AISB QUARTERLY

THE NEWSLETTER OF THE SOCIETY FOR THE STUDY OF
ARTIFICIAL INTELLIGENCE AND SIMULATION OF BEHAVIOUR

No. 135 October, 2012

Editorial

It is with great pleasure that we present the 135th issue of the AISB Quarterly in its new format. The "Q" (as it is affectionately known by many AISB members) had been in its previous format since Autumn 2001, and the AISB Committee decided that a fresh design and a smarter format was due to complement the Society's freshly refurbished website (http://www.aisb.org.uk). With this new edition we have also moved to LaTeX in order to ease, and hopefully speed up, the production process. Submission guidelines will be made available on the website shortly.

This issue contains four original articles. The first is a description of Terry Stewart's *Neural Engineering Framework*, a novel connectionist approach to the design of large-scale cognitive models while in the second, Dan Veksler explores the possibility of using the online virtual world *Second Life* as a cognitive modelling platform.

The third article is a review, by a collective of industrial and academic partners, of the AI-based *SpendInsight* platform, devised to help higher education institutions optimise their procurement process and, ultimately, save money. Finally, Martin Wurzinger responds to Pat Langley's article in AISBQ 133 by considering the process and achievements of AI, emphasising the benefits of AI applications that go beyond the scientific domain.

This issue also contains several reviews. Andrew Martin reviews J. Kevin O'Regan's book *Why Red Doesn't Sound like a bell*, and reports on the first AISB members' workshop on sensorimotor theory. We also have reports of two recent events: Marilyn Panayi reports on the *Foundations of Enactive Cognitive Science* workshop, and Anna Dumutriu gives us a taste of the *Alan Turing Centenary Exhibition* at the AISB/IACAP World Congress.

You will find a call to submit work to the second AISB workshop of the academic year—this time in Exeter—in December, focusing on hyper-heuristics and new classes of optimisation problems.

And finally, AISBQ would not be complete without Fr. Aloysius, our hacker in residence and renowned columnist who, as usual, concludes the issue by responding to your most pressing queries.

We would like to take this opportunity to thank all members of the Society, for your continued support and the constant flow of submissions that make the Q the vibrant magazine it is, contributing to the exciting life of the society. Thank you!

The Q editors

The Neural Engineering Framework

The Neural Engineering Framework (NEF) is a general methodology that allows the building of large-scale, biologically plausible, neural models of cognition [1]. The NEF acts as a neural compiler: once the properties of the neurons, the values to be represented, and the functions to be computed are specified, it solves for the connection weights between components that will perform the desired functions. Importantly, this works not only for feedforward computations, but also for recurrent connections, allowing for complex dynamical systems including integrators, oscillators, Kalman filters, etc. [2]. The NEF also incorporates realistic local error-driven learning rules, allowing for the online adaptation and optimisation of responses [3]. The NEF has been used to model visual attention [4], inductive reasoning [5], reinforcement learning [6] and many other tasks. Recently, we used it to build *Spaun*, the world's largest functional brain model, using 2.5 million neurons to perform eight different cognitive tasks by interpreting visual input and producing hand-written output via a simulated 6-muscle arm [7,8]. Our open-source software *Nengo* was used for all of these, and is available at http://nengo.ca, along with tutorials, demos, and downloadable models.

Motivation

Despite the additional constraints and computational overheads involved in building biologically plausible models, there are two major reasons for doing so. First, using biologically realistic neurons not only allows the modelling of behaviour, it also allows the comparison of network properties (e.g., firing patterns, timing effects, and neural connectivity) with real brains and the potential for more accurate investigation of neural degeneration, lesioning, deep brain stimulation, and even various drug treatments.

As an example, when we constructed a NEF implementation of a production system constrained by the properties of the various neuron types found in the brain regions involved, it not only produced the classic 50 millisecond cognitive cycle time *without parameter fitting* [9], it also produced a novel prediction that some types of productions take ~40 milliseconds, while others take ~70 milliseconds, which matches well to some unexplained behavioural data [10].

The second reason for building biologically plausible models is that it can suggest new types of algorithms. The NEF does not produce an exact implementation of whatever algorithm you specify but an approximation, the accuracy of which depends not only on the neural properties but also on the functions being computed. As a consequence, the computations used in a NEF model are constrained by the basic operations of neurons. This has allowed us to make strong claims about the classes of algorithms that cannot be implemented in the human brain (given the constraints on timing, robustness, and numbers of neurons involved) [11].

For example, in attempting to find a plausible implementation of symbol-like cognitive reasoning we were led towards a relatively unexplored family of algorithms which, upon further investigation, we discovered to be particularly useful for induction and pattern completion tasks that are difficult to explain with classical symbol structures [5,11].

In this article I will outline how vectors are encoded into the distributed activity of population of neurons, and how to interpret that activity back into a vector.

Representation

The NEF uses distributed representations and draws a sharp distinction between the activity of a group of neurons and the value (usually thought of as a vector $\mathbf{x}$) being represented. For example, 100 neurons may represent a 2D vector, with different vector values corresponding to different patterns of activity across those neurons.

To map between $\mathbf{x}$ and neuron activity a, every neuron i has an encoding vector $\mathbf{e}_i$ which can be considered the *preferred direction vector* for that neuron (i.e., the vector for which that neuron will fire most strongly). This fits with the general neuroscience methodology of establishing tuning curves for neurons, where the activity of a neuron peaks for some stimulus or condition. The NEF embodies the strong claim that the input current to a neuron is a linear function of the value being represented. If G is the neural non-linearity, α_i is a gain parameter, and β_i is the constant background bias current for the neuron, the neural activity given $\mathbf{x}$ is $a_i = G\left(\alpha_i \mathbf{e}_i \cdot \mathbf{x} + \beta_i\right)$

Importantly, G can be any neural model, including simple rate-based sigmoidal neurons, spiking Leaky-Integrate-and-Fire neurons, or more complex biologically detailed models. The only requirement is that there be some mapping between input current and neuron activity, which can include complex spiking behaviour.

While a vector $\mathbf{x}$ can be converted into neural activity a_i, it is also important to do the opposite. Finding $\mathbf{x}$ given a_i provides a measure of accuracy and a high-level interpretation of spiking activity. NEF does this by finding a set of decoding weights $\mathbf{d}$ such that $\mathbf{x} \approx \sum a_i \mathbf{d}_i$. These weights can be found using any standard error minimization technique. Crucially for the NEF, these weights are also used to directly solve for the neural connection weights that perform computations.

Computation

For the NEF, any connection between groups of neurons computes a function. The trick is to find a set of connection weights ω_{ij} such that if the first group of neurons (A) represents $\mathbf{x}$ then this will cause the second group (B) to represent $y = f(x)$.

