AISB Quarterly No.130

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

# The Turing Test

This short note is concerned with the philosophical question—Can a Machine Think? Famously, in his 1950 paper *Computing Machinery and Intelligence* [13], the British mathematician Alan Turing suggested replacing this question—which he found "too meaningless to deserve discussion"—with a simple (behaviourial) test based on an imagined Victorian-esque pastime he entitled the *imitation game.* However, in the fifty plus years since Turing's seminal paper was first published a number of different interpretations of his ideas have emerged and so, to coincide with the Loebner Prize competition held at the University of Reading (UK) on the 12th October 2008, the (British) society for the study of Artificial Intelligence and the Simulation of Behaviour sponsored a one day 'invited-speaker' symposium to present a new, formal, academic critique of issues around the Turing test (This article is based on the author's introduction to the recent Special Issue of the KYBERNETES Journal [1] dedicated to the 2008 AISB Turing symposia on the Turing Test). The day commenced with talks from three eminent speakers (Baroness Greenfield; Michael Wheeler & Selmer Bringsjord) who offered a personal context to, and their perspective on, the Turing test. These presentations were followed in the afternoon session with the invitation to four subsequent speakers (Andrew Hodges; Luciano Floridi; Margaret Boden; Owen Holland) to address specific matters related to the Turing test (e.g. definitional; adequacy; tests in other modalities; technical/computational issues). In funding the event the AISB hoped to pin down current thought on what is the best interpretation of the Turing test and so the day ended with a 'round-table' discussion on this theme and some of the other core issues raised during the day.

As early as 1941 Turing was thinking about machine intelligence [4]—specifically how computing machines could solve problems by searching through the space of possible problem solutions guided by heuristic principles. And in 1947 Turing gave what is perhaps the earliest public lecture on machine intelligence at the Royal Astronomical Society, London. Subsequently, in 1948, following a year's sabbatical at Cambridge, Turing completed a report for the UK's National Physical Laboratory on his research into machine intelligence, entitled Intelligent Machinery [12]. Although not published contemporaneously, the report is notable for predicting many core themes which eventually emerged from the yet nascent science of machine intelligence: Expert Systems; Connectionism; Evolutionary Algorithms; most intriguingly, the report offers perhaps the earliest version of the imitation game / Turing test. Turing presents this original version as follows:

*"The extent to which we regard something as behaving in an intelligent manner is determined as much by our own state of mind and training as by the properties of the object under consideration. If we are able to explain and predict its behaviour or if there seems to be little underlying plan, we have little temptation to imagine intelligence. With the same object therefore it is possible that one man would consider it as intelligent and another would not; the second man would have found out the rules of its behaviour. It is possible to do a little experiment on these lines, even at the present stage of knowledge. It is not diffcult to devise a paper machine which wil l play a not very bad game of chess. Now get three men as subjects for the experiment A, B, and C. A and C are to be rather poor chess players, B is the operator who works the paper machine. (In order that he should be able to work it fairly fast it is advisable that he be both mathematician and chess player.) Two rooms are used with some arrangement for communicating moves, and a game is played between C and either A or the paper machine. C may find it quite diffcult to tell which he is playing. (This is a rather idealized form of an experiment I have actually done.)"*

Some commentators (e.g. Whitby[2]; Shah & Warwick[10]) have suggested that Turing didn't intend the imitation game to be the specification of some fully operational procedure to be performed by future machine intelligence researchers as a yardstick with which to evaluate their wares, but merely as a thought experiment, a "philosophical ice-breaker" (ibid), "attempting to deal with the ill-definition .. of the question .. can machines think?" (Wiggins [17]). The fact that Turing personally enacted this first version of the imitation game offers perhaps partial evidence against this interpretation.

