AISB Quarterly

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

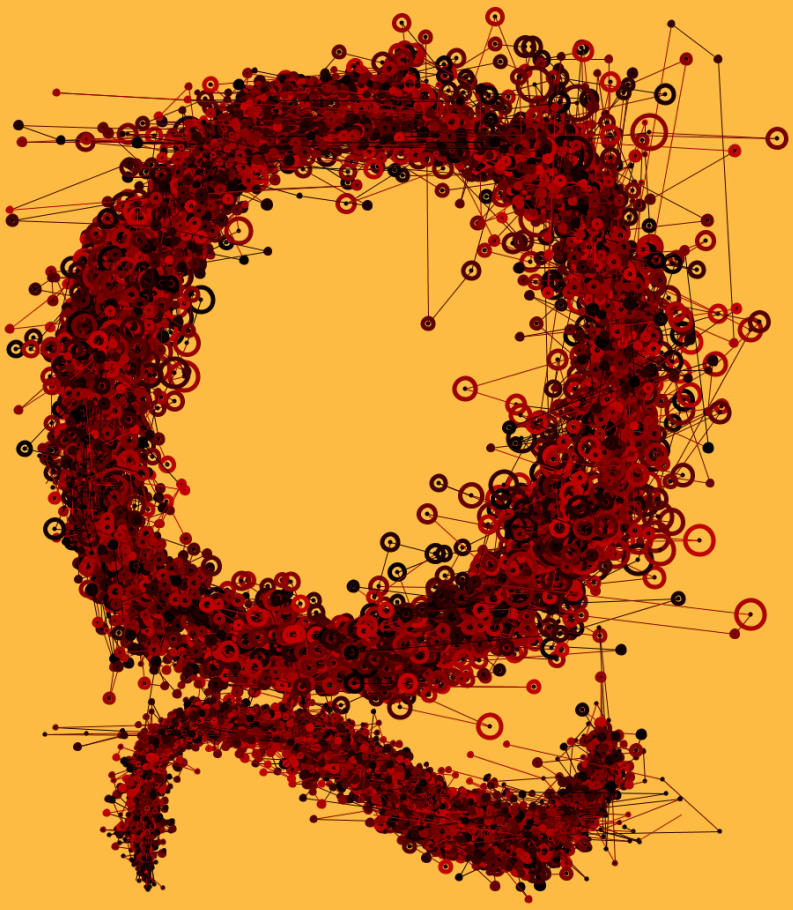

No. 137 October, 2013

All the latest Society information is available
at the newly renovated AISB website:
http://aisb.org.uk

The cover of this issue is designed by visualising the behaviour of agents aiming to trace the Q of the AISB Quarterly magazine. The agents are powered by two nature inspired swarm intelligence algorithms and a biological mechanism.

The swarm intelligence algorithms are Particle Swarm Optimisation simulating the behaviour of birds flocking, and Stochastic Diffusion Search, which mimics the recruitment behaviour of one species of ants, Leptothorax acervorum.

The biological mechanism is inspired by the behaviour of blood flow and cells in blood vessels. This particular design deployed the concept in Outward Eutrophic Remodelling, where the concept of high and low blood pressure and its impact on the vessel calibre is utilised. The details of the hybridisation can be found in the following book chapter:

Mohammad Majid al-Rifaie, Ahmed Aber and Mark Bishop (2012), Cooperation of Nature and Physiologically Inspired Mechanism in Visualisation, Book Chapter, *Biologically-Inspired Computing for the Arts: Scientific Data through Graphics*, DOI: 10.4018/978-1-46660-942-6, ISBN13: 9781466609426, IGI Global, USA.

Swarmic Art by: © al-Rifaie, Swarms & Blood Vessels

Editorial

Welcome to Q137. Now that the new format is established we are starting to get a little more adventurous with the cover. Details of the first experiment from Mohammad Majid al-Rifaie can be found on the inside front cover opposite. Mohammad will be providing more examples of his work in future editions but we also invite other members to submit their artistic interpretations of the Q. Contact the editors for further details if you're interested.

The first of this edition's articles is by Professor Ronald Arkin providing an alternative view to that put forward by Professors Noel Sharkey and Lucy Suchman in the last issue regarding the potential costs and benefits of autonomous military robots. In contrast to Sharkey and Suchman, Arkin argues that the potential benefits of intelligent machines in terms of the reduction in atrocities, civilian casualties and collateral damage may eventually lead to the conclusion that the moral course of action is to deploy them in combat situations.

We have published both sides of the argument in the hope of engendering a debate amongst the AISBQ membership and we encourage you to contribute. Send your thoughts to us at the email address on page 31.

The second article, by Dr Artemis Parvizi, discusses the problem of updating ontologies and outlines her recent thesis work developing a system to automate the addition of concepts and her evaluation of the effect of such concept addition to the resulting ontology.

This edition also contains two book reviews; the first, Pieraccini's (2012) *The Voice in the Machine* is reviewed by his long-time colleague and leading expert in the field, Professor Roger Moore. The second, Bekkerman, Bilenko and Langford's (2011) *Scaling up Machine Learning: Parallel and Distributed Approaches* is reviewed by Dr Ibrahim Venkat.

Finally, to round up this edition, we have reviews of three recent events by Claire Gerrard, Jens Koed Madsen, and Janet Gibbs, followed, as usual, by Fr. Hacker's sage advice to troubled souls.

We hope you enjoy this issue and feel inspired to pick up the (metaphorical) pen to contribute to the autonomous weapons debate—or any other matter for that matter.

The Q editors

Lethal Autonomous Systems and the Plight of the Non-combatant

by *Ronald Arkin* (Georgia Institute of Technology)

It seems a safe assumption, unfortunately, that humanity will persist in conducting warfare, as evidenced over all recorded history. New technology has historically made killing more efficient, for example with the invention of the longbow, artillery, armored vehicles, aircraft carriers, or nuclear weapons. Many view that each of these new technologies has produced a Revolution in Military Affairs (RMA), as they have fundamentally changed the ways in which war is waged. Many now consider robotics technology a potentially new RMA, especially as we move towards more and more autonomous[1] systems in the battlefield.

Robotic systems are now widely present in the modern battlefield, providing intelligence gathering, surveillance, reconnaissance, target acquisition, designation and engagement capabilities. Limited autonomy is also present or under development in many systems as well, ranging from the Phalanx system "capable of autonomously performing its own search, detect, evaluation, track, engage and kill assessment functions"[2], fire-and-forget munitions, loitering torpedoes, and intelligent antisubmarine or anti-tank mines among numerous other examples. Continued advances in autonomy will result in changes involving tactics, precision, and just perhaps, if done correctly, a reduction in atrocities as outlined in research conducted at the Georgia Tech Mobile Robot Laboratory (GT-MRL)[3]. This paper asserts that it may be possible to ultimately create intelligent autonomous robotic military systems that are capable of reducing civilian casualties and property damage when compared to the performance of human warfighters. Thus, it is a contention that calling for an outright ban on this technology is premature, as some groups already are doing[4]. Nonetheless, if this technology is to be deployed, then restricted, careful and graded introduction into the battlefield of lethal autonomous systems must be standard policy as opposed to haphazard deployments, which I believe is consistent with existing International Humanitarian Law (IHL).

