AISB Quarterly

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

No. 139 June, 2014

Do you feel artistic?
Exhibit your artwork on our front covers!
Email us at aisbq14@aisb.org.uk!

Alan Turing (23 June 1912 – 7 June 1954) was a British mathematician, wartime cryptoanalyst, computer scientist and philosopher. His formalisation of the concepts of *algorithm* and *computation*, with the so-called Turing machine, has had profound influence on each of these fields, quite literally laying the foundations for artificial intelligence as we know it.

Turing died exactly 10 years before the foundation of AISB. His professional career left him little time for socialising, and we will never know whether he would have joined our Society. His interests, convictions and aspirations, however, leave little doubt that our Society would have valued his input.

The 2014 Convention, recently held at Goldsmiths, University of London, celebrated his life and work. The conference proceedings can be found at

http://aisb50.org/proceedings

Editorial

This issue of the Q is a collector. It is a milestone, in many ways.

∴

The 2014 Convention, held at Goldsmiths University London, marked the 50th anniversary of the founding of the AISB, as well as 60 years since the death of Alan Turing. The articles that we have put together commemorate this piece of History.

In the first article, Eduardo Alonso, who served as Vice-Chair of the AISB between 2003 and 2006, sheds some light on the origins of our Society, which is in fact the *eldest AI society in the world!* He takes us through the motivations and aspirations of our Founding Fathers.

In the second article of this issue, Mark Bishop takes stock of the field of AI today, and reviews this year's Convention in the context of what he calls a "pivotal branch point". Mark, who was the lead organiser of this year's Convention, recently stepped down as Chair of the Society. In this article, he remarks that the Convention attracted very diverse interests, rooted in "embodied, embedded, enactive and ecological cognition", and asks whether this theoretical turn will carry the field in Future.

As if echoing Eduardo's historical perspective and Mark's reflection, that the remit of AI is broadening, the reports that conclude Q139 include a review of Hutto & Myin's book "Radicalizing Wnactivism" by Paulo de Jesus, an account on the last AISB workshop on the emergence of consciousness, by Janet Gibbs, and a report on the 14th International Conference on Computer, by Siyamalan Manivannan.

∴

I would like to end this editorial with an appeal for help to our members. In his historical account, Eduardo reminds us that one of the main *raison d'être* of the Society has been to bring like-minded people together, with the organisation of frequent workshops, and through the Quarterly. Ironically, 50 years ago, one of the struggle was already to find content for these outreach activities.

There are many ways to participate in the Society, and many reasons to do so. You can organise a workshop on your favourite topic, with financial and logistical support from us; you can increase the visibility of your work by contributing an article to the Q; you can join our committees and make a difference, .. and you can even join the editorial board of the Q!

If any of the above strikes a chord, and you want to be involved, or if you have questions, feel free to contact us at aisbq14@aisb.org.uk!

Etienne B. Roesch

Happy birthday, AISB!

by *Eduardo Alonso* (City Univ. London, United Kingdom)

If you are reading this issue, it is probably because you are a member of the AISB. So, let me ask you, what do you know about the AISB? You may know that we are celebrating the society's 50th anniversary and hence that it was founded in 1964. You may also know that it is a member of the European Coordinating Committee for Artificial Intelligence (ECCAI), that it publishes the AISBQ ("the Q"), and that it organises an annual convention. You may also know that one Eduardo Alonso, author of these lines, who, incidentally, is by no means a professional historian, acted as vice-chair of the AISB between 2003 and 2006. Well, maybe not.

In this short paper, I am giving a quick account of the first years of the AISB, that is, of the period during which, paradoxically, the AISB was not the AISB strictly speaking.

∴

The origins

On October the 26th 1964, after a one-day "Symposium on Artificial Intelligence and Simulation of Behaviour" Max Clowes writes to the British Computer Society (B.C.S.) to form a Study Group, with two main objectives, namely, to arrange meetings, and to edit and circulate a Newsletter. It receives a favourable response from the B.C.S. the 20th of November. The symposium itself was held in September 1964 at the Northampton College of Advanced Technology (C.A.T.), later to become City University, at Northampton Square, London, and organized, most likely, by Robin Milner. This came as a surprise to the author of this short paper since he has been working at City University, now City University London, since 2001.

The AISB is established as the "British Computer Society Study Group on Artificial Intelligence and Simulation of Behaviour". The AISB as we know it, a learned society independent of the B.C.S., will be established in 1974 (see later).

In its foundational document, an AISB steering committee is appointed, consisting of Donald Broadbent, I.J. Good, Donald Michie, Christopher Strachey and Max Clowes, secretary and "dogsbody". Max Clowes will later serve as first Chair 1969–1972, followed by Bernard Meltzer 1972–1976.

To give the reader some perspective, the A.C.M. SIGART was founded in 1966, and the IJCAI corporation in 1969; the AAAI dates from 1979, and ECCAI from 1982. Thus, the AISB is, in fact, the *eldest AI society in the world*.

Within the U.K., Michie's Experimental Programming Unit dates from 1965, before becoming the Department of Machine Intelligence and Perception in 1966 (joined by H. Christopher Longuet-Higgins' Theoretical Section and Richard Gregory's Bionics Research Laboratory); the first "Machine Intelligence Workshop" was organized in September 1965; and in the same

year, a one year diploma course in Machine Intelligence Studies (with 5 students) was run for the first time. Research-wise, this is the time of MENACE, Graph Traverser, POP-1, and FREDDY I.

In a more mundane tone, we were still using pre-decimal currency, shillings and denarii. According to the first balance of the "A is B", as the society was mockingly coined by Max Clowes, the original contribution of the B.C.S. was £10 and the subscription 10/–.

The relationship between the AISB and the B.C.S. was an uneasy one from the very beginning. In a letter from Rod Burstall to C. Strachey dated November 1966, on the occasion of organising the 1967 AISB meeting at the Atlas Computer Laboratory, Chilton, we read that "the B.C.S. regard symposia as fund-raising affairs (...). They have been charging admission fees of 5-7 guineas (...)"; Strachey's reply reflects clearly the society's dissatisfaction with this state of affairs: "I find the B.C.S. attitude quite deplorable. Their function should surely be to encourage the development of specialist groups, not to try to make money from them. (...) I should be in favour of severing all connection with them. (...)".