Subsequently, in the initial exposition of the imitation game presented in the 1950 paper [13], Turing called for a human interrogator (C) to hold a conversation with a male and female respondent (A and B) with whom the interrogator could communi-

In this issue:

**Features and Reviews**



**Society**



**Father Hacker**

# The Turing Test (cont.)

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

ISSN 0268-4179





cate only indirectly by typewritten text. The object of this game was for the interrogator to correctly identify the gender of the players (A and B) purely as a result of such textual interactions; what makes the task non-trivial is that (a) the respondents are allowed to lie and (b) the interrogator is allowed to ask questions ranging over the whole gamut of human experience. At first glance it is perhaps mildly surprising that, even after many such textual interactions, a skilled player can determine (more accurately than by chance) the correct gender of the respondents.

Turing's Victorian-esque parlour game describes a scenario perhaps not unfamiliar to that many twenty-first century video gamers encounter when participating in a large multi-user virtual world—such as World of Warcraft or Second Life—where in-game avatars controlled by real-world players may often fail to reflect the gender they claim to be; the controller may be female and the avatar male or vice versa.

Turing then asked the question - What will happen when a machine takes the part of (A) in this game? Would the interrogator decide wrongly as often as when playing the initial imitation game? In this flavour of the imitation game / Turing test - which has become known as the 'standard interpretation' - a suitably programmed computer takes the part of either player (A) or player (B) (i.e. the computer plays as either the man or the woman) and the interrogator (C) simply has to determine which respondent is the human and which is the machine.

Although it is implicit in this 1950 version of the imitation game that the interrogator knows at least one of the respondents is a machine, a subsequent version - presented in a radio discussion in 1952 [14] - describes a 'jury' of interrogators questioning a number of entities seriatim; some entities being computers, some being human. Clearly, during each interrogation in this version of the test, the jury does not know if they are interacting with a human or a machine. Similarly when Colby et al tested PARRY, they did so by assuming that the interrogators did not need to know that one or more of those being interviewed was a computer during the interrogation [3]. Copeland, in commenting on the revised 1952 test [6], argues that the 1950 version is the better, as the single interview mode is open to a "biasing effect which disfavours the machine".

However, a close reading of the 1950 paper reveals several other possible in- terpretations other than the standard version outlined above. For example it is possible to interpret Turing when he says: "We now ask the question, 'What will happen when a machine takes the part of (A) in this game?' Will the interrogator decide wrongly as often when the game is played like this as he does when the game is played between a man and a woman?" as meaning:

(a) literally what he says - that the computer must pretend to be a woman, and the other participant in the game actually is a woman (see Genova [7] and Traiger [11]);

(b) that the computer must pretend to be a woman, and the other participant in the game is a man who must also pretend to be a woman.

Towards the end of section (5) of the 1950 paper [13] Turing, perhaps rather confusingly suggests, [the computer] "can be made to play satisfactorily the part of (A) in the imitation game, the other part being taken by a man".

Although in a very literal sense, the above present valid alternative interpretations of the imitation game, the core of Turing's 1950 article (and material in other articles that Turing wrote at around the same time) strongly support the claim that Turing actually intended the standard interpretation [6,8,9].

In the 1950 paper Turing confidently predicted that by the year 2000 there would be computers with 1G of storage, (which turned out to be a relatively accurate prediction), which would be able to perform the (standard) Turing test such that the average interrogator would not have more than 70% chance of making the right identification after five minutes of questioning; the latter claim being slightly ambiguous:- did Turing intend the imitation game to be played out over five minutes of questioning in total or did he mean five minutes of questioning per respondent? Furthermore, although Turing specifically describes playing the imitation game with 'average' interrogators, some commentators - perhaps remembering Kasparov's titanic series of games against chess playing machines - hint at a 'strong' version of the imitation game; where the interrogator is an expert interrogator, the game is played as an open ended conversation and the test is for full 'human indistinguishability' (cf. Hugh Loebner's specification for the gold medal prize in his version of the Turing test).

In 1990 Hugh Loebner agreed with 'The Cambridge Center for Behavioural Studies' to underwrite a contest designed to implement a Turing-style test. Dr. Loebner pledged (i) a 'Grand Prize' of $100,000

## The Turing Test (cont.)

and a 'solid 18 carat gold medal' for the first computer program whose re- sponses were indistinguishable from a human's and (ii) an annual prize—$3,000 in 2010—and bronze medal to be awarded to the most human-like computer program (i.e. the best entry relative to other entries that year, irrespective of how good it is in an absolute sense).