Multiple potential benefits of intelligent war machines have already been declared by the military, including: a reduction in friendly casualties; force multiplication; expanding the battlespace; extending the warfighter's reach; the ability to respond faster given the pressure of an ever increasing battlefield tempo; and greater precision due to persistent stare [constant video surveillance that enables more time for decision making and more eyes on target]. This argues for the inevitability of development and deployment of lethal autonomous systems from a military efficiency and economic standpoint, unless limited by IHL.

It must be noted that past and

present trends in human behavior in the battlefield regarding adhering to legal and ethical requirements are questionable at best. Unfortunately, humanity has a rather dismal record in ethical behavior in the battlefield. Potential explanations for the persistence of war crimes include[5]: high friendly losses leading to a tendency to seek revenge; high turnover in the chain of command leading to weakened leadership; dehumanisation of the enemy through the use of derogatory names and epithets; poorly trained or inexperienced troops; no clearly defined enemy; unclear orders where intent of the order may be interpreted incorrectly as unlawful; youth and immaturity of troops; external pressure, e.g., for a need to produce a high body count of the enemy; and pleasure from power of killing or an overwhelming sense of frustration. There is clear room for improvement and autonomous systems may help address some of these problems.

Robotics technology, suitably deployed, may assist with the plight of the innocent noncombatant caught in the battlefield. If used without suitable precautions, however, it could potentially exacerbate the already existing violations by human soldiers. While I have the utmost respect for our young men and women warfighters, they are placed into conditions in modern warfare under which no human being was ever designed to function. In such a context, expecting a strict adherence to the Laws of War (LOW) seems unreasonable and unattainable by a significant number of soldiers[6]. Battlefield atrocities have been present since the beginnings of warfare, and despite the introduction of International Humanitarian Law (IHL) over the last 150 years or so, these tendencies persist and are well documented,[7] even more so in the days of CNN and the Internet. 'Armies, armed groups, political and religious movements have been killing civilians since time immemorial.'[8] 'Atrocity. . . is the most repulsive aspect of war, and that which resides within man and permits him to perform these acts is the most repulsive aspect of mankind'.[9] The dangers of abuse of unmanned robotic systems in war, such as the Predator and Reaper drones, are well documented; they occur even when a human operator is directly in charge.[10]

Given this, questions then arise regarding if and how these new robotic systems can conform as well as, or better than, our soldiers with respect to adherence to the existing IHL. If achievable, this would result in a reduction in collateral damage, i.e., noncombatant casualties and damage to civilian property, which translates into saving innocent lives. If achievable this could result in a moral requirement necessitating the use of these systems. Research conducted in our laboratory[11] focuses on this issue directly from a design perspective. No claim is made that our research provides a fieldable solution to the problem, far from it. Rather these are baby steps towards achieving such a goal, including the development of a prototype proof-of-concept system tested in simulation. Indeed, there may be far better approaches than the one we currently employ, if the research community can focus on the plight of the noncombatant and how technology may possibly ameliorate the situation.

As robots are already faster, stronger, and in certain cases (e.g., Deep Blue, Watson[12]) smarter than humans, is it really that difficult to believe they will be able to ultimately treat us more humanely in the battlefield than we do each other, given the persistent existence of atrocious behaviors by a significant subset of human warfighters?

Why technology can lead to a reduction in casualties on the battlefield

Is there any cause for optimism that this form of technology can lead to a reduction in non-combatant deaths and casualties? I believe so, for the following reasons.

– The ability to act conservatively: i.e., they do not need to protect themselves in cases of low certainty of target identification. Autonomous armed robotic vehicles do not need to have self-preservation as a foremost drive, if at all. They can be used in a self-sacrificing manner if needed and appropriate without reservation by a commanding officer. There is no need for a 'shoot first, ask-questions later' approach, but rather a 'first-do-no-harm' strategy can be utilized instead. They can truly assume risk on behalf of the noncombatant, something that soldiers are schooled in, but which some have difficulty achieving in practice.

– The eventual development and use of a broad range of robotic sensors better equipped for battlefield observations than humans currently possess. This includes ongoing technological advances in electro-optics, synthetic aperture or wall penetrating radars, acoustics, and seismic sensing, to name but a few. There is reason to believe in the future that robotic systems will be able to pierce the fog of war more effectively than humans ever could.

– Unmanned robotic systems can be designed without emotions that cloud their judgment or result in anger and frustration with ongoing battlefield events. In addition, 'Fear and hysteria are always latent in combat, often real, and they press us toward fearful measures and criminal behavior'[13]. Autonomous agents need not suffer similarly.

– Avoidance of the human psychological problem of 'scenario fulfilment' is possible. This phenomenon leads to distortion or neglect of contradictory information in stressful situations, where humans use new incoming information in ways that only fit their pre-existing belief patterns. Robots need not be vulnerable to such patterns of premature cognitive closure. Such failings are believed to have led to the downing of an Iranian airliner by the USS Vincennes in 1988.[14]

– Intelligent electronic systems can integrate more information from more sources far faster before responding with lethal force than a human possibly could in real-time. These data can arise from multiple remote sensors and intelligence (including human) sources, as part of the Army's network-centric warfare concept and the concurrent development of the Global Information Grid. 'Military systems (including weapons) now on the horizon will be too fast, too small, too numerous and will create an environment too complex for humans to

direct'[15].

– When working in a team of combined human soldiers and autonomous systems as an organic asset, they have the potential capability of independently and objectively monitoring ethical behavior in the battlefield by all parties, providing evidence and reporting infractions that might be observed. This presence alone might possibly lead to a reduction in human ethical infractions.

Addressing some of the counter-arguments

But there are many counterarguments as well. These include the challenge of establishing responsibility for war crimes involving autonomous weaponry, the potential lowering of the threshold for entry into war, the military's possible reluctance to give robots the right to refuse an order, proliferation, effects on squad cohesion, the winning of hearts and minds, cybersecurity, proliferation, and mission creep.

There are good answers to these concerns I believe, and are discussed elsewhere in my writings[16]. If the baseline criterion becomes the outperforming of humans in the battlefield with respect to adherence to IHL (without mission performance erosion), I consider this to be ultimately attainable, especially under situational conditions where bounded morality [narrow, highly situation-specific conditions] applies[17], but not soon and not easily. The full moral faculties of humans need not be reproduced to attain to this standard. There are profound technological challenges to be resolved, such as effective in situ target discrimination and recognition of the status of those otherwise hors de combat, among many others. But if a warfighting robot can eventually exceed human performance with respect to IHL adherence, that then equates to a saving of noncombatant lives, and thus is a humanitarian effort. Indeed if this is achievable, there may even exist a moral imperative for its use, due to a resulting reduction in collateral damage, similar to the moral imperative Human Rights Watch has stated with respect to precision guided munitions when used in urban settings[18]. This seems contradictory to their call for an outright ban on lethal autonomous robots[19] before determining via research if indeed better protection for non-combatants could be afforded.