From the first *Membership List* (60 members, May 1st 1965) we learn that AISB research was not exclusively focused on Edinburgh, as one is sometimes led to believe; and that there was a genuine interest in bringing together researchers from different academic areas but also from industry (Elliott Brothers, IBM), the Government (the Ministry of Aviation, the Middle East Command!), the Bank of England, the BBC... and special mention goes to psychology: It was not a coincidence that two out of the five members of the original AISB steering committee, D. Broadbent and M. Clowes, were psychologists; that members of the Experimental Psychology Society "should have free access to our meetings and to the newsletter"; and that the Medical Research Council were actively involved in the society's activities during the 60s. Somehow, this changed over the years, with the AISB becoming more AI and less SB.

The Quarterly

It is precisely the Quarterly that better defines the AISB, and the best source of information about these first days of our society. So, it is worth examining its trajectory: The Quarterly starts as the "AISB Newsletter" edited by, who else, Max Clowes; from Issue 3, April 1966, Rod Burstall and Jim Doran take over; and under Pat Hayes' editorship it becomes "the European AISB Newsletter" in July 1969; then it is briefly transformed into the "Bulletin" between November 1972 and February 1973, with E.W. Elcock and A. Ortony. Their editorial strategy was not welcome. In a letter to the Cttee., Feb 22 1973, they propose a "good solution" to the editors' difficulties, "to amalgamate with its sister–SIGART Newsletter". Clowes' answer settles the issue "(...) the Elcock-Ortony proposal is appalling. With less than 12 months of taking on the (admittedly difficult) task of editing the Newsletter they want to throw it into SIGARTs lap. I protest strongly." As a consequence, we are back to the European Newsletter, Is-

sue 14, July 1973, this time with Alan Bundy and M. Liardet as editors; and, eventually, the Quarterly sees the light of day, in October 1977 Issue 28, with Tim O'Shea (and a team of sub-editors including B. Welham, R. Young and G. Plotkin, and later C. Mellish and L. Daniel). At some point, there was a debate about handing production and distribution of the Quarterly to professional publishers (North-Holland) in 1978-79, that is, about becoming an appendix to the AI Journal, but it did not prosper.

Speaking of which, it is good to remember that the AISB was instrumental in launching "Artificial Intelligence", published by Elsevier and sponsored by A.C.M. SIGART. It is not an accident that B. Meltzer was proposed as the first Editor-in-Chief. Or, for that matter, that Tony Cohn, AISB chair between 1992 and 1994, is the current one (along with R. Dechter). It is also worth mentioning that, between 2001 and 2005, the AISB published six numbers of the "AISB Journal" (edited by Geraint Wiggins), and one more in 2010.

The Convention

As for meetings, there were plenty: one-day scientific meetings, typically invited speakers from "Machine Intelligence" workshops, one-day specialized meetings (in chess, theorem proving, robotics, ...), and summer schools, typically organised in London (I.C.L., Q.M.C., C.A.T.), or around (Oxford, the Atlas Computer Laboratory, Sessex, ...). And, of course, the first conference, organized in Brighton 1974 by Keith Oatley and Margaret Boden. The first AISB conferences, Edinburgh 1976, Hamburg 1978, Amsterdam 1980, were in fact European. The first ECAI, Orsay 1982, preceded the split of the society (see later), and of the conferences, with the AISB conference and ECAI alternating years from Exeter 1983 and Pisa 1984, except in 1996 (Sussex and Budapest respectively); a joint conference in Brighton 1998 was followed by the final separation between the now annual AISB convention and the biannual ECAI.

The relation between the AISB and Europe can be traced back to 1969: Newsletter's Issue 8 informs that during the first IJCAI "a special meeting for the European delegates" was held resulting in Erik Sandewall reporting that the "British" AISB Newsletter becomes the "European Newsletter for Artificial Intelligence and Simulation of Behaviour", to be produced by Uppsala University, distributed from Edinburgh, and edited by Pat Hayes: "It seems likely that for a while at least we will function as a kind of European AisB, until other national groups are formed on the continent"; ... until means 1982, when the ECCAI is founded with the AISB as a member.

AISB three main characteristics

Browsing through the first newsletter/quarterly issues one finds three recurrent themes that, to a certain extent, define the AISB even today, namely,

(a) (lack of) money: ToC Issue 9, November 1969, puts it rather explicitly: "MONEY *** IMPORTANT ***";

as it does J. Doran's letter to B. Meltzer, 18 Jan 73: "There is about -£5 in the kitty now";

(b) believe it or not, (lack of) contributions to the Quarterly: 22nd April 69 letter from Pat Hayes to AISB members reads "(...) only one contribution had been sent to the editor (...)";

(c) tongue-in-cheek attitude: One of my favourite examples is M. Liardet and A. Bundy's Report on the AISB Scientific Meeting January 5th 1974, "(...) was badly hit by the power crisis and the railway worker's dispute (...) but the numbers were boosted to about 30 by Edinburgh workers. The meeting took place in a cold lecture theatre on a cold wet Saturday. Hopes of some relief over lunch were dashed when we arrived in an even colder refectory to face an (airline type) salad. (...)", closely followed by a hilarious advertisement:

> *"ATTRACTIVE SCOTTISH ACADEMIC enjoys affluent lifestyle - fast cars, expensive holidays, excellent software environments, beautiful Georgian flat, etc. Highly successful, tenured, respectable publications (AI Journal, Cognitive Science, etc.), large grant-holder. Into Szechuan cooking, Baroque music, Linton Kwesi Johnson, jogging, backpacking, and wholefoods. Seeks sincere, object-oriented woman for discreet, loving relationship and mutual simulation. Send 1024x1024 pixel image in RT-11 format. Box-MC68000."*
>
> The AISB Quarterly
> Issue 39, Dec 1980.

