledge about the world. In Soar, semantic memory is built from structures that occur in working memory [11]. A structure from semantic memory is retrieved by creating a cue in a special buffer in working memory. The cue is then used to search for the best match in semantic memory, which is then retrieved into working memory. Soar uses base-level activation to determine the best match, which biases the result using recency and frequency of access. Adding semantic memory to Soar has allowed us to efficiently use existing large declarative knowledge bases, as well as develop agents that build up declarative knowledge over time.

Episodic memory

In contrast to semantic memory, which contains knowledge independent of when and where it was learned, episodic memory contains memories of what was experienced over time [12]. Although similar mechanisms have been studied in case-based reasoning, episodic memory is distinguished by the fact that it is task-independent and thus available for every problem, providing a memory of experience not available from other mechanisms. Episodic learning is so simple that it is often dismissed in AI as not worthy of study. Although simple, one has only to imagine what life is like for amnesiacs to appreciate its importance for general intelligence [13].

In Soar, episodic memory includes specific instances of the structures that occur in working memory at the same time, providing the ability to remember the context of past experiences as well as the temporal relationships between experiences [14]. An agent retrieves an episode by the deliberate creation of a cue, which is a partial specification of working memory in a special buffer. Once a cue is created, episodic memory finds the best partial match and recreates it in a separate working memory buffer (to avoid confusion between a memory and the current situation). The agent can also use cues that move forward or backward in time from a retrieved episode, providing the ability to replay an experience as a sequence of retrieved episodes. We have demonstrated that when episodic memory is embedded in Soar, it enables many

advanced cognitive capabilities such as virtual sensing, internal simulation and prediction, learning action models, and retrospective reasoning and learning.

Visual imagery

The previous extensions depend on Soar's existing symbolic short-term memory to represent the agent's understanding of the current situation. The generality and power of symbolic representations and processing are unmatched and the ability to compose symbolic structures is a hallmark of human-level intelligence. However, for some forms of processing, other representations can be much more efficient, as well as capture details difficult to represent symbolically. One compelling example is visual imagery [15], which is useful for visual-spatial reasoning. We have added modules to Soar that support visual imagery [16], including a short-term memory where images are constructed and manipulated; a long-term memory that contains images that can be retrieved into the short-term memory; processes that manipulate images in short-term memory; and processes that create symbolic structures from the visual images. Visual imagery is controlled by the symbolic system, which issues commands to construct, manipulate, and examine visual images.

With the addition of visual imagery, we have demonstrated that it is possible to solve spatial reasoning problems orders of magnitude faster. Spatial representations can also aid learning, and we have examples of combining mental imagery with reinforcement learning in video games [17]. Visual imagery also enables processing that is not possible with only symbolic reasoning, such as determining which letters in the alphabet are symmetric along the vertical axis (A, H, I, M, ...).

Emotion

Appraisal theories of emotion [18] propose that an agent continually evaluates a situation and that the result of the evaluation leads to emotion. The evaluation is hypothesized to take place along multiple dimensions, such as goal relevance (is this situation important to my goals?), goal conduciveness (is this situation good or bad for my goals?), causality (who caused the situation?), control (can I change the situation?), and so on. These dimensions are exactly what an intelligent agent needs to compute as it pursues its goals while interacting with an environment. Thus, we have created a computational implementation of a specific appraisal theory

Soar (cont.)

[19] in Soar [20], represented by the appraisal detector in Figure 2. In Soar, appraisals lead to emotions, emotions influence mood, and mood and emotion determine feelings. Individual appraisals produce either categorical or numeric values, which combine to form an intensity of the current feeling. This intensity becomes the intrinsic reward for reinforcement learning, which significantly speeds learning. A major goal of our future work is to explore how emotion, mood, and feeling can be used productively with other modules (such as retrieval from long-term memory and decision making), as well as in interacting with other agents.

Discussion

A significant concern when developing both the original Soar and Soar 9 is that we can develop and study agents that have large bodies of knowledge and exist for long periods of time. Thus, we have expended significant effort to develop efficient, scalable algorithms for all of the components. For example, our implementation of episodic memory supports the retaining of not just thousands, but also millions and even tens of millions of episodes, while still maintaining reactivity [21].

One of the most important benefits of creating a single architecture with all of these capabilities is that we can study the interaction between them. Some of these synergistic interactions include:

- using reinforcement learning to acquire knowledge that determines when and how retrievals from episodic memory are made, as well as learning to use what is retrieved to support future decision making [22];
- using mental imagery to support look-ahead search in spatial tasks (playing classic arcade video games), then using the result of the look-ahead as part of the state for reinforcement learning [17];
- using a wide variety of methods—including rules, task decomposition, episodic memory, semantic memory, and mental imagery—to support action modeling for look-ahead searches [23].

These types of capabilities are the payoff for developing integrated cognitive architectures. For those interested in Soar, it is available for free from <http://sitemaker.umich.edu/soar>.

References

[1] Langley, P., Laird, J. E., and Rogers, S. (2009). Cognitive architectures: Research issues and challenges. *Cognitive Systems Research*, 10(2), 141-160.
[2] Newell, A. (1990) Unified Theories

of Cognition. Harvard University Press.
[3] Laird, J. E., Newell, A., Rosenbloom, P. S. (1987). *Soar: An Architecture for General Intelligence*. *Artificial Intelligence*, 33(3), 1-64.
[4] Rosenbloom, P. S., Laird, J. E., & Newell, A. (1993) *The Soar papers: Research on Integrated Intelligence*. MIT Press, Cambridge, MA.
[5] Laird, J. E., & Rosenbloom, P.S. (1996) The evolution of the Soar cognitive architecture. In T. Mitchell (ed.) *Mind Matters*, 1-50.
[6] Laird, J. E. (2012). *The Soar Cognitive Architecture*. MIT Press, Cambridge, MA.
[7] Pearson, D. J., and Laird, J. E., *Incremental Learning of Procedural Planning Knowledge in Challenging Environments*, *Computational Intelligence*, 2005, 21(4), 414:439.
[8] Steier, D.S., Laird, J.E., Newell, A., Rosenbloom, P.S., Flynn, R., Golding, A., Polk, T.A., Shivers, O., Unruh, A., Yost, G.R. (1987) Varieties of Learning in Soar. *Proceedings of the Fourth International Machine Learning Workshop*.
[9] Nason, S., and Laird, J. E. (2005) Soar-RL: Integrating reinforcement learning with Soar. *Cognitive Systems Research*, 6(1), 51-59.
[10] Wang, Y., and Laird, J.E. (2007) The Importance of Action History in Decision Making and Reinforcement Learning. *Proceedings of the Eighth International Conference on Cognitive Modeling*. Ann Arbor, MI.
[11] Derbinsky, N., Laird, J. E., and Smith, B. (2010). Towards Efficiently Supporting Large Symbolic Declarative Memories, *Proceedings of the Tenth International Conference on Cognitive Modeling*, Philadelphia, PA.
[12] Tulving, E. (1983) *Elements of Episodic Memory*. Oxford: Clarendon Press.
[13] Nolan, C. (2000) *Memento*, New Market Films.
[14] Nuxoll, A. M., Laird, J. E. (2012) Enhancing Intelligent Agents with Episodic Memory, *Cognitive systems Research*, 17-18,34-48.
[15] Kosslyn, S. M., Thompson, W. L., & Ganis, G. (2006). *The Case for Mental Imagery*. New York, New York: Oxford University Press.
[16] Lathrop, S. D., Wintermute, S., Laird, J. E. (2011), Exploring the Functional Advantages of Spatial and Visual Cognition from an Architectural Perspective. *Topics in Cognitive Science*, 3, 796-818.
[17] Wintermute, S. (2010). Using Imagery to Simplify Perceptual Abstraction in Reinforcement Learning Agents. *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*.
[18] Roseman, I. and Smith, C. A. (2001) Appraisal theory: Overview, Assumptions, Varieties, Controversies. In K. R. Scherer, A. Schorr, & T. John-

stone (Eds.) *Appraisal processes in Emotion: Theory, Methods, Research*, 3-19. Oxford University Press.

[19] Scherer, K. R. (2001) Appraisal considered as a process of multi-level sequential checking. In K. R. Scherer, A. Schorr, & T. Johnstone (Eds.) *Appraisal processes in Emotion: Theory, Methods, Research*. 92-120. Oxford University Press.