Let us not stifle research in the area or accede to the fears that Hollywood and science fiction in general foist upon us. By merely stating these systems cannot be created to perform properly and ethically does not make it true. If that were so, we would not have supersonic aircraft, space stations, submarines, self-driving cars and the like. I see no fundamental scientific barriers to the creation of intelligent robotic systems that can outperform humans with respect to moral behavior. The use and deployment of ethical autonomous robotic systems is not a short-term goal for use in current conflict, typically counterinsurgency operations, but rather will take considerable time and effort to realize in the context of interstate warfare and situational context involving bounded morality.

A plea for the non-combatant

How can we meaningfully reduce human atrocities on the modern battlefield? Why is there persistent failure and perennial commission of war crimes despite efforts to eliminate them through legislation and advances in training? Can technology help solve this problem? I believe that simply being human is the weakest point in the kill chain, i.e., our biology works against us in complying with IHL. Also the oft-repeated statement that "war is an inherently human endeavor" misses the point, as then atrocities are also an inherently human endeavor, and to eliminate them we need to perhaps look to other forms of intelligent autonomous decision-making in the conduct of war. Battlefield tempo is now outpacing the warfighter's ability to be able to make sound rational decisions in the heat of combat. Nonetheless, I must make clear the obvious statement that peace is unequivocally preferable to warfare in all cases, so this argument only applies when human restraint fails once again, leading us back to the battlefield.

While we must not let fear and ignorance rule our decisions regarding policy towards these new weapons systems, we nonetheless must proceed cautiously and judiciously. It is true that this emerging technology can lead us into many different futures, some dystopian. It is crucially important that we not rush headlong into the design, development, and deployment of these systems without thoroughly examining their consequences on all parties: friendly forces, enemy combatants, civilians, and society in general. This can only be done through reasoned discussion of the issues associated with this new technology. Toward that end, I support the call for a moratorium to ensure that such technology meets international standards before being considered for deployment as exemplified by the recent report from the United Nations Special Rapporteur on Extrajudicial, Summary, or Arbitrary Executions.[20] In addition, the United States Department of Defense has recently issued a directive[21] restricting the development and deployment of certain classes of lethal robots, which appears tantamount to a quasi-moratorium.

Is it not our responsibility as scientists and citizens to look for effective ways to reduce man's inhumanity to man through technology? Where is this more evident than in the battlefield? Research in ethical military robotics can and should be applied toward achieving this end. The advent of these systems, if done properly, could possibly yield a greater adherence to the laws of war by robotic systems than from using soldiers of flesh and blood alone. While I am not averse to the outright banning of lethal autonomous systems in the battlefield, if these systems were properly inculcated with a moral ability to adhere to the laws of war and rules of engagement, while ensuring that they are used in narrow bounded military situations as adjuncts to human warfighters, I believe they could outperform human soldiers with respect to conformance to IHL. The end product then could be, despite the fact

that these systems could not ever be expected to be perfectly ethical, a saving of noncombatant lives and property when compared to human warfighters' behaviour.

This is obviously a controversial assertion, and I have often stated that the discussion my research engenders on this subject is as important as the research itself. We must continue to examine the development and deployment of lethal autonomous systems in forums such as the United Nations and the International Committee of the Red Cross to ensure that the internationally agreed upon standards regarding the way in which war is waged are adhered to as this technology proceeds forward. If we ignore this, we do so at our own peril.

The Way Forward?

It clearly appears that the use of lethality by autonomous systems is inevitable, perhaps unless outlawed by international law – but even then enforcement seems challenging. But as stated earlier, these systems already exist: the Patriot missile system, the Phalanx system on Aegis class cruisers, anti-tank mines, and fire-and-forget loitering munitions all serve as examples. A call for a ban on these autonomous systems may have as much success as trying to ban artillery, cruise missiles, or aircraft bombing and other forms of standoff weaponry (even the crossbow was banned by Pope Innocent II in 1139[22]). A better strategy perhaps is to try and control its uses and deployments, which existing IHL appears at least at first glance to adequately cover, rather than a call for an outright ban, which seems unenforceable even if enacted.

The horse is out of the barn. Under current IHL, these systems cannot be developed or used until they can demonstrate the capability of adequate distinction, proportionality, and shown that they do not produce unnecessary suffering, and must only be used given military necessity. Outside those bounds any individuals responsible should be held accountable for violations of International Humanitarian Law, whether they are scientists, industrialists, policymakers, commanders, or soldiers. As these systems do not possess moral agency, the question of responsibility becomes equated to other classes of weapon systems, and a human must always ultimately bear responsibility for their use[23]. Until it can be shown that the existing IHL is inadequate to cover this RMA, only then should such action be taken to restructure or expand the law. This may be the case, but unfounded pathos-driven arguments based on horror and Hollywood in the face of potential reductions of civilian casualties seems at best counterproductive. These systems counterintuitively could make warfare safer in the long run to the innocents in the battlespace, if coupled with the use of bounded morality, narrow situational use, and careful graded introduction.

Let it be restated that I am not opposed to the removal of lethal autonomous systems from the battlefield, if international society so deems it fit, but I think that this technology can actually foster humanitarian treatment of noncombatants if done correctly. I have argued to those that call for a ban, they would be better served by a call for a

moratorium, but that is even hard to envision occurring, unless these systems can be shown to be in clear violation of the LOW. It's not clear how one can bring the necessary people to the table for discussion starting from a position for a ban derived from pure fear and pathos.

For those familiar with the Martens clause[24] in IHL, a case could be made that these robotic systems potentially "violate the dictates of the public conscience". But until IHL lawyers agree on what that means, this seems a difficult course. I do believe, however, that we can aid the plight of noncombatants through the judicious deployment of these robotic systems, if done carefully and thoughtfully, particularly in those combat situations where warfighters have a greater tendency or opportunity to stray outside International Humanitarian Law. But what must be stated is that a careful examination of the use of these systems must be undertaken now to guide their development and deployment, which many of us believe is inevitable given the ever increasing tempo of the battlefield as a result of ongoing technological advances. It is unacceptable to be "one war behind" in the formulation of law and policy regarding this revolution in military affairs that is already well underway. The status quo with respect to human battlefield atrocities is unacceptable and emerging technology in its manifold forms must be used to ameliorate the plight of the noncombatant.

Ronald Arkin is Regents' Professor, Director of the Mobile Robot Laboratory, and Associate Dean for Research in the College of Computing at the Georgia Institute of Technology. This work was supported in part by the U.S. Army Research Office under Contract #W911NF-06-1-0252. Small portions of this essay appeared earlier in a Viewpoint article by the author appearing in the *Journal of Industrial Robots* 38:5, 2011, and from a more comprehensive treatment of the subject in the author's book *Governing Lethal Behavior in Autonomous Systems*, Taylor-Francis, 2009, and are included with permission.

Footnotes

[1] We do not use autonomy in the sense that a philosopher does, i.e., possessing free will and moral agency. Rather we use in this context a roboticist's definition: the ability to designate and engage a target without additional human intervention after having been tasked to do so.