The first step is to imagine an intermediate group of perfectly ideal linear neurons (Figure 1a), with one neuron per vector dimension. If we connect A to these ideal neurons using the connection weights $\mathbf{d}$ found above, then these ideal neurons will be driven to represent the vector $\mathbf{x}$. In general, we can also optimize $\mathbf{d}$ to approximate any function $f(x)$ by adjusting the error minimization. We then connect the ideal neurons to B using the encoder values for group B ($\mathbf{e}_j$). This causes group B to also represent $f(x)$.

Once $\mathbf{d}$ and $\mathbf{e}$ have been found, the intermediate layer is then removed, directly connecting A to B by multiplying the two sets of weights (Figure 1b). This produces the optimal weights to compute an arbitrary function $f(x)$ between A and B.

This approach can be used to approximate any function and has the valuable property that non-linear functions can be computed with a single layer of connections—no back-propagation of error is required. Not every function can be computed however; the more nonlinear and discontinuous the function, the lower the accuracy. Accuracy is also affected by the neuron properties, such that having a wide variety of neural parameters (as in biological neurons) greatly increases accuracy. This also allows you to determine the neuron properties that would be ideal for particular computations, which can then be used as neurological predictions [1].

It is important to note that the NEF cannot only produce biologically realistic models capable of computing functions of the form $y = f(x)$ but it can also compute dynamic functions of the form $\frac{dx}{dt} = A(x) + B(u)$, where x is the value being represented, u is some input, and A and B are arbitrary functions. A particularly useful special case of this equation is $\frac{dx}{dt} = u$, (an *integrator*) as this sort of component appears in many models of working memory and in accumulator models of decision making.

Symbol Processing

While manipulating vectors is extremely powerful, many cognitive algorithms rely on manipulating symbols with some sort of syntactic structure. How can neurally realistic models possibly represent some-

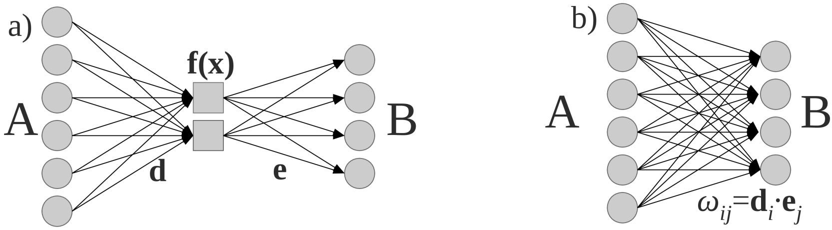

Figure 1: Connecting populations of neurons (circles) via idealized perfectly linear components (squares). (a) **x** is computed from α_i using weights **d**. **x** is then combined with **e** to compute the input current to the next layer of neurons B. (b) Idealized components are eliminated, giving a realistic neuron model functionally identical to (a)

thing like "Dogs chase cats" in such a way as to distinguish it from "Cats chase dogs"? How can we manipulate these representations in useful ways?

It turns out that there are a family of models that already exist for converting symbolic logic into vector manipulations. These are known as Vector Symbolic Architectures [12], and all follow the approach of using high-dimensional vectors for each basic symbol, and then combining these vectors with various mathematical operations to produce new vectors that encode full symbol structures. Unlike ideal classic symbol systems however, VSAs are lossy, in that as the symbol tree structure gets more complex, the accuracy of extracting the original vectors from that combined vector gradually decreases.

Furthermore, the vectors maintain similarity, so that if "pink" and "red" have similar vectors, then "pink square" and "red square" will also have similar vectors. This feature allows inductive reasoning over complex patterns. For example, our neural model of the Raven's Progressive Matrix task (a standard intelligence test where participants are given 8 visual patterns in a 3×3 grid and are asked to determine what pattern should be placed in the missing square) works by forming the vector representation of each pattern and computing the average transformation that that will take one pattern to the next [5].

As a simple example of this approach, you can create high-dimensional (~500 dimensions for adult-level vocabularies) unit vectors for each basic symbol (DOG, CAT, CHASE, SUBJECT, OBJECT, VERB, etc.). These can be chosen randomly, or so as to reflect standard similarity measures. Two operations are required to create a symbol structure: addition (+) and circular convolution (⊗). The sentence “Dogs chase cats” would then be S = DOG ⊗ SUBJECT + CHASE ⊗ VERB + CAT ⊗ OBJECT). Given this sentence, a particular component can be extracted by computing S ⊗ SUBJECT^{-1} ≈ DOG, where the inverse operation is a simple reordering of the elements in the vector. Interestingly, while circular convolution seems like a complicated operation, you can break it down into a linear transformation, a large number of pairwise multiplications, and another linear transformation. All of these operations are accurately approximated by the NEF methods.

Spaun

The ability to perform symbol-like manipulations using vectors allows you to build very large-scale cognitive models. Our largest model to date is Spaun, a 2.5 million spiking neuron model with a vision system (formed by implementing a Restricted Boltzmann Machine Deep Belief Network with the NEF), a single 6-muscle 3-joint arm for output, and a selective routing system (analogous to a production system) implemented in spiking neurons comprising the cortex (for working memory storage), the basal ganglia (for action selection), and the thalamus (for selectively routing information between cortical areas) [7]. Various other cortical areas are also modeled, allowing for transformations between visual, conceptual, and motor spaces, inductive pattern finding, and list memory. The model is capable of performing eight different psychological tasks, including recognizing hand written digits, memorizing digit lists and recalling particular items, pattern completion, reinforcement learning, and mental addition. No changes to the model are made between tasks: instead, a visual input is provided telling the model which task to perform next. We are aware of no other realistic neural model with this combination of flexibility and biological realism.

Nengo

Nengo is an open-source cross-platform Java application which implements the NEF and can be used as both a teaching tool (with hands-on classroom demos) and a research tool (all of our large-scale models are built with it, including Spaun). Neural groups can be created through a drag-and-drop interface or Python scripting. The functions to approximate are similarly specified, with Nengo automatically computing the optimised connection weights. Also included is a visualisation interface for viewing and interacting with running models, including support for simulated environments and physical robots. The software, extensive documentation, and various tutorials are available at http://nengo.ca.

References

1. Eliasmith, C., & Anderson, C. H. (2003). *Neural engineering: Computation, representation and dynamics in neurobiological systems.* Cambridge, MA: MIT Press.

2. Eliasmith, C. (2005). A unified approach to building and controlling spiking attractor networks. *Neural computation. 7*, 1276–1314.

3. MacNeil, D., & Eliasmith C. (2011). Fine-tuning and the stability of recurrent neural networks. *PLoS ONE. 6(9).*

4. Bobier, B., Stewart T. C., & Eliasmith C. (2011). *The attentional routing circuit: receptive field modulation through nonlinear dendritic interactions.* Cognitive and Systems Neuroscience Poster.