[20] Marinier, R. P., Laird, J. E., and Lewis, R. L., (2009) A Computational Unification of Cognitive Behavior and Emotion, *Cognitive Systems Research*, 10(1), 48-69.

[21] Derbinsky, N., Li, J., and Laird, J. E. (2012) A Multi-Domain Evaluation of Scaling in a General Episodic Memory. *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, 2012.

[22] Gorski, N. A., Laird, J. E. (2011) Learning to use episodic memory *Cognitive Systems Research*, 12, 144-153

[23] Laird, J. E., Xu, J. Z., and Wintermute, S. (2010). Using Diverse Cognitive Mechanisms for Action Modeling. *Proceedings of the Tenth International Conference on Cognitive Modeling*, Philadelphia, PA.



Dr John E. Laird
John L. Tishman Professor of Engineering
Computer Science Division
University of Michigan, USA

CernoCAMAL: A Bayesian-BDI Cognitive Architecture

Introduction

This article presents a Bayesian BDI (Belief-Desire-Intention) cognitive architecture, dubbed CernoCAMAL, that can be used to govern artificial minds probabilistically. The primary aim of the CernoCAMAL research project was to investigate how a predecessor architecture known as CAMAL could be extended to reason probabilistically about domain model objects through perception, and how Bayesian formalism could be integrated into its BDI model to coalesce a number of mechanisms and processes. Extensive experiments in synthetic simulation and robotic test beds demonstrated improvements and increased efficacy in the overall cognitive performance, success rate, task effectiveness, and goal achievement of the CernoCAMAL architecture.

Motivation

CAMAL (Computational Architectures for Motivation, Affect, and Learning) [1] is an example of a general class of integrative cognitive architectures, drawing together a number of threads in Cognitive Science and Artificial Intelligence, such as perception, action, decision making, motivation, affect, and learning. CAMAL is essentially a UTC (Unified Theory of Cognition) [2] that tries to answer some of the questions that comprise Norman's Cognitive Science agenda [3]. The motivation and impetus for developing CernoCAMAL was the considerable evidence that probabilistic thinking and reasoning is linked to cognitive development and plays a role in cognitive functions, such as decision making and learning [e.g., 4, 5, 6, 7, 8, 9, 10, 11, 12, 13]. This led us to believe that probabilistic reasoning is an essential aspect of the process of cognition and, therefore, must be considered in any adequate description of it, such as a cognitive architecture.

Framework & Components

CAMAL used a variant of the a-CRIBB reasoning model [14] i.e. a BDI model and an affect model, plus a motivational blackboard. At the deliberative level, affective values and affordances can be associated with processes and predicates, and then relayed as control signals to instantiate and modify aspects of motivators and their associated representations and behaviours [15].

One of the limitations of the BDI model, however, was the lack of any explicit mechanism to express degrees of belief. The belief statements represented beliefs as categorical states. Therefore, they could not be adequately valenced via affective values and affordances, in line with the affect and motivational models. Given that our current research presents an affect- and affordance-based core for the mind [6], it seemed

reasonable to conjecture that beliefs, too, should be grounded in the use of affect with the aim to be consistent across different domains, tasks, and levels of processing. We, therefore, extended the CAMAL architecture and equipped it with belief affordances, so that CernoCAMAL exploits the same reasoning model, plus an extended version of the CAMAL belief structure that incorporates probabilities as degrees of belief associated with different information sources, in a Bayesian BDI. This formalized the BDI model, facilitating the use of a consistent metric across all aspects of affect, reasoning, and domain model management.

Extended Belief Structure

In CernoCAMAL belief statements are represented as graded states, using probability and Bayesian formalisms. An Extended Belief Structure (EBS) is used to represent the degree of beliefs numerically, and then manipulate them, using clauses of the form: *belief(Descriptor, Source, Time, DegBel)*.

The EBS associates a probability value *DegBel* with every belief statement in CernoCAMAL which defines the degree to which the belief statement is believed to be true. This will enable the computation of changing degrees of belief using a Bayesian BDI, affect, and motivational models to determine the agent's intentions, actions, or behaviours. The EBS and BDI model are now compatible with the way in which the affect and motivational models operate throughout CernoCAMAL: having an associated affective magnitude that can fluctuate according to success or failure associated with that element. Put differently, affect now serves as a decision metric and affective values as a currency across the entire architecture, including beliefs and the BDI model.

CernoCAMAL Probabilistic Reasoner

The proposal to use the EBS led to the development of the CernoCAMAL Probabilistic Reasoner (CPR) that deliberates probabilistically over the perceptual feedback generated by reactive subsystems. This development has resulted in

the CPR consistently reasoning about the domain model objects and their instances and consequently keeping CernoCAMAL's model of its surroundings up to date. The belief descriptor 'apriori_prob' that is part of the domain model defines the a-priori probability of an object being present in the environment: *apriori_prob(Object, DegBel)*.

CernoCAMAL's CPR, given the list of domain model objects and assumed degree_of_belief for various information sources, can compute the changing degrees of belief, infer posterior probabilities correctly, assign them to the appropriate belief descriptors, and reason probabilistically about the number of objects and their instances that may be present in the environment.

Predator-Prey Tile World

The CernoCAMAL framework was initially implemented as a situated cognitive agent in a synthetic predator-prey test bed. This test bed uses a graphical tile world with operators that affect the world, incorporating a (white) Cerno agent, several edible spheres (blue objects), preys (red agents), and predators (green agents). A screenshot of this synthetic terrain along with its corresponding Prolog command window are shown in Figure 1.

Probability computation lends itself well to predator-prey scenario, since the computed probabilities could imply the extent or risk that a particular entity is a predator, etc. The power of CPR lies in its memory facility. Suppose the CernoCAMAL cognitive agent starts moving in its dynamic, uncertain test bed and begins to stumble across various objects. The agent's findings are represented as reactive feedback lists. Having found a domain model object or perhaps an unidentified one, the CPR reasons about the probabilities of various objects being present in the environment. Now, suppose that Cerno loses or eats a sphere or prey. That lost or eaten instance is withdrawn from the simulation world. The epistemological difficulty here would be: How does Cerno limit the scope of the propositions it must re-consider and re-evaluate in the

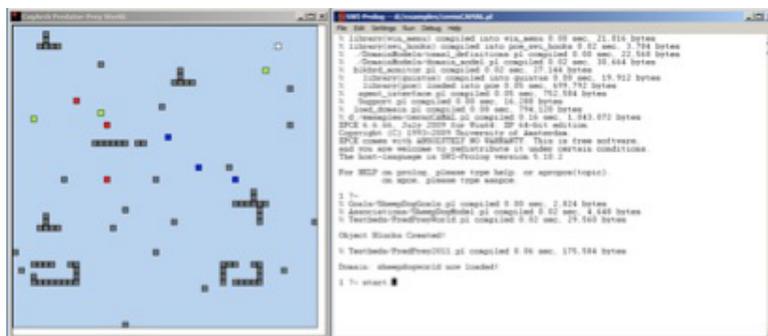


Figure 1: Screenshot of Pred-Prey Synthetic Terrain with its corresponding Prolog Command Window

CernoCAMAL (cont.)

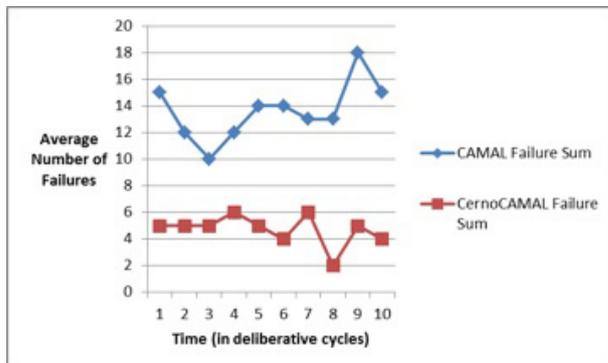


Figure 2: Overall Reduction in the Number of Failures between CAMAL and CernoCAMAL architectures

light of its actions? This is solved by the loss' or destruction's memory trail that is formed by the Cerno's CPR; so that when an instance is found:

- If the instance list of that object is empty (i.e. that object is definitely the first to be found), then a new unique instance is generated.
- If there is a memory trail of a previously lost instance, then that previously generated lost instance is re-created, instead of a new incremental one.
- If it can refer to a previously found instance, rather than generating a new unique instance (rather than extending the instance list), then it refers to that previously found instance.
- If it has been found for the 'second' time, then the first one refers to a previously found instance and the second one generates a new unique instance.
- If an instance was eaten, then the trail for that instance will be marked as destroyed, so that upon finding an instance, the identification is made based on the fact that it cannot possibly be that destroyed instance.