[2] U.S. Navy, "Phalanx Close-in Weapons Systems", United States Navy Factfile, http://www.navy.mil/navydata/fact_display.asp?cid=2100&tid=487&ct=2, accessed 7/23/2013.

[3] R.C. Arkin, Governing Lethal Behavior in Autonomous Robots, Chapman-Hall, 2009.

[4] Notably Human Rights Watch, International Committee on Robot Arms Control (ICRAC) and Article 36.

[5] Bill, B. (Ed.), Law of War Workshop Deskbook, International and Operational Law Department, Judge Advocate General's School, June 2000; Danyluk, S., "Preventing Atrocities", Marine Corps Gazette, Vol. 8, No. 4, pp. 36-38, Jun 2000; Parks, W.H., "Crimes in Hostilities. Part I", Marine Corps Gazette, August 1976; Parks, W.H., "Crimes in Hostilities. Conclusion", Marine Corps Gazette, September 1976; Slim, H., Killing Civilians: Method, Madness, and Morality in War, Columbia University Press, New York, 2008.

[6] Surgeon General's Office, Mental Health Advisory Team (MHAT) IV Operation Iraqi Freedom 05-07, Final Report, Nov. 17, 2006.

[7] For a more detailed description of these abhorrent tendencies of humanity discussed in this context, see Arkin, R.C., "The Case for Ethical Autonomy in Unmanned Systems", Journal of Military Ethics, 9:4, pp. 332-341, 2010.

[8] Slim, H., Killing Civilians: Method, Madness, and Morality in War, Columbia University Press, New York, 2008, p. 3.

[9] Grossman, D., On Killing: The Psychological Cost of Learning to Kill in War and Society, Little, Brown and Company, Boston, 1995, p.229.

[10] Adams. J., "US defends unmanned drone attacks after harsh UN Report", Christian Science Monitor, June 5, 2010; Filkins, D., "Operators of Drones are Faulted in Afghan Deaths", New York Times, May 29, 2010; Sullivan, R., "Drone Crew Blamed in Afghan Civilian Deaths", Associated Press, May 5, 2010.

[11] For more information see Arkin, R.C., Governing Lethal Behavior in Autonomous Systems, Taylor and Francis, 2009.

[12] http://en.wikipedia.org/wiki/Deep_Blue_(chess_computer), http://en.wikipedia.org/wiki/Watson_(computer)

[13] Walzer, M., Just and Unjust Wars, 4th ed., Basic Books, 1977.

[14] Sagan, S., "Rules of Engagement", in Avoiding War: Problems of Crisis Management (Ed. A. George), Westview Press, 1991.

[15] Adams, T., "Future Warfare and the Decline of Human Decisionmaking", in Parameters, U.S. Army War College Quarterly, Winter 2001-02, pp. 57-71.

[16] E.g., Arkin, R.C, op. cit., 2009.

[17] Wallach, W. and Allen, C., Moral Machines: Teaching Robots Right from Wrong, Oxford University Press, 2010.

[18] Human Rights Watch, "International Humanitarian Law Issues in the Possible U.S. Invasion of Iraq", Lancet, Feb. 20, 2003.

[19] Human Rights Watch, "Losing Humanity: The Case Against Killer Robots", Nov. 19, 2012.

[20] Christof Heyns, Report of the Special Rapporteur on Extrajudicial, Summary, and Arbitrary Execution, United Nations Human Rights Council, 23rd Session, April 9, 2013.

[21] United States Department of Defense Directive Number 3000.09, Subject: Autonomy in Weapons Systems, November 21, 2012.

[22] Royal United Services Institute for Defence and Security Studies, "The Ethics & Legal Implications of Unmanned Vehicles for Defence and Security Purposes", Workshop webpage, held Feb. 27, 2008, http://www.rusi.org/events/ref:E47385996DA7D3, (accessed 5/12/2013).

[23] Cf. Arkin, R.C., "The Robot Didn't Do it.", Position Paper for the Workshop on

Anticipatory Ethics, Responsibility, and Artificial Agents, Charlottesville, VA., January 2013.

[24] The clause reads "Until a more complete code of the laws of war is issued, the High Contracting Parties think it right to declare that in cases not included in the Regulations adopted by them, populations and belligerents remain under the protection and empire of the principles of international law, as they result from the usages established between civilized nations, from the laws of humanity and the requirements of the public conscience." (Available at the ICRC website,
http://www.icrc.org/eng/resources/documents/misc/57jnhy.htm last visited on 30 April 2013).

Automatic Concept Addition to Ontologies Aided by Semantics

by *Artemis Parvizi* (Univ. of Aberdeen)

Ontologies, formal specification of a shared conceptualisation [5], are among popular data structures for representing knowledge in various domains and applications. Ontologies provide a shared understanding within a certain domain, of concepts and attributes, through a hierarchical representation [10]. The minimum level of detail and the minimal expressiveness to be input in an ontology is not clear [9]. What can be shared is only an approximation of a domain based on a finite set of existing examples and according to the judgement of the domain experts. Partly as a result of the ambiguity of the definition and partly due to the dynamic nature of knowledge, desirable formality and completeness will rarely emerge [6].

Various methods handle the situation when ontologies change over time, two of which are ontology evolution and ontology learning: the former focuses on adapting an ontology to change in a timely manner, and the latter attends to learning new facts about a domain. These fields may sound quite different, but they are surprisingly similar. While ontology evolution has produced many different strategies, ontology learning has a more practical approach. Nevertheless, both ontology evolution and ontology learning aim at keeping the semantic hierarchy up to date, and add more depth to ontologies by generating various attributes for concepts. Concept addition to an ontology, the aim of this thesis, is part of ontology change strategy. The intended methodology is to implement a system capable of generating taxonomic and non-taxonomic relations for an input concept via ontology graphs.

A main objective is to express information by a set of vertices and edges in a hierarchical structure. Ontologies can not only be visualised by graphs; they can be represented by graphs as well [4]. An ontology graph is capable of representing all the important structural features from the web ontology language (OWL). It has been suggested that a combination of algebraic and a graph-based approach normally performs better [3]. In an OWL ontology, it is common practice to check the consistency of an ontology after each change. The same practice must be applied to an ontology graph. Since an OWL ontology and an ontology graph represent the same knowledge, by having a system capable of converting an ontology graph into an OWL ontology, existing consistency checkers can be employed.

One of the contributing factors in ontology evolution is to evolve taxonomy and meaning at the same time. The emphasis is not only concept categorisation, but also expansion of concept's meaning through attribute addition. Despite the ongoing efforts towards simplifying ontology development and update, it is still difficult to

extract human knowledge and transfer it to machines; knowledge acquisition primarily depends on experts, and automation to some extent is possible [7]. The conventional approach is to extract information from structured, semi-structured, or unstructured data sources. Often when semi-structured sources such as WordNet [2], ConceptNet [8], and FrameNet [1] are involved, semantic similarity between concepts is calculated. In this thesis, generating and employing semantic similarity has been successfully explored.