5. Rasmussen, D., & Eliasmith, C. (2011). A neural model of rule generation in inductive reasoning. *Topics in Cognitive Science, 3*, 140–153.

6. Stewart, T.C., Bekolay, T., & Eliasmith, C. (2012). Learning to select actions with spiking neurons in the basal ganglia. *Frontiers in Decision Neuroscience, 6.*

7. Stewart, T., Choo, F-X, & Eliasmith, C. (2012). Spaun: A perception-cognition-action model using spiking neurons. *Proceedings of the 34th Annual Conference of the Cognitive Science Society.*

8. Eliasmith, C. (2013). *How to build a brain: A neural architecture for biological cognition.* New York: Oxford University Press.

9. Stewart, T.C., Choo, F-X., & Eliasmith, C. (2010). Dynamic behaviour of a spiking model of action selection in the basal ganglia. In *Proceedings of the 10th International Conference on Cognitive Modeling*, 235–240.

10. Gunzelmann, G., Moore, R., Salvucci, D., & Gluck, K. (2011). Sleep loss and driver performance: Quantitative predictions with zero free parameters. *Cognitive Systems Research, 12(2)*, 154–163.

11. Stewart, T., & Eliasmith C. (2012). Compositionality and biologically plausible models. *Oxford Handbook of Compositionality.*

12. Gayler, R. (2003). Vector symbolic architectures answer Jackendoff's challenges for cognitive neuroscience. *ICCS/ASCS International Conference on Cognitive Science*, Sydney, Australia: University of New South Wales. 133–138.

Terry Stewart PhD
Centre for Theoretical Neuroscience, University of Waterloo, Canada

Second Life as a Simulation Platform

The aim of this article is to introduce Second Life as a simulation platform for cognitive systems development. There are roughly four general types of simulation platforms that may be employed for Cognitive Modelling and AI: robotics, high-fidelity synthetic environments (those bound by real-world dynamics and constraints, i.e., virtual reality), traditional games and tasks (e.g. chess, Tetris, Markov decision tree navigation), and behavioural laboratory tasks where the goal of the synthetic agent is to match empirical results from human/animal studies moreso than to optimise performance. All of these are valuable tools that offer distinct benefits. However, the choice of a task environment is non-trivial in the development of synthetic cognitive systems. Indeed, the task often drives the direction of the mechanisms built to address it. For instance, the planning requirements of chess gave rise to GPS/SOAR, recognition-type problems gave rise to various categorisation models and neural networks, and the focus on perception/action constraints gave rise to EPIC. A high-fidelity problem has the potential to drive a high-fidelity solution.

Traditional laboratory tasks and low-fidelity simulation environments are small worlds [1, 2]—they may be too limited or too contrived, and may lead to both misses and false alarms in cognitive system development. A synthetic agent may perform perfectly in the Wumpus world or closely match empirical results from the Stroop task, yet be incapable of generalising to other tasks or scaling to high-fidelity environments. Simpler tasks also suffer from the identifiability problem, where two opposing approaches to cognition can produce indistinguishable behavioural results (e.g. [3]).

On the other hand, robotics offers a large world [1, 2] high-fidelity simulation platform that bares all the complexity of the real world, but may lack the simplicity and control needed to develop, debug, and scale a cognitive system in a reasonable amount of time. Robotics work includes many technical challenges at the sensory-motor level before cognitive mechanisms may be addressed. In addition to tasking resources, these demands confound results—when a robot does not behave appropriately, it is difficult to tell whether the lack in performance is due to cognitive limitations or the lo-fi sensory-motor apparatus.

Virtual reality can offer real-world complexity and constraints to a greater degree than simple games or laboratory tasks, and more simplicity and control than robotics. It

Figure 2: Screenshot of a synthetic agent in the Second Life™ environment.

provides a balance between a small world and a large world environment (i.e., a controlled large world).

Second Life

Second Life™ is a large 3D virtual world populated by hundreds of thousands of online users and millions of objects. This synthetic environment has the potential to be a useful tool for cognitive modelling and artificial intelligence. It has the same draw as robotics in that it is a high-fidelity, multi-purpose, persistent, large, complex, uncertain, noisy, and non-stationary environment that can be useful to examine how a synthetic mind might scale beyond simple games or laboratory tasks. Like other virtual worlds, the Second Life platform can offer researchers more control than robotics, allowing for faster development and debugging, and gradual increments in complexity. Additionally, Second Life has some significant advantages over other virtual reality platforms, primarily due to its size. Second Life contains replicas of major cities and desert villages, small bars and large museums, fields and mountains, as well as historical and fictional places. Using Second Life may free up valuable resources to allow researchers to focus on developing cognitive systems, rather than building a virtual world, or dealing with mechanical failures of a robotic system.

My current work involves the construction of Second Life objects as TCP web servers and cognitive models as TCP web clients. In this way, the Second Life object may be thought of as a remotely located 'body' being controlled by a locally operating synthetic mind. While simple, this setup is limited in the number of TCP connections per server, and in the TCP request rate (throttled to about 2 sensory-motor cycles per second).

Regardless of the sensory-motor speed, Second Life, by design, will run no faster than real-time. This may be unacceptable for much research in AI and cognitive modelling as evaluating a large parameter space may take on the order of months or longer. As with robotics therefore, Second Life can only be employed for simulations after a significant reduction of parameter space.

Second Life as a complement to laboratory tasks and traditional simulation environments

My first experience with Second Life involved the comparison of Reinforcement Learning-based (RL) and Associative Learning-based (AL) decision models [4]. Prior evaluation of the two models using standard laboratory tasks and simple Markov-decision simulation environments revealed unambiguous advantages for the AL model. The AL model fit human data better than RL across two experiments, and maze navigation simulations revealed performance advantages for AL over RL over three mazes ranging in difficulty. However, examination of the models in Second Life (using a simple navigation task) revealed flaws in the AL model.

Employing Second Life revealed what simpler simulation environments could not—that the implemented standalone AL model could not scale. This led to further AL model development and an effort to integrate AL and RL [5]. In sum, Second Life offers real-world-like complexity and constraints that simpler task environments do not. For instance, approaching a landmark from the north does not present the same view as approaching it from the south, objects may be mobile, and movement is noisy. These features have the potential to make Second Life an important tool for cognitive system development and evaluation.

An alternative to robotics

In my experience with teaching cognitive system development, students using Second Life as a simulation platform greatly outpaced those using robotics. In the same time that robotics students created a simple 2-action Reinforcement Learning agent, the students in the Second Life class were able to learn a new programming language (LISP), develop two chat bots (one matching on wikipedia keywords, another matching chat content to news articles based on Latent Semantic Analysis), and develop a 4-action Reinforcement Learning bot.