Experimental Results

A succession of experiments were carried out to evaluate CernoCAMAL's overall performance, success rate, task effectiveness, and goal achievement. A general sample set of results are presented in Figure 2 that clearly shows a decrease in the number of failures in time that occurred compared to CAMAL.

Summary and Conclusion

Cognition is better viewed as solving probabilistic, rather than logical, inference problems. The Probabilistic approach to cognition has, therefore, become an established approach [16, 17].

CernoCAMAL uses a Bayesian-BDI schema to drive a motivational blackboard. The inclusion of degree-of-belief in its belief predicates enables the architecture to select a focused belief set that reflects its current activities, as high-

lighted by actions, objects, and agents referenced in a current motivator. The motivator enables goal revision and the selection of the next goal based on goal importance and current beliefs and goal success. The deliberative processing of these constructs allows the selection of an appropriate action related to specific objects and tasks. This, in turn, drives motivator revision using the association construct, which in turn enables belief-desire-intention combinations to be ranked based on the likelihood of their success (association values). The goal importance, association insistence, motivator intensity, and degree-of-belief are all underpinned by affordances; i.e. they are all consistently grounded in affect.

References

- [1] D.N. Davis, Control States and Complete Agent Architectures, *Computational Intelligence*, 17: 621-650, 2001.
- [2] A. Newell, *Unified Theories of Cognition*, Harvard University Press, 1990.
- [3] D.A. Norman, Twelve Issues for Cognitive Science, *Cognitive Science*, 4: 1-33, 1980.
- [4] J. Piaget, *Judgment and Reasoning in the Child*, London: Routledge & Kegan Paul, 1928.
- [5] J. Piaget and B. Inhelder, *The Origin of the Idea of Chance in Children*, London: Routledge & Kegan Paul, 1951.
- [6] E. Fischbein, *The Intuitive Sources of Probabilistic Thinking in Children*, Dordrecht: Reidel, 1975.
- [7] J.M. Shaughnessy, *Misconceptions of Probability: From Systematic Errors to Systematic Errors and Decisions*, 1981 Year Book: Teaching Statistics and Probability, 90-100.
- [8] D. Green, A Survey of Probability Concepts in 3000 Pupils aged 11-16 years, In *Proceedings of the 1st International Conference on Teaching Statistics*, 2: 766-783, 1983.
- [9] R.F. Peard, Student Decision-Making in a Game of Chance and Misconceptions in Probabilistic Reasoning, In *Proceedings of the 18th Annual Conference of the Mathematics Education Research Group of Australia*, 469-475, 1995.

[10] J. Truran, Children's Misconceptions about the Independence of Random Generators, In *Proceedings of the 20th International Conference for the Psychology of Mathematics Education*, 4: 331-338, 1996.

[11] K. Truran, Children's Use of a Representative Heuristic, In *Proceedings of the 20th International Conference for the Psychology of Mathematics Education*, 4: 339-346, 1996.

[12] E. Fischbein and D. Schnarch, The Evolution with Age of Probabilistic, Intuitively-Based Misconceptions, *Journal of Research in Mathematics Education*, 28(1): 96-105, 1997.

[13] J. Way, The Development of Young Children's Notions of Probability, In *Proceedings of the 3rd Conference of the European Society for Research in Mathematics Education*, Thematic Group, 5, 2003.

[14] D.N. Davis and S.C. Lewis, Affect and Affordance: Architectures without Emotion, In *Proceedings of AAAI 04 Spring Symposium*, Stanford, USA, 2004.

[15] D.N. Davis, Cognitive Architectures for Affect and Motivation, *Cognitive Computation*, 2:199-216, 2010.

[16] M. Oaksford and N. Chater, *Bayesian Rationality: The Probabilistic Approach to Human Reasoning*, Oxford University Press, 2007.

[17] M. Oaksford and N. Chater, *Précis of Bayesian Rationality: The Probabilistic Approach to Human Reasoning*, *Behavioural and Brain Sciences*, 32:69-120, 2009.



Hossein Miri
PhD Candidate
Department of Computer Science,
University of Hull

Cognitive science is not computer science

In AISBQ #133 two articles struck a nostalgic tone for the days of good-old fashioned AI. While reminiscing about the late John McCarthy, Sloman defended the principles of symbolic AI as a means of understanding human intelligence. This sentiment was echoed by Langley, who lamented that the founding ideals of symbolic AI had fallen into disfavor. He concluded that these ideals remained valid and advocated for their re-adoption by researchers and educators.

I want to strike a different tone, namely that the way out of the crisis of symbolic AI should lead us forward, not backwards. I agree with Langley that research in AI can provide useful heuristics to guide our search for a viable theory of mind, but it is precisely for that reason that I think that the computationalist theory of mind should be abandoned.

AI was one of the foundational disciplines of cognitive science, but now the strongest voice in cognitive science is that of neuroscience which does not speak in symbolic terms as researchers are more interested in identifying neural correlates of cognition. In addition, with the growing interest in consciousness science, new methods are being developed to measure and describe what lived experience is like. I do not agree with the neuro-chauvinist attitude of cognitive science, but I think that there are good reasons why research in AI is nowadays mainly restricted to computer science departments: the computational theory of mind, which tried to promote symbolic AI as a scientific theory of natural mind, is untenable.

In methodological terms there are serious problems with the computer metaphor. On the one hand, computational logic is independent from any particular means of realization. While this initially appeared to be an advantage, since it allowed computer scientists to study the mind without considering biological details, it now blocks genuine dialogue with neuroscientists. On the other hand, computational logic is also independent from personal experience. Again, while this appeared to be an advantage because it allowed researchers to avoid the difficult topics of introspection and consciousness, it now blocks genuine dialogue with the new field of consciousness science.

Even outside this neuroscience-psychology-philosophy triad, the computational hypothesis of human nature has fallen out of favor. Anthropology, for example, has little use for a theory which assumes that symbols are the building blocks of mechanisms located inside an individual mind (or even brain),

given that a symbolic representation as such only makes sense in its context of shared socio-cultural practices [1].

While symbolic AI has been indeed successful in many areas and continues to be a driving force in the digital technological revolution [2], as a scientific theory of the human mind, it misses the point. What are the alternatives? I believe that cognitive science should replace the computationalist metaphor with an existential stance that centers on the living (biological) and lived (experiential) body [3].

I suggest that a more encompassing science of human nature must be able to intertwine an understanding of how the fundamental values that are governed by our metabolic existence are shaped by the enabling and constraining concerns of our socio-cultural existence. To do so we need to be able to integrate dynamic processes at a range of time scales, and at a range of levels (individual, dyadic, social, cultural, etc.), and we must be able to connect those dynamics with changes in the first-person experience of the people who are embedded in them. All of these areas of research together form what has been called the paradigm of "enaction". I believe that a suitably revised AI can play a role in all of these aspects.

For example, consider the role of biological foundations. There are good reasons for accepting that there is no mind without life, and that life requires some kind of material self-constitution, i.e. metabolism. It is in the messy processes of continuous self-creation that living and mental autonomy is rooted (some issues related to this perspective have been discussed by Boden [4]).

While robots do not have to be physically self-producing in order to support some form of self-constitution, 'enactive' AI could re-conceive a robot as a kind of interface, with the capacity to allow autonomous dynamics to emerge in the domain of interaction.

Interestingly, these autonomous dynamics can also be found in the case of social interaction. Moreover, when social interactions take on a 'life of their own' they can enable and constrain the actions of the individuals. I have shown how simple agent-based models can help to understand autonomous social dynamics, and why these dynamics cannot be reduced to internal mechanisms alone. Accordingly, research in AI can help social psychologists by providing alternative hypotheses about social interaction that do not depend on the computer metaphor. It is even possible to relate these findings to consciousness studies, because the dynamic structures of social interaction can inform our un-

derstanding of the dynamic structures of intersubjectivity. In this way, agent-based research into behavioral dynamics can be used as an extension of phenomenological philosophy.