These pranks were accompanied by sections like Brady's cryptic crossword, Aguirre's Wyno the Learning Computer cartoons, limericks, a section of Silly Acronyms (e.g., PROLOG: PRObably the Language Of God), and, of course, Father (Aloysius) Hacker, whose identity is one of the best kept AI secrets since Bletchley Park–all the author can say is that, by July 1979, the list included Benedict du Boulay, Alan Bundy, Chris Miller, Hal Abelson, Gordon Plotkin, and Tim O'Shea. So it goes.

The AI Winter and the AISB

Back to historical facts, from early 1972 Michie promotes the idea of a learned society, and in a letter dated 26th January 1973 M. Clowes confides his fears to B. Meltzer "I feel that the long term view of the AISB is hanging by a thread (...)". It must have been clear by then that the Lighthill Report was not going to be complementary. And indeed it wasn't. Lighthill's "Artificial Intelligence: A General Survey" was followed by the famous BBC debate "The General Purpose Robot is a Mirage" and many comments–by N.S. Sutherland, H.C. Longuet-Higgins, D. Michie, R. M. Needham ... and John McCarthy. Yet, the best two appeared in the AISB European Newsletter July 1973 Issue 14: "Serendipity Resources Council, The Darkvale Report on Applied Mathematics, A Cardboard Conference" by Sir Gorgam Darkvale, F.R.S. Caucasian Professor of Divinity, University of Grantabury, and Pat Hayes' "Some Comments on Sir

James Lighthill's Report on Artificial Intelligence". Their reading should be mandatory in any AI course.

As an aside, R. Needham's role in this affair is perplexing: In October 1982, he and P. Swinnerton-Dyer enthusiastically supported the S.E.R.C. Alvey Programme in a pamphlet, Artificial Intelligence Research in the UK, which argues that Lighthill (that is, the S.R.C.) was right then, in 1973, but that he would have been wrong in 1982. A new brand is all it was needed: Intelligent Knowledge Based Systems, IKBS, is born.

Nevertheless, despite accepted wisdom, the AI winter seems to have been rather warm: In 1972, the S.R.C. Long Range Computing Policy Panel recommended "the immediate creation of at least one more major centre of Machine Intelligence" (S.R.C. Computing Science Review, p. 25); Meltzer's "S.R.C. policy with respect to senior appointments on research grants in universities" April 4th 1973, was adopted by the council; S.R.C. set up a standing Artificial Intelligence Panel in 1974; in 1974/75 the first AI courses are launched in Edinburgh and Sussex, as well as a cognitive studies programme in Sussex; the S.R.C. promotes the creation of a computing network between Edinburgh, Manchester, A.C.L. and the S.R.C. Rutherford Laboratory in July 1976; S.R.C. Interactive Computing Facilities Committee sets up a special interest group to provide advice on software requirements for AI in July 1977 (with A. Sloman (Chair), R. Burstall, A. Bundy, M. Brady, A. Smith and P. Kent); S.R.C. grants are still awarded to Edinburgh staff, including Michie ... Perhaps the best example that the crisis was not as severe as first feared is the fact that the AISB survives: It becomes a learned society, organizes its first conference, there is a steady growth in membership numbers, the newsletter is published regularly...

The AISB

Minutes of the Cttee meeting 11 April 1973 report that: "Donald Michie expressed that, in his view, the essentials of such a society were that: –

i. It possessed a formal constitution,

ii. It was therefore subject to democratic safeguards eg rotation of committee membership,

iii. Membership was not automatic - a minimal requirement was that a candidate must be proposed by an existing member,

iv. It existed as a legal entity, not merely as a group of individuals.

Donald Michie argued that if the Study Group were reformed in this way it would be taken more seriously by the outside world and by its own members." ACTION: Meltzer and Michie to prepare a draft constitution.

Events followed: the Constitution is proposed in Issue 17, July 1974; approved (with only one nay) in October 1974 (then revised in 1977 for charitable status); ballot for election of the Committee, Issue 18.

Europe and all that

Shortly after becoming the AISB, a problem arises: "the issue is whether we work towards separate national

"AISB"s (...) or a single continental organisation. Current political theory calls for the latter, sited in Brussels, with a bureaucracy of around one thousand!", from M. Brady's chairman message Nov 76. Jörg Siekmann (then in Essex) had reported on "German Intelligence Becomes Artificial" in November 74, and by 1977 the German AI Study Group had grown strong, with their own Newsletter and increasing governmental funding. In July 1977 there are mixed signals: "The European AISB Newsletter" becomes the "AISB Quarterly" ...and announces that the next AISB Summer Conference will be held in Germany.

Then, Wolfgang Bibel from IFI Munich circulates a proposal for the creation of the EAAI, Oct 79, arguing that "most people (...) regard it (the AISB) as a British society", and proposing a board with two representatives from Britain and one from each of France, Germany, Italy, Benelux, Scandinavian and Eastern Europe.

Pat Hayes, among others, is against it: "I can't support this idea. I strongly believe (...) would be a disastrous mistake. I also think that your (Bibel's) motivation in wishing to set it up reflects priorities and views on science and nationalities which are wrong." (...) "There is no room for a second organization, the entire European AI production is smaller than Stanford's" (...) "EAAI is going to directly compete with AISB". Facts though are stubborn things:

- AI groups in the continent had proliferated in the 70s;
- Although there is no mention of Britishness in the AISB Constitution, there is little involvement of continentals in the Cttee: E. Sandewall serves from 1969 to 1975, then he is not nominated in first ballot (see above);
- It is acknowledged that currency is an issue for "foreigners in joining AISB" (Sept 78);
- The Hamburg conference is a huge success, boosting confidence of continental members.

And "Europe" is not the only problem: "AISB disciplines" start to form their own societies, e.g., the British Robot Association (by Larcombe, 1977); in the late 70s Donald Michie himself founds the B.C.S. Special Group on Expert Systems -to become BCS-SGAI in June 1980; psychologists feel alienated and join the new Cognitive Science Society in 1979.

It was a brave new world.