In 2008 the organisers of the annual Loebner bronze medal Prize elected to put Turing's 1950 prediction to the test in the first - least demanding - manner, by enacting a set of five minute Turing tests for the bronze medal prize; specifically, each interrogator was allowed a total of five minutes to respond to both entities (the human and the computer). As a consequence the expected interaction time with the computer program was just two and a half minutes. However, in 2008 even this minimal 'five minute' claim proved optimistic as Elbot - evaluated as the best computer program in this competition - achieved a maximum deception rate of 25% over two and a half minutes of interaction; still 5% short of the 30% deception rate Turing had predicted in 1950. Nonetheless it seems very likely that in the next few years Turing's predictions for a 'time limited' Turing test will be met; whether that means at that juncture "general educated opinion will have altered so much that one will be able to speak of machines thinking without expecting to be contradicted" (as Turing asserted) is very doubtful as, in the 50 plus years since the paper was first published, the status of the Turing test as a definitive measure of machine intelligence and understanding has been extensively critiqued.

Perhaps the best known criticism of 'a Turing style test of machine understanding' comes from John Searle. In the Chinese Room Argument (CRA) [16] Searle endeavours to show that even if a computer behaved in a manner fully indistinguishable from a human (when answering questions about a simple story) it cannot be said to genuinely understand its responses and hence the computer cannot be said to genuinely think (for recent discussion of the CRA see Preston & Bishop [15]).

In 2008 the AISB sponsored an invited speaker symposium on the Turing test at the University of Reading in the hope of eliciting further clarity in the interpretation of the test, further insight into its implications and further reflection as to its status as a (practical) measure of machine intelligence. In 2010 the AISB hosted a second Turing symposium at the Spring Convention to continue discussion of this question (see http://www.cse.dmu.ac.uk/~aayesh/Turing-TestRevisited/Welcome.html).

Although the breadth and depth of the material presented there clearly illustrate that Turing's imitation game continues to present novel insights into mind and machine and any hope for a 'near final' word on the imitation game remains as far off as ever, it is hoped that readers may find this brief clarification of the test and its central controversies of interest.

### References


[1] Bishop, J.M. (2010) The Imitation Game. In: Bishop, J.M. (ed) (2010) Special Issue on the Turing Test, Kybernetes 39:3, Emerald, UK.

[2] Whitby, B. (1996) The Turing Test: AI's Biggest Blind Alley? In: Millican, P. & Clark, A. (eds) (1996) Machines and Thought: The Legacy of Alan Turing. Mind Association Occasional Series 1, pp. 53-62. Oxford: Oxford University Press.

[3] Colby, K.M., Hilf, F.D., Weber, S. & Kraemer, H.C. (1972) Turing-like indistinguishability tests for the validation of a computer simulation of paranoid processes, Artificial Intelligence 3: 199–221.

[4] Copeland, B.J. & Proudfoot, D. (2005) Turing and the Computer, in Copeland, B.J. (ed) (2005) Alan Turing's Automatic Computing Engine, Oxford: Oxford University Press.

[5] Copeland, B.J. (2000), The Turing Test, Minds and Machines, 10: 519–39.

[6] Copeland, B.J. (ed) (2004) The Essential Turing. Oxford: Clarendon Press: 488.

[7] Genova, J. (1994) Turing's Sexual Guessing Game, Social Epistemology, 8: 313–26.

[8] Moor, J. (2001), The Status and Future of the Turing Test, Minds and Machines, 11: 77-93.

[9] Piccinini, G. (2000), Turing's Rules for the Imitation Game, Minds and Machines, 10: 573-85.

[10] Shah, H & Warwick, K. (2010) Testing Turing's five minute, parallel paired imitation game. In: Bishop, J.M. (ed) (2010) Special Issue on the Turing Test, Kybernetes 39:3, Emerald, UK.

[11] Traiger, S. (2000) Making the Right Identification in the Turing Test, Minds and Machines, 10: 561-572.