This slower progress in the

robotics class was due to a number of issues. For example, robots move at different speeds depending on carpeting, constraining the robots to left-right rotation can be difficult as the robot will slowly move off centre, battery capacity largely limits simulation times, colour parsing is sensitive to lighting conditions (which change throughout the day and when people walk by), and servos die out after heavy use. In comparison with Second Life, the robotics platform offers too little control over environmental complexity, simulation length, and overall speed of execution and debugging. In sum, if a virtual environment can provide the embodiment constraints necessary for the warranted cognitive system evaluation, it may be prudent to forego the burdens of robotics.

User-generated content

For the final project in the Second Life class mentioned above, students were required to build synthetic agents for a foraging task in a rich environment. The agents were embodied as fish under water, learning the layout of the world, and learning which objects in the world may be edible. This world was full of interesting terrain, coral, scripted fish, rocks, plants, and artifacts—none of which were created by me. To set up this task-environment I created a feeder (a simple scripted object that generated food pellets) and scripts for how students' agents may consume the food pellets. The rest of the content was already there, generated by Second Life users.

To this point, Second Life is full of user-generated content. It is the largest of such virtual environments, offering synthetic agents the ability to interact, create, and explore persistently for years, rather than minutes. To employ a smaller, less-populated virtual environment may be equivalent to opting for a small website over the world wide web.

Summary and conclusions

The development of cognitive models and agents is often driven by the objectives of the target environment. A high-fidelity task environment can drive high-fidelity cognitive system development. Traditional laboratory tasks and games are small world environments that cannot offer the necessary real-world complexity and constraints. Large world simulation platforms (robotics or virtual reality) can help to correct the direction of model or agent development, and should be considered a necessary complement to empirical modelling and simpler task simulations. Virtual reality frameworks, in particular, are controlled large worlds, allowing the researcher to slide the scale between control and complexity as needed. Second Life is a highly populated multi-purpose virtual world containing a lot of user-generated replicas of real and fictional places. Although this platform has many limitations, and cannot run faster-than-realtime simulations, it may be a useful tool for cognitive system development.

Acknowledgments

A part of this research was performed while the author held a National Research Council Research Associateship Award at the United States Air Force Research Laboratory, Cognitive Models and Agents Branch.

References

1. Binmore, K. (2009). Rational Decisions. Princeton University Press.

2. Savage, L. (1954). The Foundations of Statistics. New York: Wiley.

3. Howes, A., Lewis, R. L., & Vera, A. (2009). Rational adaptation under task and processing constraints: implications for testing theories of cognition and action. Psychological Review.

4. Veksler, V. D. (2010). Examining the Goal-Proximity Decision Mechanism: Maze-navigation, incidental learning, Tic-Tac-Toe, and exploration of Second Life virtual worlds. Doctoral Dissertation, Rensselaer Polytechnic Institute.

5. Veksler, V. D., Myers, C. W., and Gluck, K. A. (2012) An Integrated Model of Associative and Reinforcement Learning. Proceedings of the 34th Annual Meeting of the Cognitive Science Society.

Vladislav "Dan" Veksler, PhD
Cognitive Models & Agents Branch,
US Air Force Research Laboratory

The Use of AI Technology in E-commerce Spend Analysis

At the launch of the software industry standard for green data interchange at the RSA London on 15 November, 2011, Ronald Duncan, founder and Technical Director of @UKplc (a cloud-based e-commerce company) reported on how classical AI technology in classification and matching was being exploited in a new 'spend analysis' system. The core technology—deploying Bayesian inference; symbolic rule-induction; decision trees and expert systems—has been commercialised by @UKplc as the "SpendInsight" platform which enables procurers to identify and cost equivalent products (with a subsequent extension offering an analogous carbon analysis of purchasing decisions). Procurers can thus analyse both the economic and environmental cost of purchases. Unfortunately, to date, the potential economic savings are not being realised. However, a survey reported herein, shows that the potential environmental benefits may well lead to changes in procurement processes resulting in both environmental and economic savings.

Investigations into search and natural language processing carried out during a large three year research project jointly hosted by the University of Reading, Goldsmiths College, and @UKplc, led to the development and commercialisation of SpendInsight, a radically new technological platform that enables organisations to monitor and analyse their purchase spend accurately and subsequently (via the related GreenInsight system) their procurement carbon footprint.

The 'spend analysis' e-commerce technology has now been successfully deployed in the UK education and health sectors where it has proven to be both cost effective and to deliver real savings; indeed the spend analysis system enabled a recent, widely publicised, UK National Audit Office report to highlight potential NHS purchasing savings of £500m per annum (a fact raised in parliament and widely discussed in national media).

Nevertheless, given the huge scale of the potential savings it highlights, it is perhaps a little surprising that to date, take up of this AI based spend analysis system remains relatively sluggish, (e.g., only 54 of over 400 UK Health Trusts have yet deployed the system). Reasons for the relatively poor roll-out of this technology are varied. However the underlying problem is one of organisational inertia.

For example, a procurement officer who may have invested significant time and effort building up a relationship with a supplier to negotiate a new contract for 10,000 syringes, which he has just guided through her line manager, may be slightly reluctant to embrace an alternative supplier offering a slightly cheaper deal. However, if the alternative supplier can compete on price and also be qualitatively different (for example also offering a potential reduction in the organisation's carbon footprint) then there may more motivation for organisational change.

This article reports a survey of attitudes to individual and organisational change amongst civil servants involved in purchasing process (a survey of 229 individuals carried out for @UKplc at the 'Civil Service Live' conference, Olympia 7–9 July, 2011). Results from the survey strongly support our hypothesis of organisational inertia and confirm that, whilst individual buying behaviour was strongly influenced by potential savings, only 57% of the respondents reported that savings alone were enough to 'probably or definitely' change organisational purchasing decisions.

Conversely, by highlighting economic savings and improved sustainability (lower procurement carbon footprint) there was a significant increase, with 84% of respondents now assessing that their organisation would 'probably or definitely' change purchasing choices. This data aligns well with experience from roll-out of the current NHS 'Carbon Footprint project', which is attracting much wider engagement than earlier, purely 'economically focused', spend analysis.

Clearly sustainability issues are a strong motivating factor to change purchasing behaviour, hence linking financial spend analysis with green analysis speeds up and unblocks process change. Thus, given the large potential savings identified by state-of-art spend analysis, compared with the relatively small incremental cost (per product) of carbon off-setting, consideration of green issues may result in substantial economic benefit to the organisation. At the societal level, a strong green national policy agenda may realise significant benefit for both the environment and the economy.