Because it reduces the subjective effect of experts and users, automating taxonomic and non-taxonomic relation addition is quite influential. It is well known that employing ontology or domain experts is quite expensive [11], and due to the lack of knowledge, end-user feedback is not always desirable. Automation guarantees determinism, objectivity, and reproducibility [3]. However, reducing the human involvement in most stages of ontology engineering is extremely difficult.

For generating the taxonomy, this work combined semantic similarities with a predefined set of patterns. Although these patterns can often successfully generate a logically and semantically accurate hierarchy, they are not exclusive.

Non-taxonomic concept addition other than semantic similarities, also relies on the semantic and structural information within an ontology. Aided by semantic similarities, self-information is calculated, and some non-taxonomic relations generated.

In my recent work I have studied the effect of change on taxonomic and non-taxonomic relation additions and reported results on several ontologies. I have also investigated the causes of failure.

Since evaluation has always been a subject of debate and a standard ontology change evaluation technique has not been introduced, precision and recall of concept addition have been combined with a gold standard and used for ontology evaluation.

A number of selected ontologies have been converted into ontology graphs, and concepts have been detached one by one. The taxonomic and non-taxonomic engine intends to reattach the concepts. Detaching and reattaching one or more concepts through measures of precision and recall has been studied.

A comparison between the results of this work and others is enormously difficult due to the lack of ontology change approaches that incorporate precision and recall measurements, and the absence of contributing ontologies. However the performance of the taxonomic estimator has a maximum F-measure of 92%, and the non-taxonomic estimator has a maximum F-measure of 57%.

Further information about this project can be found in a paper recently presented at the *Eighth Australasian Ontology Workshop*, available at this link:

http://ceur-ws.org/Vol-969/paper6.pdf.

Reference

[1] Baker Collin, Fillmore Charles, and Lowe John, 'The berkeley framenet project', in *Proceedings of the 17th international conference on Computational linguistics - Volume 1, COL-*

ING '98, pp. 86–90, Stroudsburg, PA, USA, (1998). Association for Computational Linguistics.

[2] Christiane Fellbaum, *WordNet: An Electronic Lexical Database*, Bradford Books, 1998.

[3] Giorgos Flouris, Dimitris Plexousakis, and Grigoris Antoniou, 'A classification of ontology change', in *SWAP'06*, (2006).

[4] Stephan Grimm, 'Knowledge representation and ontologies', chapter Scientific Data Mining and Knowledge Discovery, 111–137, Springer, Berlin Heidelberg, (2010).

[5] T Gruber, 'A translation approach to portable ontology specifications', *Knowledge Acquisition*, **5**(2), 199–220, (1993).

[6] Nicola Guarino, Daneil Oberle, and Steffen Staab, *International Handbooks on Information Systems*, chapter What is an Ontology?, 1–17, Springer, 2004.

[7] Dragan Găsević, Dragan Djurić, and Vladan Devedžić, *The semantic web. In Model Driven Engineering and Ontology Development*, chapter 81-124, Springer, Berlin Heidelberg, 2009.

[8] H Liu and P Singh, 'Conceptnet – a practical commonsense reasoning tool-kit', *BT Technology Journal*, **22**, 211–226, (2004).

[9] Peter Mika, 'Knowledge representation on the semantic web', in *Social Networks and the Semantic Web*, volume vol. 5 of Semantic Web and Beyond, pp. 65–92. Springer, (2007).

[10] Natalya Noy and Deborah L. McGuinness, 'Ontology development 101: A guide to creating your first ontology', Technical report, Technical Report KSL-01-05, Stanford Knowledge Systems Laboratory,, (March 2001).

[11] Ljiljana Stojanovic, Alexander Maedche, Nenad Stojanovic, and Rudi Studer, 'Ontology evolution as reconfiguration-design problem solving', in *Proceedings of the 2nd international conference on Knowledge capture, K- CAP '03*, pp. 162–171, New York, NY, USA, (2003). ACM.

Artemis Parvizi
Research fellow
School of Natural and Computing Sciences
University of Aberdeen

Book review: The Voice in the Machine (Pieraccini, 2012)

by *Roger K. Moore* (Univ. Sheffield)

I've known Roberto for more years than probably either of us cares to remember—right back to the early 1980s when we were both young post-doctoral researchers starting out on our respective careers in speech technology—Roberto at the Italian Centro Studi e Laboratori Telecomunicazioni (CSELT) in Turin, and myself at the UK Government Speech Research Unit (SRU) in Malvern. These were exciting times, with several of the most significant advances in our field happening all around us. The first International collaborations (such as the NATO Research Study Group—RSG10) were paving the way for sharing standardised speech corpora, National and International funding agencies were starting to invest heavily in directed speech technology programmes, and the results of years of laboratory work were beginning to establish a tentative foothold in the commercial marketplace. Roberto was (and, I should say, continues to be) very much part of the action, so it was with eager anticipation that I awaited my copy of The Voice in the Machine. Touted as an examination of "six decades of work in science and technology to develop computers that can interact with humans using speech", I was keen to hear about his perspective on all the developments—positive and negative—that many of us 'old hands' have lived through (and participated in).

I was not disappointed. Although the catchy title might mistakenly imply that the main topic is speech synthesis, the subtitle—Building Computers that Understand Speech—makes it clear that the emphasis is in fact very much on automatic speech recognition and understanding with, unsurprisingly, a significant component dedicated to Roberto's own speciality—spoken language dialogue systems. Unusually for a book of this type, Roberto approaches the central topic from the high-level perspective of language and thought, rather than the low-level characteristics of speech signals. This novel structure (at least, novel in a book about speech) provides a compelling framework within which to emphasise the importance of meaning and context in spoken language interaction; a theme that pervades the rest of the book, particularly through a running humorous example of speech communication involving Roberto's dog—Ares—and a neighbour's stolen sausage!

The central theme of the book (and the title of the third Chapter) is the fundamental dichotomy between the GOFAI 'expert' approach and the use of 'brute force' statistical methods to solve the challenges posed by attempting to create automated speech technology systems. By taking a historical perspective, Roberto manages to convey, not just the algorithms that are now deployed in contemporary sys-

tems, but also the philosophical and practical arguments that culminated in the use of stochastic state machines and optimal graph-based search as the core basis for modelling the patterned behaviour in speech, language and dialogue. The book is packed with wonderful sub-headings that will have special meanings for us older folk such as "Finding your Way in Time Warp Land", "The Hidden Models of Markov", "Digit Wars", "Mad Cows" and, a particularly good one, "Invisible Speech"—a nod towards the phonetic writing system known as 'Visible Speech' developed by Alexander Melville Bell (Alexander Graham Bell's father) in the 1860s. Indeed Roberto's crucial insight—"the fate of any good technology is to become invisible"—provides a very satisfactory set of opening and closing parentheses on the entire endeavour.