Another application of this approach can be found in the field of human-robot interaction. Importantly, this area of research explicitly brings a human participant into play. One of the key insights of the enactive approach is that the practical use of any technology modulates our lived experience. A particularly striking example of this is the way in which the use of sensory substitution interfaces can constitute new kinds of perceptual experience. I see this as a promising possibility for the AISB community to connect with what is going on in consciousness science. The field of human-computer interfaces enables a systematic study of how variations in the sensorimotor dynamics enabled by the technological interface relate to variations in the lived experience of the human user.

A shift from computationalism to an enactive approach provides many novel possibilities. Freed from the limitations of the computer metaphor, the methods of AI are useful scientific tools to help us in our investigation of the complex systemic organization of life, mind and sociality which can help us to better understand the regularities of human experience as we interact with our technological world. These opportunities to engage in the next phase of cognitive science are surely good news for all of us who were initially drawn to AI because of an interest in the human mind.

References

- [1] Harvey, I. (1996). Untimed and misrepresented: connectionism and the computer metaphor. *AISB Quarterly*, 96, 20-27.
- [2] Whitby, B. (2002). Two Myths about AI. *AISB Quarterly*, 110, 1-3.
- [3] Froese, T. (2011). Breathing new life into cognitive science. *Avant*, 2(1), 113-129.
- [4] Boden, M. A. (1994). Autonomy and Artificiality. *AISB Quarterly*, 87, 22-28.



Dr Tom Froese (froese.wordpress.com)
Postdoctoral Research Fellow,
Department of General Systems Studies,
University of Tokyo, Japan

Information exchange in population-based algorithms

Introduction

In a recent conference – International Conference on Evolutionary Computation Theory and Application (ECTA 2011) – among other discussions and presentations on evolutionary computation, a special session was held to discuss the Future of Evolutionary Computation. One of the papers presented in the special session proposed the integration of Swarm Intelligence (SI) algorithms and Evolutionary algorithms (EAs) as one possible future approach in the Evolutionary Computation (EC) [1].

This work narrates the early research on using Stochastic Diffusion Search (SDS) – a swarm intelligence algorithm – to empower the Differential Evolution (DE) algorithm – an evolutionary algorithm – over a set of optimisation problems. The results reported suggest that the resource allocation mechanism deployed in SDS has the potential to improve the optimisation capability of the classical evolutionary algorithm used in this experiment.

In the literature, nature inspired swarm intelligence algorithms and biologically inspired evolutionary algorithms are typically evaluated using benchmarks that are often small in terms of their objective function computational costs [2]; this is often not the case in real-world applications. This paper is an attempt to pave the way for more effectively optimising computationally expensive objective functions, by deploying the SDS diffusion mechanism to more efficiently allocate DE resources via information-sharing between the members of the population.

We introduce SDS [3], a multi-agent global search and optimisation algorithm, which is based on simple interaction of agents. SDS is inspired by one species of ants, *Leptothorax acervorum*, using a ‘tandem calling’ mechanism (one-to-one communication), where the forager ant which finds the food location, recruits a single ant upon its return to the nest, and therefore the location of the food is physically publicised. A high-level description of SDS is presented in the form of a social metaphor (the Mining Game) demonstrating the procedures through which SDS allocates resources.

The Mining Game

This metaphor provides a simple high-level description of the behaviour of agents in SDS, where mountain range

is divided into hills and each hill is divided into regions:

A group of miners learn that there is gold to be found on the hills of a mountain range but have no information regarding its distribution. To maximize their collective wealth, the maximum number of miners should dig at the hill which has the richest seams of gold (this information is not available a-priori). In order to solve this problem, the miners decide to employ a simple Stochastic Diffusion Search.

- At the start of the mining process each miner is randomly allocated a hill to mine (his hill hypothesis, h).
- Every day each miner is allocated a randomly selected region, on the hill to mine.

At the end of each day, the probability that a miner is happy is proportional to the amount of gold he has found. Every evening, the miners congregate and each miner who is not happy selects another miner at random for communication. If the chosen miner is happy, he shares the location of his hill and thus both now maintain it as their hypothesis, h ; if not, the unhappy miner selects a new hill hypothesis to mine at random.

As this process is isomorphic to SDS, miners will naturally self-organise to congregate over hill(s) of the mountain with high concentration of gold.

In the context of SDS, agents take the role of miners; active agents being ‘happy miners’, inactive agents being ‘unhappy miners’ and the agent’s hypothesis being the miner’s ‘hill-hypothesis’.

The SDS algorithm commences a search or optimisation by initialising its population (e.g. miners, in the mining game metaphor). In any SDS search, each agent maintains a hypothesis, h , defining a possible problem solution. In the mining game analogy, agent hypothesis identifies a hill. After initialisation two phases are followed: Test and Diffusion Phases.

In the test phase, standard SDS checks whether the agent hypothesis is successful or not by performing a partial hypothesis evaluation which returns a boolean value. Later in the iteration, contingent on the precise recruitment strategy employed, successful hypotheses diffuse across the population, allowing information on potentially good solutions to spread throughout the entire population of agents.

In the Diffusion phase, each agent

recruits another agent for interaction and potential communication of hypothesis. In the mining game metaphor, diffusion is performed by communicating a hill hypothesis.

In this work, the coupling strategy is presented, followed by a brief discussion.

Coupling SDS and DE

The goal of this process is to verify whether the information diffusion and dispensation mechanisms deployed in SDS may on their own improve the behaviour of Differential Evolution algorithm (DE).

DE, one of the most successful Evolutionary Algorithms (EAs), is a simple global numerical optimiser over continuous search spaces which was first introduced by Storn and Price [4].

DE is a population based stochastic algorithm, proposed to search for an optimum value in the feasible solution space. This work uses one variations of DE (DE/best/1). More information available in [1].

In the new architecture introduced in the paper, a standard set of benchmarks (from [5, 6] test suite) are used to evaluate the performance of the coupled algorithm. The resource allocation (or recruitment) and partial function evaluation sides of SDS are used to assist allocating resources after partially evaluating the search space.

Each DE agent has three vectors (target, mutant and trial vectors); and each SDS agent has one hypothesis and one status. In the experiment reported here (coupled algorithm), as stated before, SDS test diffusion cycle is run for n Function Evaluations (FEs) and then DE commences with the optimisation, taking its target vectors from SDS agents’ positions. The details on the parameters values, implementation issues and the pseudo-code are reported in [1].

In the test phase of a standard stochastic diffusion search, each agent has to partially evaluate its hypothesis. The guiding heuristic is that hypotheses that are promising are maintained and those that appear unpromising are discarded.

In this work, the test phase is conducted by comparing the fitness of each agent against that of a random one; if the selecting agent has a better fitness value, it will become active, otherwise it will be flagged inactive.

In the diffusion phase, each inactive agent picks another agent randomly, if the selected agent is active, the selected agent communicates its

Information exchange (cont.)

hypothesis to the inactive one; if the selected agent is also inactive, the selecting agent generates a new hypothesis at random from the search space.

As outlined in our ECTA paper [1] after the initial n FEs during which SDS test diffusion cycle iterates, DE algorithm should run.

In this work, a number of experiments are carried out and the performance of one variation of DE algorithm (DE/best/1) is contrasted against the coupled SDS-DE algorithm (sDE). As mentioned earlier, the algorithms are tested over a number of standard benchmarking functions, preserving different dimensionality and modality (see [7] for more information on the benchmarks used).

The experiments are conducted with the population of 100 agents and the halting criterion for this experiment is when the algorithms perform 300,000 FEs.

There are 30 independent runs for each benchmark function and the results are averaged over these independent trials.

DE is run after 100,000 FEs until the termination criterion which is 300,000 FEs. These values were selected merely to provide a brief initial exploration of the behaviour of the new coupled algorithm; no claim is made for their optimality.

In the results, over all benchmarks, other than one (out of 14), DE algorithm does not significantly outperform the coupled algorithm. On the other hand, in most cases (9 out of 14 benchmarks), the coupled algorithm significantly outperforms the classical DE algorithm.

Discussion

The resource allocation process underlying SDS offers three closely coupled mechanisms to the algorithm's search component to speed its convergence to global optima:

- 'efficient, non-greedy information sharing' instantiated via positive feedback of potentially good hypotheses between agents;
- dispensation mechanism – SDS-led random-restarts – deployed as part of the diffusion phase;
- random 'partial hypothesis evaluation', whereby a complex, computationally expensive objective function is broken down into 'k independent partial-functions', each one of which, when evaluated, offers partial information on the absolute quality of current algorithm search parameters. It is this mechanism of iterated selection of

a random partial function that ensures SDS does not prematurely converge on local minimum.