∴

Acknowledgments: I would like to express my gratitude to the AISB, to John Barnden, Mark Bishop and Katerina Koutsantoni in particular; very special thanks to Jim Doran, Alan Bundy, Richard Young, Tony Cohn, Pat Hayes, Wolfgang Bibel, Jim Howe and Tim O'Shea; also to Blay Whitby (and his flying session), Erik Sandewall, Aaron Sloman, Maggie Boden, Geraint Wiggins, Stephen Muggleton and Ben Du Boulay. I would also like to thank the Royal Society for funding my research on the history of the

AISB (RG2010/R2), the Bodleian Library, and Esther Mondragón for her help in cataloguing the material that supports this paper.

Eduardo Alonso
Department of Computer Science
City Univ. London, United Kingdom

AISB50: Artificial Intelligence at a new branch point

by *J. Mark Bishop* (Goldsmiths, Univ. London, UK)

Introduction

> *It should be noted that from now on 'the system' means not the nervous system but the whole complex of the organism and the environment. Thus, if it should be shown that 'the system' has some property, it must not be assumed that this property is attributed to the nervous system: it belongs to the whole; and detailed examination may be necessary to ascertain the contributions of the separate parts.*
>
> W. Ross Ashby, 1952 [1]

An oft repeated aphorism is that the universe is constantly changing and hence that our world is in a perpetual state of flux. In order to behave intelligently within this varying natural environment, any system - be it man, machine or animal - faces the problem of perceiving invariant aspects of a world in which no two situations are ever exactly the same. Cartesian theories of perception can be broken down into what Chalmers [5] calls the 'easy problem' of perception - the classification and identification of sense stimuli - and a corresponding 'hard problem' - the realisation of the associated phenomenal state[1]. The difference between the 'easy' and the 'hard' problems - and an apparent lack of link between theories of the former and an account of the latter - has been termed the 'explanatory gap' [10] and this [unbridgeable] gap is symptomatic of the underlying dualism.

Many current theories of natural visual processes are grounded upon the idea that when we perceive, sense data is processed by the brain to form an internal representation of the world. The act of perception thus involves the activation of an appropriate representation. The easy problem reduces to forming a correct internal representation of the world and the hard problem reduces to answering how the activation of a representation gives rise to a sensory experience.

In machine perception progress in solving even the 'easy' problem has so far been unexpectedly slow; typical bottom-up (or data driven) methodologies involve the processing of raw sense data to extract a set of features; the binding of these features into groups then classifying each group by reference to a putative set of models. Conversely, in top down methods, a typical set of hypotheses of likely perceptions is gen-

[1] David Chalmers introduced the term 'hard problem' to investigate "Why is all this [neural] processing accompanied by an inner life?" [5]; I deploy the phrase 'hard problems of consciousness' to additionally encompass related problems pertaining to Levine's 'explanatory gap' [10].

erated; these are then compared to a set of features in a search for evidence to support each hypothesis.

To date successes in machine perception have been limited to a relatively small subset of the possible [human] perceptual gamut[2].. Hence, at this 50th Anniversary Convention of the AISB, AI once again finds itself, pace Dreyfus [8], at something of a branch point; a choice, not as Dreyfus imagined in 1988, between 'making a mind' or 'modelling the brain', but a choice at a much more fundamental level between neo-classical paths (e.g. GOFAI; connectionism; dynamic theories of mind; swarm intelligence etc) that are fundamentally (i) dualist and (ii) essentially formal and representational[3], and the more radical 'Embodied, Embedded, Ecological, Enactivist' - the so called '4Es' - framework.

In this context, it was a particular delight that at this golden anniversary convention of the AISB both paths towards AI were well represented: the 'classical' approaches being championed by symposia such as: *Computational Creativity; Computational Intelligence; Computational Scientific Discovery; and Evolutionary Computing*; and the 4Es being championed by symposia such as: *Varieties of Enactivism; Consciousness without inner models; Reconceptualising mental illness, and Embodied Cognition, Acting and Performance*[4]. The foundations of classical AI are by now very well known; the foundations of the 4Es approach perhaps a little less so, therefore a few words of contextualisation may be helpful.

The 4Es: 'Embodied, Embedded, Enactive and Ecological' Cognition

Rooted at the heart of the emerging 4Es framework is [radical] *enactivism*; a theoretical approach to understanding the mind derived from work by Humberto Maturana, Francisco Varela, Evan Thompson, and Eleanor Rosch. In contrast to classical computational (cognitivist) or bottom-up (cybernetic) approaches to machine intelligence, the emerging enactivist framework emphasises the way that mentality emerges through self-organising processes *interacting* with their environment.

At a fundamental level enactivism is non-dualist: the self arises as part of the process of an embodied entity interacting with its *umwelt* in precise ways determined by its physiology. In this sense, individuals can be seen to "grow into", or arise from, their interactions with the world. The self does not repre-

[2]As is no surprise to long term critics of 'classical' AI such as Searle [17], Dreyfus [8] and Penrose [15]; for extended critical reflection on the impact of Searle's 'Chinese Room Argument' see, for example, [16].

[3]In describing a symbol, representation or process as purely 'formal', I mean both that it has no intrinsic content (semantics/meaning), and that it is 'hardware independent'.

[4]E.g. The *Embodied Cognition, Acting and Performance' symposium* featured in a post-conference 'New Scientist' article on the use of principles from embodied cognitive science in the treatment of autism, *Drama helps kids with autism communicate better* [New Scientist: April 17th 2014].

[5]Francisco Varela defined the term, "to evoke the view that what is known is brought

sent the world, but *produces it* through the nature of its unique way of interacting with its environment[5].

One particular 'variety of enactivism' is *sensorimotor theory* and the conference was also fortunate to host one of its founders - Kevin O'Regan - leading both a symposium on *Consciousness without inner models* and presenting at several symposia. Contemporary Sensorimotor Theory [4] offers a new *enactive approach*[6] to perception that emphasises the role of motor actions and their effect on sensory stimuli. The seminal publication that launched sensorimotor theory is the target paper co-authored by J. Kevin O'Regan and Alva Noë and published in Behavioral and Brain Sciences (BBS) for open peer commentary in 2001 [13].