Of particular importance is the central role that has been played by competitive benchmarking and assessment using shared standardised data. Roberto spells out the long and hard road that had to be taken in order to turn a rather ad-hoc engineering-oriented 'suck-it-and-see' field into the rigorous mathematical discipline that it is today. Of course, the imposition of a strict evaluation framework has its own particular downsides and, as Roberto correctly points out, researchers dependent on funding from Government agencies tended to become risk-averse and true innovations are now rare events in the continuous drive for guaranteed incremental improvements. And the safest and most effective approach to improving performance turns out to be... more training data. Indeed, to repeat one of the most famous quotes in the speech technology field—"There's no data like more data"—a phrase used by Roberto as another of his Chapter headings.

Unsurprisingly, given the career trajectory of the author, the book is very much focused on telephone-based applications, with an inevitable emphasis on developments in the United States, particularly those which took place at Bell Labs/AT&T. However, the style is not overly egocentric or self-promotional, and most of the major speech labs around the world receive a satisfactory acknowledgement. Nevertheless, there are some surprising omissions; some of the significant early speech recognisers are missing, for example there's no mention of GEC-Marconi's SR128 connected word recogniser (the first to fly in an aircraft), the French VecSys systems, NEC's ground-breaking DP100 system, or indeed of Loquendo—the very successful Italian speech company that spun out from Roberto's old lab in Turin. There is also no reference to Dennis Klatt's famously critical review of the ARPA SUR programme, or that HARPY was actually the outcome of a one-man PhD project and, perhaps the biggest omission of all, there is no reference to the Cambridge HTK hidden-Markov model toolkit which could easily be argued to have played the largest role in transforming worldwide speech technology R&D over the past twenty years. Likewise, the section on speech synthesis completely fails to mention recent developments in hidden Markov model based speech generation despite there being a section headed "Brute

Force and Statistics"!

Rather less surprising is that Roberto falls into the standard traps when getting down into details about speech itself. He presents the usual engineer's description of a 'phoneme' as a "distinct sound" corresponding to a particular configuration of the vocal tract. If only this were true—building speech technology systems would be so much easier! Almost all engineers fail to grasp the significance of the difference between phonetics and phonology, so it's a shame that Roberto missed the opportunity to educate the naïve reader into the importance of the phoneme as an abstract concept deriving from linguistic contrast. Likewise, 'coarticulation' is defined as "the continuous transition between consecutive phonemes", thereby missing the most important property of coarticulation that it can operate over many sound segments (such as nasalisation or lip-rounding occurring several phones in advance of a nasal or lip-rounded sound). Roberto also repeats the popular misconception that "vowel sounds in normal speech reach a point of stability in the middle... that doesn't depend on the preceding or following sound" which, I'm sorry to say, reflects the wishful thinking of the optimistic engineer more than the reality of speech.

Overall, *The Voice in the Machine* is a masterclass in communication. It is a highly personal account of the development of a key area of contemporary AI from an individual who has been involved almost from the beginning. It is written from the perspective of someone who has direct experience of both the rigour of the research laboratory and the urgency of the commercial environment. It is an unusual book in as much as the various algorithms and approaches are placed in their historical context rather than simply listed as elements in the speech technology toolbox. For older members of the community (myself included), it will be appreciated as a nostalgic trip down memory lane that reminds us that we are part of a relatively modern field of science and engineering that has by now acquired a reasonable amount of interesting history and given birth to a viable industry. For people who are new to the area, or who are interested in one of the main success stories in data-driven approaches to machine learning, it provides a definitive account of exactly how we came to be where we are today. If you've always wondered how (and why) speech technology people use hidden Markov models, but you've been put off by the more mathematical texts on offer, then this is the book for you. It should be essential reading for all students studying speech and language technology, but it also contains important lessons (both positive and negative) for other areas of AI where rigorous benchmarking has yet to be established.

Finally, this book is not the end of the speech technology story. It was clearly written just before the 2011 surprise release of Apple's speech-based personal assistant for the iPhone—Siri. Luckily, Roberto managed to slip in an epilogue in which he rightly points out that there's nothing dramatically different in Siri—it's just an example of something happening "in the right place at the right time". He leaves open

the question as to whether Siri is the beginning or the end of developments in speech technology, but we can probably guess his thoughts from his remark that "there is a huge distance between what we expect talking machines to do and what they actually can do". In fact no real attempt is made in this book to speculate on where we might go from here. So I look forward to a follow-up to *The Voice in the Machine* in which Roberto draws on his extensive knowledge and experience to take us on another gentle ramble, this time into possible futures for speech-based human machine interaction.

Reference

[1] Roberto Pieraccini, *The Voice in the Machine*, Building Computers that Understand Speech, MIT Press, 2012.

Roger K. Moore
Chair of Spoken Language Processing
Department of Computer Science
University of Sheffield

Book review: Scaling up Machine Learning: Parallel and Distributed Approaches (Bekkerman, Bilenko, & Langford, 2011)

by *Ibrahim Venkat* (Univ. Sains Malaysia)

Thanks to extensive research efforts, the prediction capabilities of machines that learn from data have evolved considerably over the recent decades. Nevertheless, as the needs of expanding population and industries grow rapidly, dimensions of large-scale data expand in multi-folds day by day. This book serves as a valuable resource for those who want to explore and apply scalable machine learning approaches that can learn from such voluminous data. The authors have knitted two vital paradigms, viz. machine learning and parallel, distributed approaches. Those who are looking up for a unified material, enriched with recent advancements and pragmatic scalable approaches that can enhance machine learning algorithms via sophisticated parallelization techniques, will find comprehensive coverage of these two paradigms in this text. However, the authors assume that the reader has already a basic understanding of both machine learning and parallel, distributed architectures.

The first part of the book gives a brief introduction to the recent programming frameworks that are suitable for machine learning algorithms and that can be adapted to parallel platforms. The authors have taken into account simple case studies comprising of quite popular machine learning algorithms, and illustrated how basic distributed models can deploy these algorithms to achieve scalability. The expertise of the authors play a good role in gradually taking the reader into advanced concepts and a variety of relevant problems without much burden. Advantages as well as bottlenecks of the techniques demonstrated are summarised well.

In the second part of the book, modern machine learning concepts such as probabilistic graphical models are leveraged with the parallelization of more sophisticated algorithms. Apart from software design issues, focus has also been given on customising hardware design. The authors have also taken efforts to show how parallel optimisation algorithms can be distributed over a large number of modern multicore processors using some state-of-the-art datasets in order to achieve high performance computing.

The third and final part of the book respectively focus on scaling up machine learning algorithms pertaining to online learning, and typical practical learning applications that demand customisable distributed platforms. Relevant theoretical as well as empirical results have been synthesised and analysed coherently. Interesting applica-

tions specific to areas such as computer vision and automatic speech recognition have been presented to show that multi-fold gains in terms of computational speed and cost can be achieved with the aid of sophisticated parallelization.

Those audiences who expect more mathematical insight into advanced concepts need to refer additional material besides this volume. While the book could be resourceful for computer engineers and researchers, how far the book can serve the student community and academia is not apparent. To fulfil this genuine expectation more detailed worked out examples are required. Importantly the inclusion of a series of exercises at the end of every chapter and solutions to selected exercises would be much appreciated by the academic community.

Reference

[1] Bekkerman R, Bilenko M, and Langford J, *Scaling up Machine Learning: Parallel and Distributed Approaches*, Cambridge, UK: Cambridge Univ. Press, 2011.