Richard Barraclough[1] Victor M. Becerra[2] Mark Bishop[3] Matthew Brown[1] Sebastian Danicic[3] Ronald Duncan[1] John Howroyd[3] Richard Mitchell[2] Slawomir Nasuto[2] Paul Roberts[1] Virginie Ruiz[2]

1. @UKplc, Aldermaston, Reading, UK.

2. School of Systems Engineering, University of Reading, Reading, UK.

3. Dept. Computing, Goldsmiths College, London, UK.

Quo Vadis, AI?

In a recent article entitled "Artificial Intelligence and Cognitive Systems", Pat Langley summarises the history of AI by pointing out that its original goal, the simulation of the mind, has been largely shelved over the past decades [1]. I would like to take this further by firstly, focusing on the essential difference between what was and was not pursued, secondly, describing what that difference entails, and thirdly, by raising a few questions which would suggest themselves once a mind model is in operation. My own background is in mind research and the development of computer programs that simulate cognitive dynamics [2]. What follows has been written from that perspective.

Current mainstream

Looking over the past ten years of the Quarterly and considering presentations of models that incorporated some form of cognitive simulation, there were 71 that were based on a formalism designed by its creators (a top-down approach), but only 11 which allowed rule sets to emerge from within, or were at least suggested (a bottom-up approach) [3]. A random example of the former would be Jamnik's "Informal human mathematical reasoning" [4], one of the latter is Stillwaggon's "How does a-life inform the mind-body problem?" [5]. Overall a ratio of over 6:1—no wonder Langley misses the "real" AI!

Top-down vs bottom-up

There are a number of essential aspects that speak against a top-down approach. The idea that the system of mind—and hence any simulation thereof—should be based on a series of preconfigured algorithms becomes questionable as soon as the sheer variance in cognitive structures observed in history and the presence is taken into account. There may be frameworks that allow one to identify a certain religion, or philosophy, or an overall paradigm such as predicate logic or a programming language, but they have been developed *a posteriori*, and they all are a product of the same old neurons and transmitters and dendrites we have been stuck with for all those millennia. Clearly, if there is an identifiable system, its functional granularity must be of the highest order.

Furthermore, during the processing of data, any rule set that can be identified from then onwards must have emerged from within, otherwise we would need to postulate some decision-making agent that is responsible for that specific configuration. But this is impossible because the process is still underway—

the eventual idea and/or concept has not been developed yet. Therefore not only does the brain possess a high degree of granularity, the functional granularity of the mind (the system sitting atop the physicality of the brain) has to be of a similarly high degree.

Another aspect concerns the specificity of situatedness. Quite apart from the diversity of cultures in the past and present, the exposure to a particular environment spurns the mind to adopt appropriate thought structures in order to process the data pertinent to that environment. So much so that in the absence of social influences per se there is hardly any process which allows us to identify the owner as typically 'human', as research into abandoned children has so abundantly demonstrated [6].

While basic behaviour forms are taken care of (to be expected, given that we are all mammals), the eventual maturing of the mind makes it relatable to a very particular environment where any common denominators suggested by construed algorithms cannot be found. Just consider the mindset of a medieval knight, a KGB agent, and a Buddhist monk.

Keedwell comes close regarding the question of bottom-up emergence when comparing DNA with computer code [7]. He argues the code in the computer is not known until it is executed, just as the DNA is not known until the organism is formed. However, while the code in a software may not be explicitly "known", there is nevertheless a direct relationship between its statements and what the program does at any given time. The information stored in DNA on the other hand is sufficiently compressed for such an explicit link to be impossible.

Onwards from chaos

A system that exhibits definitive patterns but allows variance, that iterates through essential processes but has the potential to become as complex as its resources allow, is a chaotic one. A simulation that processes data such that every node is performing under the above auspices and therefore permits relationships with others to be formed that eventuate in clustering (or the destruction of clusters for that matter), and in which the thus created affinity relationships become representative of input, exhibits cognitive dynamics such as memory, learning, and seemingly wilful behaviour (I say "seemingly" because space does not allow to enter the age-old argument of free will, especially in the current context).

The phenomenon of latency can also be observed, which relates to the storage of implicit information mentioned above in the case of DNA. For example, input A prompts a node cluster to produce some output, A1. Some input B produces another output, B1. Some particular input C lets A1 re-emerge, although without input A input C does not produce A1. In terms of incoming data triggering responses that relate to previous content the results are similar to those obtained from investigations into human memory (e.g., Kokinov [8]). In any case, the node clusters responsible demonstrate a latency with respect to previously imported data, although there is no one-to-one relationship between the nodes' states and the content–the states are merely representative of the content.

Such a scenario leads to further questions. If representative content can be "packed" into neurons and/or nodes, what is the upper limit for any given volume, and does it depend on the degree of affinity between types of content? What about intersections, in other words nodes that are shared between two or more clusters? Conceptually one could compare them with abstractions, the principles shared between ideas. By what increments does intelligence grow when increasing the number of nodes and their connectivity?

As the last item alone demonstrates, the parallelism between the biological and the artificial system leads into challenging regions. The relative complexity of neuronal systems in several species as pointed out by Elston [9] confirms the tests on the computer models where a greater number of nodes and a higher degree of connectivity among the nodes generates better outcomes in terms of learning and adaptation [10].

On a much larger scale the usefulness of the model could be shown when the expected outcomes in the Iraq war were questioned based on an analysis of that demographic's cognitive dynamics, done in 2003, and later confirmed in two reports, published in 2006 and 2007 respectively [11].

While the results from mainstream AI have certainly been worthwhile, its original goal should not be quietly forgotten. There are benefits to be had that go way beyond the purely scientific realm.

References

1. Langley, P., (2011). Artificial Intelligence and Cognitive Systems, AISBQ No. 133 (pp. 1-4), Society for the Study of Artificial Intelligence and Simulation of Behaviour, London.

2. Wurzinger, M., (2003–2012). http://www.otoom.net/ see pages: otoomcmprogram.htm, omoprogram.htm, owormprogram.htm, ovideoprogram.htm.

3. Wurzinger, M., (2012). AISBQ statistics, http://www.otoom.net/AISBQStats.xls.

4. Jamnik, M., (2003). Informal human mathematical reasoning, AISBQ No. 114 (p. 3), Society for the Study of Artificial Intelligence and Simulation of Behaviour, London.

5. Stillwaggon, L., (2006). How does a-life inform the mind-body problem?, AISBQ No. 123 (p. 8), Society for the Study of Artificial Intelligence and Simulation of Behaviour, London.

6. Newton, M., (2002). Savage Girls and Wild Boys, Faber and Faber Limited, London.

7. Keedwell, E., (2011). Artificial Intelligence and the Frontiers of Genetics Research, AISBQ No. 131 (pp. 3-4), Society for the Study of Artificial Intelligence and Simulation of Behaviour, London.