[12] Turing, A.M. (1948) Intelligent Machinery. National Physical Laboratory Report, 1948, in Copeland, B.J. (ed) (2004) The Essential Turing, Oxford & New York: Oxford University Press.

[13] Turing, A.M. (1950) Computing Machinery and Intelligence, Mind 59, pp. 433–460.

[14] Turing, A.M., Braithwaite, R., Jefferson, G. & Newman, M. (1952) Can Automatic Cal- culating Machines Be Said To Think? in Copeland, B.J. (ed) (2004) The Essential Turing. Oxford: Clarendon Press: 487–506.

[15] Preston, J. & Bishop, M. (eds) (2002) Views into the Chinese Room: New Essays on Searle and Artificial Intelligence, Oxford & New York: Oxford University Press.

[16] Searle, J. (1981) Minds, Brains, and Programs, Behavioral and Brain Sciences, 3: 417-57.

[17] Wiggins, G. (2007) Personal communication, Geraint Wiggins, Professor of Computational Creativity, Goldsmiths, University of London, UK.



Prof. Mark Bishop
Goldsmiths College
University of London
mas02mb@gold.ac.uk


# Conference Report: AISB 2009 Symposium on New Frontiers in Human-Robot Interaction

Human-robot interaction (HRI) research studies the communication and interaction between humans and robots. Of particular interest to artificial intelligence as a whole, much of HRI looks at interaction with robots having some degree of artificial intelligence. Several interesting projects were presented in the 2009 AISB symposium on New Frontiers in Human-Robot Interaction. This report contains a summary of three projects from the symposium that discuss collaborative control, action prediction, and trust in robots, respectively.

The first project involves collaborative control of a robotic wheelchair as described in the paper by Carlson et al. [1]. The wheelchair has sensors to observe the environment as well as an eye-tracking system to monitor the user's gaze. While a joystick is the primary input for controlling the direction of the wheelchair, eye gaze also gives feedback to the robot. Carlson tested the system with a task involving navigation through a narrow doorway with and without the collaborative controller. Three people participated in the test; none of the participants was informed about how the collaborative controller operates.

The results show that eye movement increased and was more chaotic for collaborative control when compared to full human control. The authors explain the increased eye movement as a result of giving more attention to the navigation task and computer while using collaborative control. One proposed reason for this result is the lack of mental models for how the system works. This idea is supported by evidence from a fourth participant who was informed about the operation of the system prior to the test. The lesson learned from this project is that physiological measures can provide very useful feedback both in the evaluation of a system and as part of controlling a system.

The second project looks at modeling the way humans predict the actions of other humans in a joint task as discussed in the paper by Bicho et al. [2]. Teammates seem to be more effective when they can anticipate the actions of each other and when all team members know the task goals. One current theory is that humans have motor simulation capability to support action prediction. This project uses Dynamic Neural Fields (similar to recurrent neural networks) to predict actions as well as provide an abstraction for goals.

Bicho validated the system with a joint assembly task, where a robot and human assemble a simple toy vehicle. The robot can assemble parts of the vehicle. Using computer vision, the robot can determine where objects are and what piece the human is grasping or reaching for. Based on current subgoals and the particular way a human grasps an object, the robot can infer the human's intent. In addition to requesting parts or giving parts to the human, the robot can detect error in human intent or action. The resulting behavior of the robot seems to work well; the robot seems to know the task well and seems to understand what the human is doing in the context of the task.

The third project, by Desai et al. [3], discusses the role of trust in HRI. Trust is important in automated systems, because the wrong level of trust leads to disuse or overuse of automation. Much research has been done regarding automated systems and trust. The presence of error in automation leads to distrust, even when performance is usually high, although feedback about what caused the error leads to more trust. Such things as good automation etiquette can increase the amount of trust a user has for an automated system.

Desai et al. point out that previous research in automation has rarely involved intelligent robots, and in several cases the lessons learned in general automation do not clearly apply to robots. Historically, automation has been used in factory or process settings (e.g. water treatment), where almost everything is known about the environment. In robotics, little of the world is known, so robots must use imperfect sensors and approximation to observe the world. Simulation studies are more complicated with robots, since the uncertainty of the world must be included in the the simulation to yield reliable results.