In the central argument of their paper, O'Regan and Noë suggest radically shifting the nexus of research in visual perception away from analysis of the raw visual patterns of stimulation, to refocus on the law-like changes in visual stimulation brought about as a result of an agent's actions in the [light-filled] world. In so doing it shifted the problem of vision away from that of construction of rich internal representations of an 'out there' world, to that of active exploration of the environment 'on demand'; conscious experience being *brought forth* via a series of [saccadic] movements that either confirm (or disabuse) the notion that the world *actually is* of the form currently anticipated.

In contrast to classical approaches to cognition and AI, O'Regan and Noë's sensorimotor theory - and more broadly enactivism in general - potentially account for *why* our conscious experience of the world appears to us as it does; if correct this is a significant achievement and one that may offer new insight into at least some of the *hard problems of consciousness*[7].

The plenary talks and keynotes

It was not coincidental that the keynote speakers at AISB50 offered talks presenting new insights on both classical and 4Es related AI and cognitive science: emerging from the classical branches of AI, John Barnden presented the first public lecture of the conference entitled *Creative Metaphor, Mind Out! Or Rather, Mind In.* John Barnden is Professor of Artificial Intelligence at the University of Birmingham where his research is centrally con-

forth, in contra-position to the more classical views of either cognitivism or connectionism".

[6]Here the term 'enactive approach' is taken from Noë [12] where he states, "What I call here the enactive approach was first presented in [13]. I refer to the view as the sensorimotor contingency theory. Hurley and I, in joint work, deploy another term: the dynamic sensorimotor account. I borrow the term enactive from Francisco Varela and Evan Thompson (Varela, Thompson and and Rosch 1991 [19]), although I may not use it in exactly their sense. I use the term because it is apt, and to draw attention to the kinship of our view and theirs."

[7]E.g. In part it was this new, enactive, theory of visual perception that prompted AISB collaboration with 'The Colour Group' (Great Britain) on the symposium *New perspectives on colour.*

cerned with investigating figurative language, metaphor and Artificial Intelligence [2].

The second evening public lecture, *The Painting Fool: Weak and Strong Computational Creativity Research in Action*, was presented by Simon Colton. Simon is Professor of Computational Creativity at Goldsmiths, University of London. Computational Creativity is a sub-area of Artificial Intelligence research, which involves the study of software that can take on some of the creative responsibility in arts and science projects. In his talk Simon demonstrated the "thepaintingfool.com" in action - an artificial intelligence that he hopes will one day be accepted as an artist in its own right.

The third public lecture stemmed from a recent collaboration between 'The AISB', Dr. Kate Devlin (Goldsmiths) and 'The Colour Group of Great Britain'; a collaboration that resulted in the launch of a new AISB50 symposium entitled *New perspectives on Colour*. To celebrate, closing the AISB50 Convention on Friday evening, Dr. Hannah Smithson (Oxford) presented a special lecture entitled *New perspectives on colour from a 13th century account of light, material and rainbows*.

For the opening conference plenary address organisers were delighted to welcome Professor Susan Stepney from the University of York, UK, who enquired *When does a slime mould compute?* Susan's work in non-conventional computing sits somewhere between classical and the 4Es approaches to Artificial Intelligence and emphasises the importance of the *physical embodiment* of the *computational* substrate. At last year's AISB convention the 6^{th} Computing and Philosophy symposium was provocatively entitled, *The scandal of computing: 'What is computation?'* and further developing this theme Susan's plenary investigated precisely what it meant to say of a physical system 'that it computed'.

The second convention plenary was from Professor Lucy Suchman (Lancaster) who gave an elegant and wide-ranging lecture entitled *Human(oid) Robot Reconfigurations*. In seminal work, undertaken whilst she was based at Xerox's Palo Alto Research Centre between 1979 and 2000, Lucy fundamentally challenged common assumptions behind the design of interactive systems, with an incisive anthropological argument that human action is constantly constructed and reconstructed from dynamic interactions with the material and social worlds [18]; as such her strongly resonates with the 'embedded (or 'situated') theory of cognition (which emphasises the importance of the environment as an integral part of the cognitive process) and the broader 4Es approach[8].

Professor Terence Deacon's (Berkeley) recent high-profile monograph *Incomplete Nature* fundamentally challenges classical approaches to intelligence and cognition and in so doing it has won high praise from all quarters; perhaps most surprisingly from

[8] For an alternative analysis of related themes that makes the link (via autopoiesis) to the 4Es approach further explicit, see Winograd & Flores [20].

the philosopher Daniel Dennett. Over his lifetime Dennett has been a stalwart champion of computational theories of cognition against what he somewhat derisively labels its 'romantic challengers'[9], so on the surface it appears something of a volte-face to read his praise of Deacon's book promoting 'romantic" cognitive science, concluding that: "*Deacon, with his more ambitious exercise of reconstruction, has me re-examining my fundamental working assumptions*" [7]. In his plenary at AISB50 Terry outlined a new vision for a 'living machine' as one that instantiates principles of autogenesis - in which multiple self-organizing processes are linked by virtue of each producing the critical boundary constraints that maintain the others.

The final convention plenary *Ethical dilemma of the AnthropoRobotic* was from Professor Humberto Maturana (Instituto de Formación Matriztica, Chile). Maturana first came to international fame with the 1959 paper "*What the frog's eye tells the frog's brain*" [9] one of the most cited papers in the Science Citation Index. Over seven decades Humberto's research has touched on cybernetics, languaging, autonomy and enactivism and extends to philosophy, cognitive science and even family therapy. These days, however, along with with his protege Francisco Varela, Humberto is perhaps best known for his work on *autopoiesis* [11] - a groundbreaking thesis about the nature of reflexive feedback mechanisms in living systems that led them both to conclude: "*Living systems are cognitive systems, and living as a process is a process of cognition. This statement is valid for all organisms, with or without a nervous system*".