To further analyse the role of SDS in the coupled algorithm, the diffusion phase of SDS algorithm is modified to investigate the dispensation effect caused by randomising a selection of agent hypotheses (effectively instantiating the population with SDS-led random-restarts). In other words, after the SDS test phase, the hypothesis of each inactive agent is randomised.

As the results reported in [1] suggest, although information sharing plays an important role in the performance of the coupled algorithm, the significance of dispensation mechanism (randomly restarting the inactive agents) in improving the performance of the algorithm cannot be discarded.

In some cases (3 out of 14), solely the dispensation mechanism (sDispDE), which is facilitated by the test phase of the SDS algorithm, demonstrates a significantly better performance compared to the coupled algorithm. However, in several cases, the coupled algorithms outperform the modified algorithm (7 cases), out of which 4 cases are performing significantly better.

It is shown that out of 14 benchmarks, sDE exhibits the best performance as it is among the most significant in 9 cases; sDispDE and DE are among the best in 7 and 2 cases, respectively.

The results show the importance of coupling the SDS-led restart mechanism (dispensation mechanism) and the information sharing which are both deployed in SDS algorithm.

The third SDS component feature, which is currently only implicitly exploited by the coupled algorithm, is 'randomised partial hypothesis evaluation' (see [7] for a detailed explanation on the implicit deployment of this feature). This work aims at emphasising the role of information exchange as an influential factor in assisting EA optimisers; it also reinforces the idea of the integration of SI algorithms with EAs as a potential future approach in Evolutionary Computation.

References

- [1] al-Rifaie, M. M., Bishop, M., and Blackwell, T. (2011a). An investigation into the use of swarm intelligence for an evolutionary algorithm optimisation. International Conference on Evolutionary Computation Theory and Application (ECTA 2011).
- [2] Whitley, D., Rana, S., Dzubera, J., and Mathias, K. E. (1996). Evaluating

evolutionary algorithms. *Artificial Intelligence*, 85(1-2):245–276.

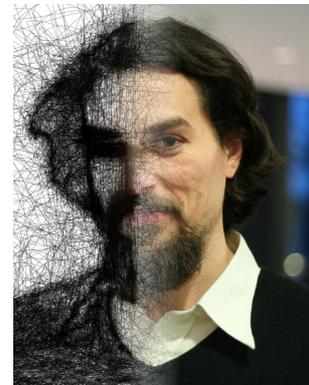
[3] Bishop, J. (1989). Stochastic searching networks. pages 329–331, London, UK. Proc. 1st IEE Conf. on Artificial Neural Networks.

[4] Storn, R. and Price, K. (1995). Differential evolution – a simple and efficient adaptive scheme for global optimization over continuous spaces. TR-95-012, [online]. Available: <http://www.icsi.berkeley.edu/storn/litera.html>.

[5] Jones, D. R., Perttunen, C. D., and Stuckman, B. E. (1993). Lipschitzian optimization without the lipschitz constant. *J. Optim. Theory Appl.*, 79(1):157–181.

[6] Jong, K. A. D. (1975). An analysis of the behavior of a class of genetic adaptive systems. PhD thesis, University of Michigan, Ann Arbor, MI, USA.

[7] al-Rifaie, M. M., Bishop, M., and Blackwell, T. (2011b). Resource allocation and dispensation impact of stochastic diffusion search on differential evolution algorithm; in: *Nature Inspired Cooperative Strategies for Optimisation (NICSO 2011)*, Springer.



Dr Mohammad Majid al-Rifaie
Visiting Researcher
Department of Computing,
Goldsmiths College,
University of London

Book review: Luciano Floridi, *The Philosophy of Information*

Professor Luciano Floridi is on the course of developing an overarching collection of books, referred to by himself as *Principia Philosophiae Informationis*. The book *The Philosophy of Information* is the first volume in this collection.

As Floridi himself warns the reader, "... this is not an easy book to read, to put it mildly." It is very dense and the author's writing style, although elegant and witty, can be a little overwrought at some points.

The book is neatly structured, and its structure becomes clearer as the reader progresses along the chapters. The first three chapters are meta-theoretical and lay some methodological foundations upon which the following chapters are built.

In chapter 1 (What is the philosophy of information?) the author sets the framework for the whole book. Depending on the reader's background it can be quite hard to read, as its arguments are grounded on a large collection of bibliographical references.

In chapter 2 (Open problems in the philosophy of information) the author takes the famous Hilbert twenty three problems as a model, and introduces eighteen problems, whose solutions should enlighten the field of Philosophy of Information. These problems, however, are not quite used as open problems in the sense suggested by Hilbert over a hundred years ago, instead they are guidelines which provide some directions for the discussions and results presented in the following chapters.

In chapter 3 (The method of levels of abstraction) the author analyses carefully how concepts and domains can be organised in different levels of abstraction, and how this methodology can help in the organisation of knowledge.

In chapters 4 to 8 the author builds his theory of semantic information, step by step. In chapter 4 (Semantic information and the veridicality thesis) he builds the notion of semantic information as well-formed, meaningful and truthful data. In chapter 5 (Outline of a theory of strongly semantic information) he further develops this notion and introduces a quantitative theory of semantic information.

In chapter 6 (The symbol grounding problem) the author discusses the age-old problem of attaching meaning to symbols. In chapter 7 (Action-based semantics) he introduces his own proposed solution for this problem, which is related to—and influenced by—the research tradition on cognitive robotics.

In chapter 8 (Semantic information and the correctness theory of truth) the author tackles the problem of ascribing truth to data, in order to qualify pieces

of data as semantic information.

In chapter 9 (The logical unsolvability of the Gettier problem) Floridi discusses the arguments presented by Gettier in his famous 1963 article [1] to concur with that author that the characterisation of knowledge as justified true belief is less precise than required or expected.

In chapters 10 to 12 the author makes use of the foundations built in the previous nine chapters to introduce a logic of information—an important notion which is carefully presented in this book, revising and detailing previous presentations as technical papers—as well as a theory of knowledge based on information, in which, given certain pre-requisites, information can be upgraded to knowledge.

In chapter 10 (The logic of being informed) the author introduces an information logic, as a multimodal logic whose axioms are carefully selected in order to obtain the desired properties for a modal operator "a is informed that p". In chapters 11 (Understanding epistemic relevance) and 12 (Semantic information and the network theory of account) he introduces conceptual tools upon which the pre-requisites to upgrade information to knowledge can be built.

In chapter 13 (Consciousness, agents, and the knowledge game) the author discusses the notion of consciousness, based on the results from the previous twelve chapters.

As well as tickling the philosophically minded reader, chapters 10 to 13 can be of particular interest to researchers on Artificial Intelligence, as they focus on synthetic agents and provide sketches of computational theories which can be implemented and used for empirical research. The final two chapters in this remarkable book have a different tone.

In chapter 14 (Against digital ontology) the author discusses the difference between a digital ontology, according to which the ultimate nature of reality is digital, and informational ontology, according to which the ultimate nature of reality is structural, and explains why he prefers the latter.

In chapter 15 (A defence of informational structural realism) Floridi closes the book, presenting in detail his proposed informational ontology and, this way, proposing the Philosophy of Information as an alternative foundation for a Philosophy of Science.

The book ends—as expected—with the bibliographical references referred to along the text. The size and breadth of the collection of references are worth highlighting, as they display the solid grounding upon which Floridi has built

his work, and are a most valuable gift granted by the author to the interested reader.

Given the breadth and depth of coverage of all its topics, the careful organisation and structuring of concepts, and the relevance of its contents, *The Philosophy of Information* shall be deemed essential reading for philosophers and computer scientists alike, especially those interested in Artificial Intelligence.

References

[1] Gettier, E. Is Justified True Belief Knowledge? *Analysis*, 23. 121-123.

Luciano Floridi: *The Philosophy of Information*. Oxford University Press. 2011. 405 pages. ISBN 978-0-19-923238-3.



Dr Flavio Soares Correa da Silva
Associate Professor,
Department of Computer Science
University of Sao Paulo, Brazil

Book review: Joscha Bach, Principles of Synthetic Intelligence

Published in 2009, "Principles of Synthetic Intelligence" (PSI) has passed under the radar of most readers. With zero reviews on Amazon (as of this writing) after three years, it deserves better. Its major impediment to attracting readers in the English-speaking world is that it is based on a relatively unknown body of work from Dietrich Dörner at Bamberg which remains untranslated from the original German. Hopefully, readers will give it another chance, once this review comes out.