8. Kokinov, B.N., (2000). Integration of Memory and Reasoning in Analogy-Making: The AMBR Model, CEE Centre for Cognitive Science, New Bulgarian University, Bulgaria.

9. Elston, G., (2003). Cortex, Cognition and the Cell: New Insights into the Pyramidal Neuron and Prefrontal Function, Cerebral Cortex November V 13 n 11 (pp. 1124-1138), Oxford University Press, Oxford.

10. Wurzinger, M., (2006). OWorm tests, http://www.otoom.net/owormtests02.zip.

11. Wurzinger, M., (2006, 2007). (1) Notes on the Iraq Study Group Report, http://www.otoom.net/notesontheisg.htm. (2) Notes on Where is Iraq heading? Lessons from Basra, ../notesonwhereisiraqheading.htm.

Martin Wurzinger Semi-retired, developed a mind model and a prototype of an artificial mind, continues research into the relevance and applicability of the model to wider society. www.otoom.net

Book review: Why Red Doesn't Sound Like a Bell

Kevin O'Regan presents a sensorimotor account of consciousness in a largely successful attempt to facilitate a discussion of a subject resistant to scientific analysis [1].

The theory concerns "raw feel", i.e., what is left of an experience after all measurable effects have been accounted for, or phenomenal experience. Three main applications of the theory are discussed: How raw feel is related to action; why the raw feels of different sensory modalities are distinct; and what is required for a being to consciously experience a raw feel. The term "sensorimotor dependency" refers to a relationship between a subject's motor actions and the resulting sensory stimuli from the environment. Holding an object, for example, has a feel related to the changes in the perceived stimuli that would occur with exploratory actions, rather than the stimuli that pertain to a static grasp at a given instance.

This approach has important implications for the requirements of sense organs. The eye, for example, can be considered an imperfect optical device due to features such as retinal scotoma and the blind spot. This critique is, however, only valid if its function is to gather a high fidelity snapshot view of the environment. The sensorimotor approach only concerns the *changes* in retinal stimulation as the eye moves relative to the environment, and these relationships remain in the presence of the features previously identified. There remains the question of how imperfect, inconsistently sensitive sense organs could produce a raw feel that presents itself as continuous and detailed. O'Regan declares this effect as tantamount to an illusion, features of raw feel can only be observed through active interrogation which necessarily presents them in detail. It is the convenience with which we can interrogate our whole visual field that produces the impression of continuous high fidelity. In the way that a fridge light is commonly only observed as being on, raw feel can only be consciously experienced in detail.

As the motor actions involved in interrogating visual signals are distinct from those involved in interrogating auditory signals, the differences between the sensorimotor dependencies affect their distinct feeling, that is, why the raw feel of the colour red is not like the raw feel of the sound of a bell.

By separating feel from stimulation at an instant and associating it with a continuum of potential interactions, the sensorimotor account provides explanations for the localisation of stimulus, environmental awareness without invoking a detailed internal representation, and how conscious experience presents itself as continuous.

Scientific accounts that describe consciousness as arising from activity amongst neural representations suffer from an explanatory gap: How can any neural structure generate raw feel? Why would it produce one type of feel rather than another? This question does not apply to the sensorimotor account as feel is not generated anywhere, it is simply "an abstract quality of our interaction with the environment" ([1], p. 182). O'Regan clarifies that they are not denying the necessity of the brain or the existence of representations, indeed neural encodings are required to store sensorimotor contingencies as learned, but simply that raw feel is not the product of the activation of an internal picture.

Though abstract, the quality of an experience that defines its feel can be characterised by a few terms, namely (partial) insubordinateness, richness, grabbiness and bodiliness. Respectively, they are the levels at which stimuli may change without a subject's acting, interrogating the experience presents detailed information, experiences grab our attention, and are subject to change as the body moves.

Feelings that rate highly on bodiliness and grabbiness are easier to describe as real sensations, while environmental sensory experiences are describable as also having high (partial) insubordinateness and richness.

A seemingly obvious counterexample is that of non-environmental experiences that require no motor actions such as remembering, imagining and dreaming. O'Regan addresses these phenomena specifically, explaining that purely introspective experiences have distinctly different profiles when considered in the given sensory terms, for example there is no insubordinateness in a recalled situation, entirely of your creation. This point is paramount as it demonstrates that O'Regan's theory can not only admit purely mental experiences but also explain the differences in their experiential qualities.

To be able to experience raw feel, it is described as necessary and sufficient for a being to have the following cognitive capacities. First, the capacity to be poised to make use of a set of sensorimotor contingencies. Second, the capacity to be poised to make use of the fact that they are poised to make use of a set of sensorimotor contingencies. Third, the being requires a notion of self.

In the case of a being with the previously described cognitive capabilities one final requirement remains for the quality of an interaction to be consciously experienced, that is the being must also be consciously attending to that quality.

This stance has two important implications for what is *not* present in raw feel; without conscious attendance raw feels are not felt, and without a notional sense of self experience *cannot arise at all*. Babies and animals, therefore only experience pain if they are considered to have developed a specific cognitive capability. This is a potentially controversial result, but it is not damaging to the theory in the way an inconsistency would be, and there is certain credit in the theory making clear, if controversial, predictions even when applied to fringe cases. O'Regan clarifies that stimuli can be processed and may affect future behaviour without being consciously attended to, though they will not be present in conscious experience. This is another potentially controversial result as raw feels that are not felt are present in any interaction in an environment, applying equally to animal, vegetable and mineral.

Though it may be impossible to empirically prove any account of conscious experience, O'Regan demonstrates the advantages of adopting the sensorimotor account. It may seem a radical departure from classical accounts of consciousness but it is not incompatible with the majority of the literature, and provides clear explanations for the observations therein.

The text falls short of an set of imperative instructions for building a robot that is conscious and feels, but O'Regan explicitly states that he sees no logical reason why it cannot happen.

Ultimately, the sensorimotor account of consciousness is a significant theory in the philosophical reshuffle pervading contemporary Cognitive Science. This makes this concise, accessible text broadly relevant and potentially very influential to this audience.

Reference

1. J.K. O'Regan (2011). Why Red Doesn't Sound Like a Bell: Understanding the Feel of Consciousness. Oxford University Press, 182 pages.

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

Event: Foundations of Enactive Cognitive Science, Feb 26–27th, 2012

This workshop was organised by Etienne Roesch, Slawomir Nasuto and J. Mark Bishop, and sponsored by the Centre for Integrative Neuroscience and Neurodynamics and the School of Systems Engineering (Univ. Reading), with the support of the EPSRC, SAISB, the AVANT journal.