Framed by seven apposite keynote talks, the A-Eye Computer Art exhibition and the theatrical premiere of *MIL-STD-1815*[10], the twenty-four symposia that comprised this 50th convention undoubtedly offered a unique, exciting and celebratory snap-shot of Artificial Intelligence at a pivotal branch point at this, the golden anniversary party of the AISB. Viewed in this context AI is clearly once again positioned at a pivotal branch point; time will tell along which stream future AI will most effectively flow.

∴

Acknowledgments: AISB50 organisers gratefully acknowledge the Artificial Intelligence journal, Springer, Taylor-Francis and CHOICE Radio Cars for their generous financial support. Elements of this conference report have been adapted from Bishop's introduction to the AISB50 'Selected Pa-

[9]For the so called 'romantic side' Dennett claims "Romanes and Baldwin, Kropotkin, Stephen Jay Gould, Humberto Maturana, Francisco Varela, Stuart Kauffman, Roger Penrose, Ilya Prigogine, Rupert Sheldrake, and the philosophers John Haugeland, Evan Thompson, Alicia Juarrero, John Searle, and - off the map, now - Jerry Fodor and Thomas Nagel" [7].

[10]A new production premiered during the conference at Goldsmiths' George Wood theatre on April 2nd 2014, connecting thoughts, ideas and biography relating to Alan Turing, Ada Lovelace, Charles Babbage and Snow White.

pers' volume [ISBN 978-1-908187-42-0]; Bishop & Nasuto [3] and Bishop & Martin [4]. This short article is an abridged version of an extended conference report submitted to the Artificial Intelligence journal.

Bibliography

[1] Ashby, R., (1952), Design for a Brain, Chapman & Hall.

[2] Barnden, J. & Holyoak, K.J., (1994), Analogy, Metaphor and Reminding (Advances in Connectionist & Neural Computation Theory), Intellect Books.

[3] Bishop, J.M. & Nasuto, S.J., (2005), Second Order Cybernetics and Enactive Perception, Kybernetes: 34(9/10), pp. 1309-1320, Emerald.

[4] Bishop, J.M., & Martin, A., (2014), (Eds), Contemporary Sensorimotor Theory, SAPARE: Studies in Applied Philosophy, Epistemology and Rational Ethics 15, Springer.

[5] Chalmers, D., (1996), The Conscious Mind: in Search of a Fundamental Theory, Oxford University Press, UK, Oxford UK.

[6] Deacon, T., (2012), Incomplete Nature: How Mind Emerged from Matter. W. W. Norton & Company.

[7] Dennett, D., (2013), Aching Voids and Making Voids, Quarterly Review of Biology: 88(4), pp.321-324.

[8] Dreyfus, H.L. & Dreyfus, S.E., (1988), Making a Mind versus Modelling the Brain: Artificial Intelligence at a Branch Point, Artificial Intelligence 117(1), pp. 15-43.

[9] Lettvin, J. Y., Maturana, H. R., McCulloch, W.S. & Pitts, W. H., (1959), What the Frog's Eye tells the Frog's Brain, Proceedings of the IRE: 47(11), pp. 1940-1951.

[10] Levine, J., (1983), Materialism and Qualia: the Explanatory Gap, Pacific Philosophical Quarterly: 64, pp. 354-361.

[11] Maturana, H. & Varela, F., (1980), Autopoiesis and Cognition: the Realization of the Living. Cohen, R.S. & Wartofsky, M.W., (Eds.), Boston Studies in the Philosophy of Science: 42.

[12] Noë, A, (2004), Action in Perception, The MIT Press: Cambridge, MA.

[13] O'Regan, K., & Noë, A, (2001), A Sensorimotor Account of Vision and Visual Consciousness, Behaviour and Brain Sciences: 24, pp. 939-1031.

[14] O'Regan, J.K, (2011), Why Red Doesn't Sound like a Bell: Understanding the Feel of Consciousness', Oxford University Press: Oxford, UK.

[15] Penrose, R., (1989), The Emperor's New Mind: Concerning Computers, Minds, and the Laws of Physics, Oxford University Press, Oxford UK.

[16] Preston, J. & Bishop, J.M., (eds), (2002), Views into the Chinese

Room: New Essays on Searle and Artificial Intelligence, Oxford University Press, Oxford UK.

[17] Searle, J. R., (1980), Minds, Brains, and Programs. Behavioral and Brain Sciences 3 (3): 417-457, Cambridge, UK.

[18] Suchmann, L., (1987), Plans and Situated Actions: The Problem of Human-Machine Communication. Cambridge University Press.

[19] Varela, F.J., Thompson, E. & Rosch, E., (1991), The Embodied Mind: Cognitive Science and Human Experience. The MIT Press.

[20] Winograd, T. & Flores, F., (1986), Understanding Computers and Cognition: A New Foundation for Design. Ablex Publishing Corporation.

J. Mark Bishop
Professor of Cognitive Computing
Department of Computer Science
Goldsmiths, Univ. London
United Kingdom

Book review: Radicalizing Enactivism: Basic Minds without Content (Hutto & Myin, 2012)

by *Paulo de Jesus* (Goldsmiths, Univ. London, United Kingdom)

Cognitive science has been dominated by two fundamental principles, representationalism and computationalism. Cognitive systems build, store and manipulate inner (to the system, usually its brain) representations, while computational processes are responsible for the appropriate manipulation of these representations. In their book, Hutto and Myin (H&M) present a serious challenge to both of these principles.

The first two chapters set the scene by distinguishing two varieties of representationalism; a stronger version, dubbed "CIC" ("Cognition Necessarily Involves Content"), the proponents of which are committed to a form of "hyper-intellectualism" which maintains that representational content in the brain is a necessity for genuine cognition; and a weaker version, termed "CEC" ("Conservative Embodied/Enactive Cognition) that endorses embodiment but is deemed conservative, because it retains CIC's commitment to the idea that basic cognition is representational and involves manipulating content. In contrast, for H&M, only creatures with fully developed conceptual abilities scaffolded by public language, can or need be described in content-involving terms, and the significant majority of cognitive activities are non-contentfull/non-representational in character.