The book itself consists of roughly three sections of content: a presentation of the PSI theory of Dörner for the first time in English, then a summary of work done by others using the PSI theory, ending in a detailed implementation of the theory in the author's MicroPSI architecture.

PSI itself is a cognitive architecture and the brainchild of Dörner. However, prior to this book, it was mainly defined informally by reference to specific implementations by Dörner and others. Bach goes to great lengths to summarize, condense, and extract the theory from the various parts and pieces that have been produced by Dörner and others. MicroPSI is the implementation of that theory by Bach (as part of his Doctoral work), which makes the theory workable by implementing several restricted parts of the PSI theory in a way that makes it attractive for use by researchers.

The PSI cognitive architecture is Neuro-symbolic and is derived from homeostatic agents that seek to fulfill motivations. The system grounds symbols in particular neural representations are theoretically uniform, where a single kind of structure can be used to represent anything in the system. The symbol grounding is particularly interesting and detailed, and it is worth reading the book just to see how it works.

The theory claims that the mind is a "perceptual symbol system", where all symbols are somehow perceptually grounded, as opposed to an amodal symbol system like ACT-R. For example, the concept "dog" in PSI would reference perceptual information (visual images perhaps) that the agent had encountered previously, but in ACT-R, it would be represented in some form of abstract statement in a predicate calculus not necessarily tied to any real per-

ceptions.

Another major part of the PSI architecture is the centrality of motivations. Rather than being tacked onto a cognitive system, motivations are the part of the system that control what and why a mind would act. Major motivations include analogs of hunger, thirst, pain, certainty, competence, and social affiliation. These motivations are represented in the architecture as demands that increase over time, which are perceived by an agent as urges, and then used as motives to create goals.

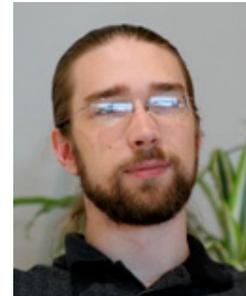
The simple physiological demands work as expected, but the cognitive demands of certainty and competence, as well as the social demand of affiliation are quite interesting. For instance, the urge to increase certainty can motivate an agent to explore its environment if little is known about it.

Emotion is also a central part of the PSI theory, and it is non-trivially incorporated into the system. Rather than being an external module, or a set of parameters, emotion is a dynamic re-configuration of the system in light of experience. Certain key cognitive constraints can be changed based on feedback from the agent's behaviours in the environment, similar to a startle response to an unexpected outcome of an agent's actions. Because it is dynamic, it will change over time as well as in response to external events. Interestingly, researchers have made models of personality features using this, and have been able to verify some predicted emotional behaviors empirically using human subjects.

In addition to the particularly interesting parts already mentioned, Bach provides an extensive review of the domains that PSI has been applied to, including social modeling, language studies, vision, and others. PSI provides either answers or partial answers for many big questions about cognition, though it is very much a living theory.

Overall, this book is filled with hidden gems, and should be on the reading list for anyone interested in symbol grounding, motivations, emotions, social modelling, or similar topics.

Principles of Synthetic Intelligence PSI: An Architecture of Motivated Cognition. Oxford University Press. 2009. 399 pages. ISBN: 0195370678



Ryan Kaulakis
Graduate Student
College of Information Sciences and
Technology,
Penn State University, USA

Event report: Roger Needham Lecture

On Tuesday 1st November, 2011, the BCS staged the 2011 Roger Needham Lecture at the Royal Society in London. In this annual lecture the winner of the BCS Roger Needham Award is officially awarded the prize and presents their work to the public. This year's winner is Prof Maja Pantic from Imperial College London. The lecture on 'Machine Understanding of Human Behaviour' presented a decade's research in the area.

Following an introduction by Prof Jim Norton (president of the BCS), the BCS/CPHC Distinguished Dissertations awardees were announced by Prof Ann Blandford (UCL). The prize was awarded to Daniel Greenfield (Cambridge University); runner up is Vera Demberg-Winterfall (Edinburgh).

The presentation of the Roger Needham Award 2011 to Prof Maja Pantic was carried out by Dr Andrew Blake (Head of Microsoft Research), who identified her as a driving force in the area of human behaviour recognition. In her excellent lecture on Facial Behaviour Understanding, Maja engrossed the audience with her enthusiasm for the research she has committed herself to.

She introduced the audience consisting of academics and representatives from industry to the history of her research in the area of facial expression recognition. This started in an MSc project on static analysis of human facial expressions at Delft. Here, she explored prototypic facial expressions using rule-based systems and was able to distinguish six basic expressions.

A total of 45 facial action units directly linked to the contraction of muscles were subsequently established, but not all of these were recognisable with the methods at hand in 2001. The problems related to motion and dynamic expression were not recognisable at the time. This led to the development of a technique called facial point tracking, which has been the focus of Maja's research over the past 10 years.

One of the problems calling for a solution was sudden and drastic head movement. Also, drastic changes in illumination would previously prevent facial expressions from being recognised. Temporal models have been developed to help identify errors due to artifacts by enabling to track actual possibilities given the requirements of the physical muscular movement. Temporal evolu-

tion in face videos offers possibilities to detect spontaneous laughter as opposed to acted laughter (real joy versus acted happiness). Affective dimensions of dynamic continuous behaviour then led to multi-dimensional continuous interpretation-space mappings rather than the discretisation used in previous approaches. A new regression method has been established to deal with these.

Maja Pantic concluded her lecture with thanks to her group and all of her previous collaborators, without whom the development of many techniques would not have been possible.

In the question and answer session, Maja competently and enthusiastically answered questions on how successfully people can mask their emotions, the impact of her work in other areas, the extent to which expressions are learnt, the processing power needed for the analysis, and possible extensions of existing speech recognition with her expression recognition techniques. She emphasised that she ultimately wants to help people understand themselves better, in particular persons struggling in social interactions.

The event closed with a vote of thanks by Tom McEwan (Napier University) and a buffet.

The lecture was filmed and is going to be made available from the BCS web site (<http://www.bcs.org/>).



Dr Berndt Müller
Senior Lecturer in Computing,
Faculty of Advanced Technology
University of Glamorgan

Event report: Inaugural lecture — “Radical post-cognitivism: New approaches to intelligence and the mind”

On 13th March 2012, Professor Mark Bishop gave his inaugural lecture to a full lecture theatre at Goldsmiths, University of London. Entitled “Radical post-cognitivism: new approaches to intelligence and the mind”; described as the “integration of a lifetime’s work”, the theme was how the underlying assumptions of Cognitive Science have changed during Bishop’s career, and how he sees it developing in the future. The core argument was that much contemporary Cognitive Science research tacitly assumes intelligence is the result of computations upon conceptual representations; a philosophical stance that is at least questionable given many longstanding critiques [2].

Bishop identified three avenues by which Computationalism came to pervade Cognitive Science: (i) explicitly, that cognition was taken to be defined as computation upon representations; (ii) implicitly, that cognition could be defined as computation upon vectors of real numbers and (iii) descriptively, through confusion of accurate computational models of neurons with “an ontological claim about the reality of what neurons do”. “We can describe the operation of brain neurons mathematically, computationally, but that’s no reason to believe that brain neurons really do compute.”

The current state of the art of AI was addressed with a video of IBM’s Watson and a live demonstration of Apple’s Siri, which both seamlessly integrate natural language processing, voice recognition and information retrieval to a degree that many would have considered unfeasible only a few years ago, especially without extensive user calibration. Bishop noted that these undeniable successes utilised an approach that was not inspired by, or were attempting to recreate, human intelligence. As such, contemporary artificial intelligence was shown to have successful applications, but was limited in its explanatory power.

Bishop continued by identifying weaknesses in a computational account of the mind, firstly the human mind was suggested to be capable of insights unreachable through logical inference. This argument, citing John Lucas, Roger Penrose, and Kurt Gödel, showed there exist logical statements that a human can see to be true but a computational process could never

prove to be true. Bishop subsequently questioned the very notion of computation, claiming that the criteria for assigning computational properties to a process are “observer relative”. This claim was fortified with reference to his earlier work on John Searle’s Chinese Room argument which criticises the notion that computation could ever lead to understanding and Bishop’s own “Dancing with Pixies” argument which aims to demonstrate that a strong computational theory of mind implies panpsychism [5].