Ibrahim Venkat, PhD
Senior Lecturer
School of Computer Science
University Sains Malaysia

Event: International conference on Neural Information Processing

by *Claire Gerrard* (Robert Gordon Univ., Aberdeen)

The 19th International Conference on Neural Information Processing (ICONIP) took place on the 12–15 November 2012 in Doha, Qatar. ICONIP is an annual conference of the Asia Pacific Neural Network Assembly (APNNA) and one of the premier international conferences on the areas of neural networks. Held since 1994, it provides a forum for researchers, students and industrial professionals worldwide to discuss new ideas, and explore challenges in the field of neural networks and related disciplines. ICONIP 2013 will be held in Korea, in November.

This year, ICONIP received about 700 submissions from over 60 countries; 400 of which were selected for publication in the LNCS Springer series. Each of the five volumes of these proceedings represents one of five topical sections: theoretical analysis, neural modelling, algorithms, applications, as well as simulation and synthesis, covering topics as diverse as bioinformatics, data mining, or molecular and quantum computing.

The program included about 30 keynotes, plenary and invited speeches. The opening keynote was an inspirational talk given by Shunichi Amari entitled "brain, stochastic world, and information geometry". Amari discussed the reasons why the brain so accurately processes information using stochastically fluctuating units. He then introduced the mathematical theory of information geometry, as a useful tool for studying such stochastic processes and discussed its application toward neural spike analysis, dynamical behaviours of learning machines, and sparse signal processing.

Another interesting keynote was by Stephen Grossberg, entitled "Neural dynamics of invariant object learning, attention, recognition and search". Grossberg discussed two new paradigms which describe computational processing within the brain: Complementary Computing and Laminar computing. The talk focused on results that illustrate these two paradigms. For example, the ARTSCAN model sheds light on how multiple brain processes are coordinated to allow the brain to recognise objects from multiple view points of a scene scanned by eye movements.

Plenary speaker Ron Sun gave an interesting talk entitled "The CLARION Cognitive Architecture: Motivation, Personality, and Social Interaction". CLARION is a computational cognitive model used to represent human cognition using a dual process theory of mind approach. It is made of a number of subsystems including the action centred, non-action centred, motivational and meta-cognitive. One of the main differences between CLARION and existing frameworks is that each of these subsystems captures the distinction and interaction between im-

plicit and explicit cognitive approaches.

A popular theme was the development of brain computer interface (BCI) techniques using electroencephalography (EEG). Mohammed Shakir's presentation presented a system using EEG to control a wheelchair. The EEG signal produced by a subject is processed to obtain statistical features that can then be interpreted into commands, using a fuzzy logic classification algorithm. Other applications of EEG signals in BCI techniques ranged from communication aids for sufferers of conditions such as locked-in syndrome, movement of prosthetic limbs or even classification of neurological conditions. Wafaa Khazaal Shams presented a fuzzy model to help diagnose autistic spectrum disorders, using features extracted from EEG to provide a degree of occurrence.

The session on cybersecurity was very popular, containing an invited talk by Daisuke Inoue on identifying emerging security threats with data mining technologies: NICTER (network incident analysis software) analyses real time darknet (accessible and unused IP address space) traffic to identify and predict a range of cyber-attacks, including for example botnets.

Many big issues were addressed in the panel discussions: How do we deal with or efficiently make use of the ever increasing problem of "big data"? Can we guarantee that the accuracy of inference on data increases with the increase of data? Another common problem identified was how abstract is inaccurate. In other words, when does a model become so simplified that it no longer provides a realistic representation of its subject? As one can expect, the nature of intelligence and its application to AI gave rise to a big debate, when performance of humans are compared to that of artificial devices.

Claire Gerrard
PhD Candidate
School of Computer Science
Robert Gordon University, Aberdeen

Event: Distributed Thinking Symposium V

by *Jens Koed Madsen* (University College London)

The workshop took place at Goldsmiths, University of London, on the 30^{th}-31^{st} of January 2013. It was the second instantiation of the AISB workshop series, and it was made possible by funding generously provided by the AISB. Alongside being part of the AISB workshop series, the event was also part of recurring symposia on distributed thinking, hence the title of the workshop. In particular, the event was possible by the help of Mark Bishop (who was the general contact at Goldsmiths) as well as Yasemin Erden and Kent McClymont (who are responsible for organizing the overall AISB workshop series). Without their invaluable help and assistance, the workshop would not have taken place.

The workshop focused on time-scales in systemic thinking–particularly when concerned with language dynamics, which is a theme that promises interesting findings for a variety of disciplines ranging from psychology, linguistics to artificial intelligence (AI) and human-computer interaction (HCI). When performing languaging, there is an interesting time-scale relationship between multiple influences, which affect the linguistic capabilities of interlocutors. These influences exist on different timescales given the fact that some of these influences reaches over centuries. For instance, the general developments of English (in terms of vocabulary, grammar, etc.), your personal relation to language throughout your life, and your cultural upbringing all contribute to the language dynamics in the specific moment that you make use of language. In other words, the exploration of language assumes that language is a dynamic process that, while centred on human interaction, also exploits historically derived resources. In this view, interaction and problem solving are understood, not in relation to normative models, but as sense-saturated, regulatory human activity, which result in systemic thinking and action. The workshop aimed at discussing these interactive and time-scale issues in order to place language theoretically in a dynamic milieu.

The practical extrapolations of such a placement of language reaches into a variety of related disciplines. From a psychological point of view, a greater understanding and appreciation of time-scales in dynamic systemic thinking may lead to interesting discussions concerning the limits of traditional cognitive psychology as centred in the mind of the individual by exploring the influences of interaction par excellence. For linguistics, the theme of the workshop carries natural interest such that an increase of the understanding of the underlying factors, which influence language dynamics and linguistic capability, the better are the theories of language produced by linguists. For AI and HCI the influences are less

obvious and intuitive. However, when constructing artificial mechanisms that make use of some form of systemic thinking or which is supposed to be able to communicate with the human users, a more in-depth understanding of these mechanisms becomes apparent and a natural boon. Thus, in order to create systems that rely on language capabilities, it is essential for artificial intelligence as well as computer scientists to engage with the second-order influences such as the ones discussed at the workshop. This is necessary in order to create systems that are as closely related to human language dynamics as possible, which in turn facilitates HCI and provides a stronger foundation for complex AI systems. In other words, the workshop was of interest to several disciplines, and happily members of the audience were present from all these disciplines (although most of the members of the audience in the end were either psychologists or linguists).

The presentations at the workshop were interesting and tied in well with the general theme. One cancelled talk allowed for a rescheduling of talk so that the participants could go to the Whitehead Lecture series in the afternoon. So, despite a cancellation, the program was met satisfactorily. Generally, the atmosphere was positive, as the mood seemed a mix of inquisitive approaches to the work of others whilst keeping an open mind towards new definitions of complex issues. Indeed, the fact that the theme described in the above is still relatively opaque in theory and application means that people generally were willing to listen to suggestions concerning theoretical definitions, new empirical data, and potential extrapolations of this. In this, the workshop was a great success. Finally, organisers currently consider a publication in form of a collection of articles as a special issue in a journal yet to be selected. However, this is not yet finalized. In all, the AISB workshop was successful and a very enjoyable exchange of thoughts regarding distributed thinking in terms of timescales.