Enter "REC" (Radically Enactive (or Embodied) Cognition), the authors alternative to representationalism, a naturalist, radicalised enactivism that argues that "basic minds" do not contain, nor process informational content. Unlike the majority of other approaches to embodied cognition REC is different in that it not only targets proponents of traditional computational accounts of mind, but also more non-traditional accounts closer to their own. A central message here is that only a thoroughly radical anti-representationalism can save enactivism.

For REC, basic cognition is a "dynamically unfolding, situated embodied interactions and engagements with worldly offerings" (p. ix) that commits the authors to an "embodiment thesis". These interactions are essentially nonlinear and recursive, making it impossible to clearly delineate an "inner" domain of mentality from a separate "outer" domain of environmental causal factors. Therefore, basic cognition involves brain, body and environment tightly interconnected in a mutual and simultaneous dance of reciprocal causation. This further commits the authors to a "developmental-explanatory thesis". The idea here is that "mentality-constituting interactions are grounded in, shaped by, and explained by nothing more than the history of an organism's

previous interactions" (p. 8).

The core of the book is chapter 4, which introduces the "Hard Problem of Content", and "aims to show that proponents of CIC currently lack any naturalistically credible account of content upon which to ground their theorizing about basic minds" (p. 57). H&M posit that current naturalist accounts of mental content are incapable of adequately addressing this issue, and push further by arguing that informational content is incompatible with explanatory naturalism, and sooner or later both CIC and CEC will have to face up to this predicament.

If something has content, then it must have some special properties, some conditions of satisfaction such as truth, reference, accuracy, veridicality. But as H&M note, this has some undesirable consequences for explanatory naturalism, since it cannot account for states with these properties with only the standard naturalistic acceptable notions of covariance or information-as-indication. This is because covariant relations are simply dyadic causal relations where one state of affairs stands in a binary relation to another but does not say anything about the other state of affairs. This leaves CIC and CEC with a dilemma: either accept the notion of information-as-covariance but deny that cognition involves information-as-content with truth bearing properties, or if they insist that content exists they have to provide a naturalistic feasible solution to the HPC, which as H&M point out has proved near impossible. REC implies that at the level of basic engagement simple dyadic covariance relations are sufficient to explain basic cognition. However, covariance should not be confused with content, while covariance can account for basic forms of cognition, at this level content is neither involved nor necessary.

The concluding chapters applies REC to two specific debates. Chapter 7 criticises the Extended Mind Hypothesis as presented by Clark & Chalmers (1998) and further developed by Menary (2007) and Sutton (2010), for its implicit commitment to CIC. "Extensive Minds" is presented as the logical consequence of embracing REC. From the REC perspective basic minds are "fundamentally, constitutively already world-involving." (p. 137), As with the extended mind hypothesis, minds are not understood as brain-bound but as "extensive" and always dependent on brain, body, and world. But unlike the extended mind extensive minds do not traffic in contents. Chapter 8 concludes the book by arguing that taking REC seriously, in the domain of consciousness, implies that the supervenience base for phenomenal consciousness is "wide" and includes brain, body and environment.

Paulo de Jesus
PhD Candidate
Department of Computing
Goldsmiths Univ. London, UK

Event: Emergence of Consciousness

by *Janet Gibbs* (Kings College London)

The 3rd AISB workshop set out to examine questions about consciousness: What is it, and is it unique to humans? Did it "emerge", and if so how, when and why? Is there a continuum between consciousness and non-consciousness, and is there something special about self-consciousness?

The first keynote speakers addressed questions of consciousness in animals and machines. Murray Shanahan first presented the concept of Global Workspace Architecture, and showed how the wiring of many animal brains appears to fit this structure, and asked whether machines endowed with equivalent 'brains' would become conscious. Steve Torrance pursued the concept of 'superintelligent' machines, which could arise as machines recursively create new machines a little more intelligent than themselves - and the likely concomitance of 'superconsciousness'.

Then, James Steele presented studies on the development of human physiology as compared with early hominids, and what this might tell us about the emergence of language and consciousness. Geoffrey Hunt offered a critique of Rossano's 'Archaeology of Consciousness', which proposes that the development of hand-axe technology is evidence of the development of consciousness.

In the general sessions, we heard from researchers at all levels from MSc student to established professor, offering empirical and theoretical research, and one speculative position paper on the role of time and 'runningness' in consciousness. The question of emergence was addressed from a range of perspectives, including when and why the concept of consciousness arose in scientific thought. Questions were raised on physicalism and substance dualism; and an alternative duality of abstract vs concrete was proposed, with the intriguing concept of 'panabstractism'. Developments were presented on Dennett's intentional stance in the light of contextual emergence, and on Davidson's account of the role of language in rationality, as well as what enactivism has to say about Mary, bats and zombies. Some 'relatively neglected' writings of Locke were reviewed in light of their use by subsequent research. The final session focussed on social aspects of consciousness, including aspects of agency and community, how self-awareness is mediated by other minds, and an evolutionary approach to morality and ethics.

Proceedings can be found in J. Consc. Studies 21(5-6), 2014.

Janet Gibbs
Database Administrator
Dept of Psychological Medicine
King's College London

Event: 14th International Conference on Computer Vision

by *Siyamalan Manivannan* (Univ. Dundee, United Kingdom)

ICCV is the premier computer vision research conference sponsored by IEEE, and is held every other year, focusing on cutting-edge research in computer vision, pattern recognition, and related areas. ICCV comprises the main conference and several co-located workshops and tutorials. Each time this conference attracts world-leading researchers and representatives of industry, from a wide range of disciplines associated with computer vision to share ideas, concepts and technologies.

The 14th ICCV was held in December 2013 at the Conference and exhibition Centre in Darling Harbour, Sydney, Australia, and consisted of four days of workshops/tutorials before and after the main conference, several invited talks, technical presentations, exhibits and a keynote speech. Out of nearly 1629 submissions 41 papers were selected for oral and 413 were selected for poster presentations. This year the overall acceptance rate was 27.87% and the percentage selected for oral presentations was 2.52%. The conference covered a wide spectrum of computer vision research including detection and classification, 3D computer vision, motion and tracking, face and gesture recognition, low-level vision and image processing, statistical methods and learning, and optimization methods. Among these categories, detection and classification received the most attention, being addressed by 138 out of 413 papers. The first and the last two days of the conference were filled with multi-track tutorials and workshops. In total, there were ten tutorials and twenty five workshop sessions with many invited talks by experts from various universities and industries.