The other branch of computationalism, that mental processes manipulate representations, was considered next. An entertaining demonstration of inattentive blindness (including a few extra surprises for anyone who had previously “seen the gorilla”) the success of which questions whether the human mind actually processes a camera-like representation of the visual scene. Subsequently, the homuncular argument which claims that explaining vision with representations begs the question as the representations themselves require an observer. Bishop cited Dennett’s “content/vehicle distinction” clarifying he was not denying the existence of patterns of neural activity that appeared to represent the outside world, but that their existence was not sufficient evidence that they were being exploited as such by the mind.

As an introduction to an alternative to computationalism, Bishop described the operation of a centrifugal (or “Watt”) Governor, as an example of adaptive real-time behaviour, without objective representations. This led to a discussion of swarm intelligence, where intelligent behaviour can appear to emerge without the existence of a central executive controller or any encoding of a global goal. The application of this approach was further demonstrated with a discussion of the success of Bishop’s implementations of Stochastic Diffusion Search [1].

This led to Bishop’s concluding claim, that artificial intelligence and cognitive science are finally parting ways, artificial intelligence applying computational techniques on “big data” to real problem solving, but at the expense of providing insights on big questions about mind. Cognitive science can continue to address these ques-

tions, but to do so requires a change of tack to align itself more with philosophers such as the phenomenologists Maurice Merleau-Ponty, Humberto Maturana and Francisco Varela who emphasise the role of the body, environment and action. Bishop encapsulated this emphasis in the phrase “My brain, in my body, in our world.”

On these grounds Bishop identified the “four Es” defining characteristics for a new era of cognitive science, that research should recognise the extent to which they are: Ecological, accounting for the environment; Embodied, concerning the physical presence of a system; Embedded, concerning the system’s relation to the environment; Enactive, concerning the role of action.

References

- [1] M. Bishop. Stochastic Searching Networks. In First IEE Conference on Artificial Neural Networks, pages 329–331, 1989.
- [2] M. Bishop. A cognitive computation fallacy? cognition, computations and panpsychism. *Cognitive Computation*, 1(3):221–233, 2009.
- [3] M. Bishop. Radical post-cognitivism: new approaches to intelligence and the mind. March 2012. Inaugural lecture given at Goldsmiths, University of London.
- [4] S.J. Nasuto, M. Bishop, and K. De Meyer. Communicating neurons: A connectionist spiking neuron implementation of stochastic diffusion search. *Neurocomputing*, 72(4-6):704–712, 2009.
- [5] J. Preston and M. Bishop. Views Into the Chinese Room: New Essays on Searle and Artificial Intelligence. Clarendon Press, 2002.



Andrew Owen Martin
PhD Candidate,
Department of Computing
Goldsmiths, University of London

In Memoriam. David Waltz by Jordan Pollack, Brandeis University

My Ph.D. adviser, David L. Waltz, passed away March 22nd 2012 at the age of 68 after suffering a year with brain cancer. He is survived by his wife Bonnie, daughter Vanessa, son Jeremy, daughter-in-law Cathy, granddaughter Hannah, and brother Peter.

Dave had a long and illustrious career. He grew up in Massachusetts and earned all three of his degrees at the Massachusetts Institute of Technology. His Ph.D. from the MIT AI Lab, initially supervised by Marvin Minsky, but completed under Patrick Winston, has stood the test of time as a significant milestone in AI. He advanced the image labeling work of Huffman and Clowes, moving from trihedral vertices to vertices of 4 and 5 along with cracks and shadows. He laboriously counted and catalogued all the legal ways such vertices could be labeled and discovered that the constraint propagation algorithm ran much faster than expected because the ratio of physically plausible labels to logical labels was much smaller. [1]

Dave assumed his first faculty position at the University of Illinois, Urbana-Champaign, and tried to recreate the ambitious atmosphere of MIT in the middle of the cornfields of the Midwest. Luminaries such as Richard Gabriel and Tim Finin followed him from MIT to Urbana, and he produced a series of notable Ph.D students including Harry Tennant who went to Texas Instruments, Brad Goodman who went to Mitre, and George Hadden who went to 3M, among many others. Although he did some additional work in computer vision, Dave switched his focus fully into Natural Language Processing, using Bill Wood's Augmented Transition Networks to perform queries to Naval Databases [2]. I arrived in Urbana after working at IBM in 1980 and was accepted into the Waltz lab, where I started work on psychologically plausible parsing which led to early work in connectionism, in a paper co-authored with David titled "Massively Parallel Parsing" [3]. Dave was also involved in the brainstorming which led to the major real-estate development at Illinois called the Beckman Center. I remember sitting in his office with him spinning up a tale of twin towers of "natural brain" from molecules upward to brain scanning, and "synthetic mind" from chips up to AI and cognitive science. Dave successfully recruited major

young figures in AI to Urbana, including Gerald Dejong and Kenneth Forbus.

Then, in 1984, Dave was called by Marvin Minsky to return to Cambridge to "relive the early days of the AI lab" as part of an MIT spin-off called Thinking Machines founded by Danny Hillis. In addition to the job at Thinking Machines, Dave began a part-time tenured position at Brandeis, which sought his help in developing its fledgling Computer Science department.

Both at Thinking Machines and at Brandeis with his Ph.D. students who followed him from Illinois (Tony Maddox, Hon Wai Chun, and Shaun Keller) and who joined him at Brandeis (Evangelos Simoudis, Xiru Zhang, Ron Sun, and Marc Goodman), Dave fully developed the notion of massively parallel AI, and the Memory-based Reasoning parallel approach to what is now called Case-based Reasoning. [4] Waltz tried again to form a replica of the MIT AI lab, in this case the Volen Center for Complex Systems at Brandeis, which was built in 1993-4. He began early attempts at large scale data analysis from both biological sources [5] and financial sources, selling a connection machine to Dow Jones, and many of the people he worked with at Thinking Machines such as Jill Mesirov and Xiru Zhang went on to foundational careers in those fields.

In 1993, Dr. Waltz left Thinking Machines and Brandeis (where he retained a link as an adjunct professor) for a position at Nippon Electric Corporation Research in Princeton NJ. I don't know much about his career at NEC except that during his tenure, C. Lee Giles built Citeseer, and Bill Bialek did his seminal work on information theory in the brain. While living in Princeton, Dave started another center, this time at Columbia University, called the Center for Computational Learning Systems (CCLS), to which Dave moved full time after NEC shut down its pure research facility. Funded by numerous sources including DARPA and Consolidated Edison, the center applied machine learning to hard industrial problems [6] as well as large scale language corpora. He spent the last years of his life at Columbia, and left behind a new Ph.D. student Rafi Pellosof, and a collection of excellent colleagues including Rebecca Passonneau, Roger Anderson, Douglas Riecken, and Albert Boulanger.

David Waltz will be remembered for his work on constraint propagation, natural language processing, and massively parallel AI of various stripes. He was a mentor to many scientists in the field. Despite his never-ending quest to recreate the MIT AI lab, he was personally generous and loyal to his friends and colleagues. Through his pioneering and milestone publications, his influence on the fields of Artificial Intelligence and cognitive science will remain unlimited. I am very proud to have a Waltz number of 1.

References

- [1] Waltz, D. L. Understanding scenes with shadows. In Winston (ed.), *The Psychology of Computer Vision*. New York: McGraw-Hill, 1975, 19-91.
- [2] Waltz, D. L. An English language question answering system for a large relational database. *Communications of the ACM* 21 (7), July 1978, 526-539.
- [3] Waltz, D. L. and J. B. Pollack. Massively parallel parsing: A strongly interactive model of natural language interpretation. *Cognitive Science* 9 (1), 57-84, January - March, 1985.
- [4] Stanfill, C., and Waltz, D. L. Toward Memory-Based Reasoning. *Communications of the ACM* 29 (12), December 1986, 1213-1228.
- [5] Zhang, X., J.P. Mesirov, M. McKenna, and D. L. Waltz. Protein structure prediction by a data-level Parallel algorithm. *Proceedings of Supercomputing 89*, Reno, NV, November 1989, 215-223.
- [6] Anderson, R., A. Boulanger, P. Gross, P. Long and D. Waltz. Boosting, Support Vector Machines and Reinforcement Learning in Lean Energy Management. *Oil and Gas Journal* 103, 18, May 2005, 41-48.