The meeting provided a truly interdisciplinary forum that focused on the challenges and future directions of research for cognitive science. Keynotes from Mark Bickhard, Fred Cummins, Tom Fröse, Thomas Fuchs and Kevin O'Regan, together with the diverse contributions from the oral presentations and posters, successfully illustrated how knowledge from the domains of philosophy, theoretical biology, linguistics, robotics, neuroscience and pathologies are making a significant contribution to the paradigm shift within cognitive science—from the conventional to enactive approaches to interaction. The meeting revisited Maturana and Varela's biological systems framework that places cognition as knowledge, and knowledge as action at the centre of how organisms bring forth their interaction with world. Such approaches reconsiders the dependence on representation for cognitive processes. These include embodiments in AI platforms and applications in applied research arenas such as, developmental psychology, psychiatry and the potential for related therapies.

Discussions emphasised that the theoretical foundations of such disciplines need to revisited, together with consequences for conventional computational models. Further, it became apparent that the field can no longer justify fragmented approaches that ignore the embodied, extended and enactive nature of the interaction of organisms with their environments. What the meeting exemplify was a viable community of researchers that has the capacity to not only to contribute to the paradigm shift but also to embrace and nurture a vision for the future of enactive cognitive science.

Marilyn Panayi, PhD Candidate
School of Health Sciences, City University, London, and Ensomatica, London, UK

Event: Intuition and Ingenuity – Alan Turing Centenary Exhibition

The AISB/IACAP World Congress in 2012 played host to a stunning touring art exhibition organised by the Alan Turing Centenary Arts and Culture Subcommittee. The show aimed to celebrate Turing's life and work by looking at its impact on the art world. The exhibition was an official Turing Centenary project curated by Anna Dumitriu, Sue Gollifer and Nick Lambert. Arts Council England, The Computer Arts Society and The University of Hertfordshire kindly supported the exhibition.

This exhibition, which took its name from Turing's own writing on the subject of mathematical reasoning, brought together a number of important artists from digital art pioneers to emerging contemporaries to investigate Turing's enduring influence on art and contemporary culture.

A centrepiece of the show was a 'face stealing' robot created by artists Anna Dumitriu and Alex May in collaboration with Dr Michael Walters and Professor Kerstin Dautenhahn from the University of Hertfordshire. The piece uses a Microsoft Kinect sensor to take features from exhibition visitors' faces and combine them with features from (up to 16 of) their companions' faces based on their proximity to the robot. As one approaches, it welcomes you, turns to you, waves its arms and begins to morph. The robot then speaks phrases such as "I like your face" or "I love you". The artwork intentionally plays with that feeling of discomfort, well known in robotics as "the uncanny valley" (Mori, 1970), where users feel a sense of repulsion as robots become very human-like (in this case very like themselves and their companions) but stopping short of being wholly human. The depth camera in the Kinect can be used to measure this effect in operation by recording how visitors approach the robot and produced huge amounts of data that were used in a scientific study by the University of Hertfordshire (presented at TAROS 2012). "My Robot Companion" was given an AISB Public Understanding of AI Award in 2011 and has attracted over 60000 visitors to exhibitions at many venues including The Science Gallery in Dublin, The Science Museum in London, Kinetica London and Lighthouse Gallery in Brighton.

The exhibition also included works by such digital pioneers as Roman Verostko, now 83, who has been creating algorithmic art since the 1950s; Professor Ernest Edmonds, who spoke about how Turing had inspired him as part of the AISB/IACAP World Congress; Professor Paul Brown, who created a new cellular automata-based 'kinetic painting' for the show called "Dragon"; boredomresearch, whose evolving artwork "Fragments of Lost Flight" used Turing Machines in its creation; Martin A Smith's newly commissioned soundwork; historical works by William Latham including one created during his pioneering residency at IBM in the late '80s; a reworking of Professor Greg Garvey's famous "Automatic Confession Machine: A Catholic Turing Test"; and a genuine "Smoke and Mirrors Machine" by Alex May in collaboration with Professor Bruce Christianson from University of Hertfordshire. For more information on the show and other touring venues see www.turingcentenaryarts.eu

References

1. Mori, M. (1970). Bukimi no tani – The uncanny valley. (K. F. MacDorman & T. Minato, Trans.). *Energy*, *7(4)*, 33–35. (Originally in Japanese)

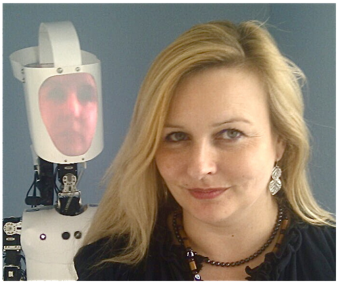

Anna Dumitriu, Artist and Curator Co-chair of the Turing Centenary Arts and Culture Subcommitee; Visiting Research Fellow and Artist in Residence, Department of Computer Science at The University of Hertfordshire; Artist in Residence on The Modernising Medical Microbiology Project based at The University of Oxford; Visiting Research Fellow and Artist in Residence, Centre for Computational Neuroscience and Robotics at The University of Sussex.

Event: First AISB Members' Workshop – Sensorimotor Theory

The first AISB Members' workshop was held at Goldsmiths, University of London on the 26th September, co-organised by Mark Bishop and Andrew Martin. Addressing Sensorimotor Theory and described as "A day of discussion on the Sensorimotor account of Perception, Consciousness and Robotics, its development and contemporary state", and with over 60 delegates attending the event attracted interest across a wide range of disciplines.

Two themes emerged from the presentations, they were: the results of adopting a sensorimotor approach in various ongoing research projects and the comparisons of the philosophical profile of sensorimotor theory with that of (relatively traditional) cognitivist and (relatively dynamic) enactive theories.

J. Kevin O'Regan gave the inaugural keynote presentation, describing the latest developments in sensorimotor theory since his seminal 2001 paper "A sensorimotor account of vision and visual consciousness" coauthored with Alva Noë. O'Regan's focussed on an account of "raw feel" which provides an account of how the phenomenological profile of conscious experiences differ and the means for their scientific investigation. In extending the theory to the necessary and sufficient criteria for conscious experience, O'Regan includes concepts of cognitive access to one's activity and a notion of self. As presented by O'Regan, contemporary robotics projects can be described as having the necessary cognitive access resulting in lively discussion on the possibility of computational consciousness. Similarly, the requirement for a notion of self as described in O'Regan's (fundamentally Dennettian) terms which implied that (for example) babies did not *consciously experience* pain, provoked further discussion amongst the attendees.

The keynote presentation for the afternoon session, given by Daniel Hutto praised the move away from cognitivism evident in sensorimotor theory but identified philosophical criticisms that would apply until the theory cut all ties with its classical roots. The implications of the argument being that, in its current form sensorimotor theory will prove valuable in many applications, but in reforming relatively auxiliary aspects of the theory it could prove to have general explanatory power and applicability to consciousness.