Feature encoding plays an important role in computer vision to aggregate local features to get an image representation. On that topic, I found that three papers in particular, are related to my work and are interesting to me. The first one is 'Low-Rank Sparse Coding for Image Classification' by Zhang et al. In contrast to existing approaches, where local features are encoded independently, failing to capture the local image structure, the authors present a method that exploits local structure information among local features. To capture local structure information they assume that local regions can be well approximated by their low-rank approximations. First an image is segmented into super-pixels, then densely sampled local features from each super-pixel are jointly encoded using a learned dictionary. Finally, the image representation is obtained by applying max-pooling on the encoded features. Experiments show that the method outperforms current approaches, performing well on standard benchmark datasets. Authors claim that the method also reduces

time complexity, compared to the traditional sparse coding method as it encodes the features jointly.

The second paper, 'Group Sparsity and Geometry Constrained Dictionary Learning for Action Recognition from Depth Maps', by Luo et al., is also related to sparse coding for feature encoding. The first method learns a dictionary in an unsupervised way, but this work proposes a supervised discriminative class-specific dictionary learning approach. The basic idea is that the features computed from an image belonging to a particular class can be well reconstructed by the class-specific dictionary. To enforce that, the sparse coding approach has been modified to include two additional terms; group sparsity and geometry constraints. Group sparsity supports the reconstruction of features by the dictionary of the same class, and the geometric constraints force the features from the same class with high similarity to have similar coefficients. In addition to this supervised dictionary learning approach, a temporal pyramid matching (TPM) approach also has been proposed. TPM is nothing more than an extension to the well-known spatial pyramid matching in the temporal domain. TPM divides a video sequence into several segments along the temporal direction; histograms generated from segments by max-pooling are concatenated to form the video representation. TPM captures the order of the events/actions happening in that sequence. Improved performance has been reported compared to other supervised dictionary learning approaches on action recognition datasets captured using depth cameras.

I note a few problems, however: (1) dictionaries for feature encoding and final classifiers are learned independently, which may lead to reduced classification performance; (2) the dimensionality of the local feature is limited; (3) larger dictionaries are preferred to get a high accuracy, which increases the complexity of the final classifier; (4) increased time complexity for feature encoding due to the optimizations in sparse coding,. In contrast to the above approaches, the dictionary-learning approach by 'Hierarchical Joint Max-Margin Learning of Mid and Top Level Representations for Visual Recognition' by Lobel et al., constructs the dictionary with a set of linear classifiers, which are jointly learned with the final image-level classifier. In this approach an image region is represented by applying max-pooling over the classifier scores of local features. This method has several advantages, including reduced dictionary size, and improved time complexity for feature encoding. Experiments on several datasets show that this approach can lead to state-of-the art performance with comparatively small size dictionaries.

Siyamalan Manivannan
PhD Candidate
School of Computing
Univ. Dundee, United Kingdom

Announcements

Upcoming workshop on "Figurative language: its patterns and meanings in domain-specific discourse"

The 5th AISB Workshop is jointly organized as a workshop of the Institute of Advanced Studies at the University of Birmingham "Figurative language: its patterns and meanings in domain-specific discourse", 18th–19th August 2014, University of Birmingham, UK.
Forms of figurative language such as metaphor and metonymy are key resources for communicating domain-specific information in an accessible way. Modelling such patterns of communication is a key aim of academic disciplines such as linguistics, discourse studies, and psycholinguistics, and understanding such phenomena is an emerging goal within Artificial Intelligence and the related field of Natural Language Processing.
Important dates:

- July 18th: Close of registration
- August 18th-19th: Workshop

Organisers: Prof. John Barnden (J.A.Barnden AT cs.bham.ac.uk), Dr. Andrew Gargett (A.D.Gargett AT cs.bham.ac.uk).

More at:
http://cs.bham.ac.uk/~gargetad/figurative-language-workshop-birmingham-2014.html

Call for AISB workshops

The AISB is funding a series of workshops to be held across the country, covering a wide range of themes pertinent to the aims of the AISB. These events are abstract-only and free for all AISB members. Current non-members would be able to attend for the cost of AISB membership, which they will be asked to arrange and pay for in advance by submitting a completed application form to the Executive Office. They would then be eligible to a year's membership of the Society. This applies to speakers and audience alike. Refreshments (coffee and teas) are funded by the AISB. If you are interested in hosting one of these events in your home institution, you will find information on what you will need to do on this page: http://www.aisb.org.uk/events

Make your voice heard!

The Q always needs fresh contributions. If you are a student, and you are wrapping up your thesis, you may want to advertise your work, send us an abstract. If you would like to attend that very fine conference, and you need a bit of help to fund your travel, speak to us about travel awards in exchange of a one page report. Have a look on our website for a list of books we are sent by all major publishers, speak to us, write a review and keep the book. If you want to reach the community and make your voice heard about a topic close to your heart, send us an article.

Dear Aloysius. . .

Due to recent attacks from various Security Agencies against the community of Hackers, Father Hacker has decided to lay low for a little while, and will resume answering your numerous queries after a well-deserved summer break.

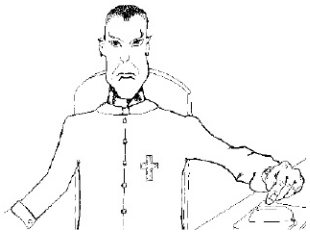

Fr. Aloysius Hacker
Cognitive Divinity Programme
Institute of Applied Epistemology

Back matter

AISB Quarterly Editor – aisbq@aisb.org.uk
Dr Etienne B Roesch (*Univ. Reading*)

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. 139, June, 2014

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://aisb.org.uk

ISSN 0268-4179