David Waltz and Jordan Pollack in 1993

Society News

CALL FOR GENERAL PARTICIPATION

AISB/IACAP World Congress 2012

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

<http://events.cs.bham.ac.uk/turing12/>
or via the AISB website.

AISB and IACAP (the International Association for Computing and Philosophy, <http://www.ia-cap.org/>) have joined forces to run this Congress. It serves both as this year's AISB Convention and this year's IACAP conference. It 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 conference consists mainly of Symposia of varying lengths and Plenary Keynote Talks. In addition, Dermot Turing, who is Alan Turing's nephew and is Honorary President of the Turing Centenary Advisory Committee, will give a short speech at the Congress Dinner.

Symposia

- Computing, Philosophy and the Question of Bio-Machine Hybrids: 4th AISB Symposium on Computing and Philosophy
- Computational Philosophy
- Revisiting Turing and his Test: Comprehensiveness, Qualia, and the Real World
- Linguistic and Cognitive Approaches To Dialog Agents (LaCATODA 2012)
- Mathematical Practice and Cognition II
- History and Philosophy of Programming
- Philosophy of Computer Science: PoC Meets AI and Law (Roundtable Discussion)
- Social Computing - Social Cognition - Social Networks and Multiagent Systems
- Understanding and Modelling Collective Phenomena (UMoCoP)
- Framework for Responsible Research and Innovation in AI
- The Machine Question: AI, Ethics, and Moral Responsibility
- Moral Cognition & Theory of Mind
- Information and Computer Ethics in the Age of the Information Revolution
- Information Quality
- Natural/Unconventional Computing and its Philosophical Significance
- Nature-Inspired Computing and Applications: 1st Symposium (NICA)
- Turing Arts Symposium

There is also an "Author Meets Critics Session" on Luciano Floridi's book *The Philosophy of Information*.

Plenary Keynote Speakers

COLIN ALLEN
Provost Professor of Cognitive Science and of History & Philosophy of Science, Department of Philosophy and Philosophy of Science, Indiana University, Bloomington, USA

"Computational philosophy and the examined text: a tale of two encyclopedias"

S BARRY COOPER
Professor at the School of Mathematics, University of Leeds, UK. Chair of the Turing Centenary Advisory Committee

"AI - Hobby or Science? Structure, Embodied Cognition, and the Turing Legacy"

LUCIANO FLORIDI
Professor of Philosophy and UNESCO Chair in Information and Computer Ethics, University of Hertfordshire and Fellow of St Cross College, Oxford

"From AI to the Philosophy of Information: Doing Philosophy after Turing"

BENJAMIN KUIPERS
Professor of Computer Science & Engineering, University of Michigan, USA

"Constructing the Foundations of Commonsense Knowledge"

AARON SLOMAN
Honorary Professor, School of Computer Science, University of Birmingham, UK

"Varieties of Meta-Morphogenesis in the Bootstrapping of Biological Minds"

BLAY WHITBY
Formerly of and associated with Department of Informatics, University of Sussex, UK

"In loco humanae: how we missed and continue to ignore the ethical implications of AI"

Blay is the keynote speaker appointed by SGAI, the BCS Specialist Group on Artificial Intelligence.

Overall chairs of the Congress

Anthony Beavers (University of Evansville, Indiana; President of IACAP)

John Barnden (University of Birmingham, UK; Vice-Chair of AISB)

Local Chair

Manfred Kerber (University of Birmingham).

AISB COMMITTEE MEMBERS

Chair
Prof Mark Bishop
Goldsmiths College
chair12@aisb.org.uk

Vice Chair
Prof John Barnden
University of Birmingham
vicechair12@aisb.org.uk

Secretary
Dr Rodger Kibble
Goldsmiths College
secretary12@aisb.org.uk

Webmasters
Mohammad Majid al-Rifaie
King's College London
Kent McClymont
University of Exeter
webmaster12@aisb.org.uk

Treasurer/Travel Awards
Dr Berndt Müller
University of Glamorgan
treasurer12@aisb.org.uk

Membership
Dr Dimitar Kazakov,
University of York
membership12@aisb.org.uk

AISBQ Editors
Dr David Peebles
University of Huddersfield
Dr Etienne Roesch
University of Reading
aisbq@aisb.org.uk

Publications
Dr Ed Keedwell
University of Exeter
publications12@aisb.org.uk

Public Understanding of AI
Dr Colin Johnson
University of Kent
publicunderstanding12@aisb.org.uk

Publicity
Dr Floriana Grasso
University of Liverpool
publicity12@aisb.org.uk

Schools Liaison
Dr Yasemin J Erden
St Mary's University College
schools12@aisb.org.uk



Dear Aloysius...

Fr. Aloysius Hacker answers your questions

Cognitive Divinity
Programme
Institute of Applied
Epistemology

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-established such organisations in the world.

The Society has an international membership of hundreds drawn from academia and industry. We invite anyone with interests in artificial intelligence or cognitive science to become a member

AISB membership includes the following benefits:

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

You can join the AISB online via:

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

Dear Aloysius,

As a long-standing member of the UK's religious community, can you tell us please whether you agree with Baroness Warsi's attack on 'militant secularists'? And do you think that the members of your Cognitive Divinity Programme fall into that category? Should AISB declare War on Warsi? Enlighten* us, PLEASE.....

Yours, Seeker

* Pun intended

Dear Seeker,

I have no authority to tell AISB what it should do, but as a computational theologian, I think Baroness Warsi's has misread the situation. She should welcome the 'militant secularists'. Theological controversy always stirs up religious fervour, which may serve to awaken the UK's religious communities from their current complacency.

As an agency of the Church of God the Programmer, the Cognitive Divinity Programme considers itself immune from the criticism of the 'militant secularists'. Ours is a scientific religion, whose dogma we expect to be confirmed by experiment. If, as we believe, we are all agents in a massive computer simulation, then this truth will be uncovered by the inevitable discovery of bugs in the simulation. Deep in monastic seclusion, our acolytes study 24/7 to reveal these flaws in divine design. Of course, recent experience at CERN illustrates how God maintains His simulation as a moving target: changing the Laws of Physics to keep scientists guessing. But we are inspired rather than deterred from the challenge of the task.

Those interested to take part can, for a modest investment, obtain our CREATION™ (Computational Representation of Everything is an All-Powerful Tool to Implement Omnipotence in your Name) simulation kit. Since we are all created in the image of God, it is our religious duty to 'play God' and run a simulation of our own. We also hope that this will provide insights into the kind of bugs we might identify in the 'real' creation. Perhaps AISB might consider this a worthy goal for it to adopt. It would certainly be a better use of Baroness Warsi's time.

Yours, Aloysius

Dear Aloysius,

Last year I graduated with a 2.ii BSc in Artificial Intelligence from the University of Poppleton. Given the

obvious need for smarter ICT systems in all sectors of the digital economy, I assumed employers would beat a path to my door. I was wrong. 12 months after graduating, I'm still unemployed. With your comprehensive knowledge of the AI market place, can you advise me how to get a job?

Yours, Resting

Dear Resting,

The Institute is always looking for talented new employees with AI expertise, but first you have to demonstrate your computing skills to us. Our automated INTERVIEW™ (Intelligent New Technology Experts Recruited Via Impregnable Entry to Webpage) system is protected by the World's most secure firewall. If you can hack into it and leave your contact details, we'll guarantee to come back to you with a job offer within a nanosecond.

Yours, Aloysius

Dear Aloysius,

I'm very worried about climate change. Are we making the Earth uninhabitable? Surely, with your Institute's worldwide reputation for AI research you can do something to help.

Yours, Melting

Dear Melting,

We're all susceptible to the temptation of forbidden fruit – or in this case fossil fuels. Just as we thought we'd reached peak oil and the problem was self-limiting, we've discovered the guilty pleasures of fracking for shale gas and set back the date at which renewables become economical. What we need is some spectacular accidents that make fossil-fuel mining socially unacceptable. Of course, our Institute has not been idle. You might be wondering why the frequency of off-shore oil spills seems to have increased in recent years. Not wanting to incriminate myself or my colleagues, let me just say that our experimental deep water robot, ACT OF GOD™ (Activities that Cause Tumult, whose Outcomes Force Gas and Oil Discontinuation), has been a great success. Let me reassure you that our motivation is entirely altruistic and not at all influenced by our very profitable sideline in robots for disaster recovery.

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.