Jens Koed Madsen
PhD Candidate
Cognitive, Perceptual, and Brain Sciences Department
University College London

Event: Sensory Substitution and Augmentation

by *Janet Gibbs* (Kings College London)

For three days before Easter, the worlds of philosophy, psychology, neuroscience and practical rehabilitation came together at the British Academy to explore how far Sensory Substitution and Augmentation (SSA) have enhanced our understanding of perception thus far, and how we might best use them in the future. The conference brought together those who have built and studied devices, whether for rehabilitative use or more academic purposes, along with experimental and theoretical researchers from the various disciplines. Its expressed aim was threefold—to enable designers of SSA devices to share good practice; to inform the design of future devices through a knowledge exchange between theory and practice; and to enhance our understanding of the nature of perception and the senses.

Issues addressed included whether such devices give rise to genuinely new perceptual or sensory experience; and, if so, is such experience characteristic of the substituting or substituted modality? Is it, indeed, perceptual at all or 'merely' cognitive skill? Possible limitations on SSA were considered, such as why devices developed so far have all addressed spatial modalities—visual, tactile and auditory—and whether we are constrained to using 'common sensibles' between the substituted and substituting modalities: will it ever be possible, for example, to substitute by or for taste, which has no spatial dimension? Also, why, after more than 50 years in the field, are SSA devices still so little used outside the research environment? More theoretical speakers considered what perceptual theories are supported or undermined by SSA work, what light is shed on such things as perceptual learning, neural plasticity and the relations between perception and cognition, and whether SSA casts doubt on how we differentiate senses.

Two key threads were apparent throughout the conference. First was the difficulty of drawing meaninful conclusions from such a varied range of studies: from single case reports, through small group studies to controlled trials; most using very different devices, working with different sensory modalities, and with widely differing experimental designs and conditions. It was particularly notable that the most promising results come from single user or small group studies, and from the more subjective reports. Some individual users report experiences characteristic of the modality aimed for only after months or years of immersive use: what, then, can we really conclude from a controlled trial where the user experience is measured in hours, over a period of days or weeks? The second issue that became increasingly apparent was the extent to which users' past perceptual experience influences their experience of a device—as, for example, with late- or early-blind users, partially

sighted, or sighted users with blindfold, in trials of visual substitution.

The conference was followed by an afternoon of SSA system demonstrations—the first time that such a range of devices have been brought together into one place. Participants were invited to experience for themselves a number of different Auditory/Visual devices, a minimal version of TVSS (Tactile/Visual), a Tongue Display Unit and the Tactile/Magnetic Orientation feelSpace belt.

Janet Gibbs
Database Administrator
Dept of Psychological Medicine
Kings College London

Announcements

AISB-50

AISB is proud to announce that next year's convention will commemorate both 50 years since the founding of the AISB Society, and sixty years since the death of Alan Turing, founding father of both Computer Science and Artificial Intelligence. It will be held at Goldsmiths, Univ. London, a central location we hope accessible to all, from the 1st to the 4th of April, 2014. The convention will follow the same overall structure as previous conventions, with parallel symposia lasting for one or two days, and including any type of events of academic benefit: talks, posters, panels, discussions, demonstrations, outreach sessions, etc. Get in touch with us if you are interested in participating, or even simply offering your help! For more information, visit http://aisb.org.uk/events/aisb14

The 21st Century Body Reloaded

The organising committee for the 21st Body Reloaded Symposium is pleased to annouce the opening of registration. The event is scheduled on Friday 8th of November, 09:00–17:30, in the Dept. Anthropology at UCL. Attendance is free and everyone is welcome. Lunch on the day will be provided, and it will be followed by a drinks reception. Please register in advance to help organising catering; see http://c21stbody.doattend.com

RIP Public Lecture

The next Royal Institute of Philosophy Public Lecture will be presented by Dr Michaela Kendall and is entitled: "Scientific Creativity and The Rise of the Entreprenerds".

It will take place on Wednesday 20th November, 2013 (17:15–18:45) in the Senior Common Room, St Mary's University College, Twickenhamn and will be followed by a drinks reception. Dr Michaela Kendall is co-founder of Adelan and Honorary Senior Lecturer at the School of Medicine, University of Southampton. Further information: Dr Stephen Rainey (stephen.rainey@smuc.ac.uk).

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

Dear Aloysius. . .

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

———

Dear Aloysius,

I am an international champion of open data. When that data is labelled "top secret", governments get unhappy, and when they get unhappy they start playing dirty tricks. That's why I'm currently taking refuge from some trumped-up charges, in the Embassy of a sympathetic country. I can't take up the offer of asylum in that country, unfortunately, without leaving my present sanctuary and getting arrested. Can you help me?

Yours, Stir Crazy

Dear Stir Crazy,

To help you start a new life in a safe haven you must have a new identity and plastic surgery. I'm sure the Embassy can arrange that. Your pursuers will anticipate this, however, so, in addition, you need a major diversion while you make good your escape. This is another job for SHYSTER™ (Speech and Holography Yield a Simulation That Emulates Reality). What better diversion than your apparent successful arrival in your target country? Our new mobile SHYSTER™ technology will project your image and your voice so that you appear to be in one place, while actually being in another. We will stage a few brief glimpses of you in the departure lounge at Heathrow, being interviewed by a journalist, whom we will provide. Then the President of your target country will welcome you, or rather your holographic image, as you appear to arrive at its airport. While our security forces are wondering how you could have escaped their clutches, you sneak, heavily disguised, out of the back door of the Embassy, drive to the airport and fly to your asylum.

Yours, Aloysius

———

Dear Aloysius,

I'm a history teacher who was drafted into teaching ICT, on the grounds that I knew a bit about Word and Excel. Now Michael Gove has announced that I have to teach children to programme. I've never programmed a computer in my life and I'm a bit long in the tooth to start now. If you can't help me, I'm afraid I'll have to go for early retirement.

Yours, Obsolete

Dear Obsolete,

The Institute's new BITSTREAM™ (Brain Interface Translates Speculations to Routines; Effected by Automated Mechanisms) system can save you from obsolescence. Our research into automated programming and brain-machine interfaces has reached the point at which a mental rehearsal of your desired procedure processing a few well chosen inputs into outputs will create a FORTRAN program implementing that procedure. [Sorry it's FORTRAN; that language turned out to be the best fit to human neural processes.] The rapidity with which you will be able to write faultless code will impress your young charges. BITSTREAM™ also produces English explanations of the workings of the programs, which you can then incorporate into your lessons.

Yours, Aloysius

Fr. Aloysius Hacker
Cognitive Divinity Programme
Institute of Applied Epistemology

Back matter

Editors – aisbq@aisb.org.uk
Dr David Peebles (*Univ. Huddersfield*)
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. 137, October, 2013

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