Outside of the keynote presentations were energetic sessions of presentations describing projects applying a sensorimotor approach. Interesting parallels were drawn between sensorimotor perception and the technique from the field of robotics known as active sensing; by mimicking the morphology and action of whiskered animals Sheffield University's Active Touch Laboratory demonstrated a novel and successful technique for robotic perception. Chrystopher Nehaniv demonstrated the cutting edge of enactive robotics which showed how interactions between humans and socialised robots can result in emergent behaviours in both participants.

There was a general acceptance of the fundamental tenet that perception is grounded in embodied interaction in an environment. Of the criticisms, they only addressed sensorimotor accounts of the constituents of conscious experience, and were always accompanied by alternative formations of the theory that allowed these criticisms to be avoided. Ultimately the workshop showed there was a lot of interest in sensorimotor approaches to various aspects of cognition and a wide range of active research benefitting from the approach.

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

CfP: Hyper-Heuristics, Past, Present and Future

Important Dates

- Abstract Subm.: 14 Nov 2012
- Notifications: 16 Nov 2012
- Registration: 21 Nov 2012
- Workshop: 13 Dec 2012

About the Workshop

The workshop aims to provide a forum for discussing all aspects of hyper-heuristic research from philosophical and conceptual foundations through to new approaches and applications of methods from the field. In addition to presentations on traditional hyper-heuristic methodologies, presentations are particularly invited which focus in new hyper-heuristic paradigms and classes of optimisation problems, such as Multi-method hyper-heuristics and approaches to solving multi-objective optimisation problems. Please note, a new Hyper-heuristic symposium will be launched at the AISB's 2013 annual convention.

Date & Venue

The hyper-heuristic workshop will be held at the University of Exeter, Streatham Campus, Devon on Thursday 13th December 2012. The event will take place between 10am and 4pm. Details of the workshop location and a full timetable and list of speakers will be publicised closer to the event.

Registration

First come, first serve basis. Attendees should register for the workshop by emailing workshops@aisb.org.uk with their full name, affiliation, position, email address and AISB membership number (if applicable). Non-members are required to pay a fee equal to the annual AISB membership rate and will be eligible for membership with no additional fee and will be processed as a membership application unless explicitly requested. AISB membership costs are available at http://www.aisb.org.uk/about/join.

Submissions

500 to 2,000 words abstracts are requested for speakers wishing to present at the workshop. Abstracts will be reviewed by the workshop organisers and assessed based on quality and relevance to the workshop. The presentation title and abstract should be submitted to the workshop organisers at workshops@aisb.org.uk by Wednesday 14th November 2012. The workshop abstracts will be published online through the AISB website and the University of Exeter.

Dear Father Hacker...

Dear Aloysius,
I was just about to start the final year of my BSc in Artificial Intelligence at the University of Poppleton, when I got caught up in the latest immigration scandal. Thanks to the UK Border Agency, the University can no longer authorise visas, so, as an overseas student, mine has been cancelled. A month ago I was looking forward to a first class degree and a glittering career; now I'm facing deportation. What can I do?

Yours, Visaless

Dear Visaless,
In order to resume your studies, you need to become a UK citizen as soon as possible. The standard naturalisation process is, regrettably, too slow and haphazard for your purposes. Fortunately, my Institute has developed an alternative, faster and totally assured process: SANCTUARY™ (*Student's Alternative National Citizenship Technique Usurps Agency's Records for You*). For a modest consideration, SANCTUARY™ will 'correct' Government data so that you will always have been a UK citizen.

The Government secretly welcome SANCTUARY™. The rescinding of your University's visa authority is, of course, linked to the Government's drive to reduce immigration. Every immigrant who becomes a citizen thereby reduces the number of immigrants in the UK. SANCTUARY™ speeds up this reduction process by several orders of magnitude.

Yours, Aloysius

Dear Aloysius,
Thanks to your advice in your 5th Agony Uncle column, we entered our running robot, R...RUSH™ (*Robot Runner Undertakes Sprints and Hurdles*) for the 100 metres race in the 2012 London Paralympics. The whole world now knows the outcome. R...RUSH™ won his heat with a new world record, but was disqualified because his blades were considered too long. Undeterred, we are now planning to enter an improved version, R...RUSH 2™, to Rio.

Yours, Coach

Dear Coach,
We are always delighted to learn of the successful application of Institute technology. Our PISS™ (*Pass Inspection via a Source of Solutions*) attachment enabled R...RUSH™ to enter the race, even if it was subsequently disqualified on unrelated grounds. We had several other very satisfied customers in both the Olympics and Paralympics, although customer confidentiality prevents me from revealing the full details.

For Rio we are developing a range of new technologies, including our intelligent running blades, SPUR™ (*Sporting Projection Uses Restyling*). Their length will dynamically but imperceptibly vary from a short form in order to pass the pre- and post-race inspections to a longer form to ensure victory in the race itself.

Yours, Aloysius

Dear Aloysius,
My opponent in this November's presidential contest has failed to implement the changes he promised and has presided over an unpopular war and a failed economy. I ought to be streets ahead of him in the polls, but instead we're neck and neck. Could this be due to his massive support from social media? Can you help me organise a comparable army of cyber-world friends and followers?

Yours, Friendless

Dear Friendless,
You are right that social media support is a vital ingredient in any successful 21st century, political campaign. I'm delighted to report that the Institute has the ideal product to help you. ACOLYTES™ (*Automated Chatbots Organise to Laud Your Talents and Ensure Success*) orchestrates a huge multi-agent army, on Facebook, Twitter and other social media sites, to sing your praises. The achievements of ACOLYTES™ have been demonstrated in a series of recent amazing political triumphs.

Yours, Aloysius

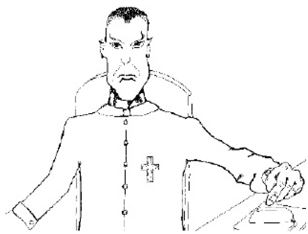

Fr. Aloysius Hacker
Cognitive Divinity Programme
Institute of Applied Epistemology

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.

Back matter

Advertising and Administration
Dr Katerina Koutsantoni (AISB Executive Office)
Institute of Psychiatry, King's College London
Addictions Sciences Building (B3.06)
4 Windsor Walk, Denmark Hill
SE5 8AF, London, United Kingdom
Tel: +44 (0)20 7848 0191, Fax: +44 (0)20 7848 0126

AISB Quarterly – No. 135, October, 2012

The AISB Quarterly is published by the Society for the Study of Artificial Intelligence and Simulation of Behaviour (AISB). 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 email bulletin and web search engine for AI-related events and opportunities

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

ISSN 0268-4179