Desai et al. also point out that another large difference between traditional automation and robotics is that the level of autonomy for robots can change dynamically. The primary contribution of Desai's paper is that research needs to be done on trust with physical robots, and the paper highlights several potential areas for future work.

In summary, this report discussed three projects from the 2009 AISB conference that showcase interesting research for collaborative control, action prediction, and trust in robots. These projects are examples of another application of artificial intelligence. The future of HRI and artificial intelligence looks exciting indeed.

## References


1. T. Carlson and Y. Demiris, Using visual attention to evaluate collaborative control architectures for human robot interaction, Proceedings of AISB 2009 Symposium on New Frontiers in Human-Robot Interaction, April 8-9, Edinburgh, Scotland.

2. E. Bicho, L. Louro, N. Hipólito, and W. Erlhagen, A dynamic field approach to goal inference and error monitoring for human-robot interaction, Proceedings of AISB 2009 Symposium on New Frontiers in Human-Robot Interaction, April 8-9, Edinburgh, Scotland.

3. M. Desai, K. Stubbs, A. Steinfeld, and H. Yanco, Creating trustworthy robots: lessons and inspirations from automated systems, Proceedings of AISB 2009 Symposium on New Frontiers in Human-Robot Interaction, April 8-9, Edinburgh, Scotland.



J. Alan Atherton
Brigham Young University
tron@byu.edu

# Modeling the Performance of Children on the Attentional Network Test

Recent research in attention indicates ithat it involves three anatomical networks concerned with alerting, orienting and executive control (cf. Posner & Fan, 2007). The Attentional Network Test (ANT) provides a behavioral measure of the efficiencies of these three networks within a single task (Fan, McCandliss, Sommer, Raz & Posner, 2002). Alerting performs the function of achieving and maintaining a vigilant state; orienting refers to selective visual-spatial attention; and executive control involves monitoring and resolving conflict in the presence of conflicting information. Fan and colleagues (2002) designed the Attentional Network Test (ANT) that measures the efficiencies of all three networks in a single behavioral task. ANT is a 30 minute reaction-time based task combining cueing experiments (Posner, 1980) and flanker effects (Eriksen & Eriksen, 1974). ANT-C is the revised version illustrated in figure 1 which is a child friendly version administered to children. The paper (Hussain & Wood, 2009b) that was presented recently at the International Conference of Cognitive Modelling (ICCM) held in Manchester, UK from July 26-29, 2009, adapts an ACT-R 6.0 (Anderson & Lebeire, 1998) model of adult performance on ANT (Hussain & Wood, 2009a) to model the performance of children (aged 6, 7, 8, 9 and 10) on a child-friendly version of the task (Rueda, Fan, McCandliss, Halparin, Gruber, Lercari, Posner, 2004).

The child study (Rueda et al, 2004) reported that latency and accuracy improve over age, up to adulthood. The efficiency of the alerting network is much higher in children up to 9 years with no significant change across age. By age 10 and for adults alerting efficiency significantly reduces. The orienting network seems to be relatively stable up to 10 years with no change. Rate of development of executive control seems to reduce significantly from ages 6 to 7, but after that seems to stabilise up to adulthood with no significant change. Results are similar for 10 year olds and adults on both ANT and ANT-C. This paper compares the human data with model results and statistics show a veridical simulation of the human data as given in table 1.
The modeling work reported in simulates performance of various age groups on a child-friendly version of ANT (ANT-C) projecting the trajectory of development of various attentional networks. The sequence of models simulates the child study findings well. The model fitting process in the light of relevant child development literature helps explain some of the observed age differences: (1) the overall increased latencies are accounted for by slowing down the rule firing times of all productions, which means that children take more time to process in general and tend to make more mistakes; children make more commission errors, the ones due to confusion and distraction (2) alerting network efficiency is slower than that found in healthy adult studies simulated by slowing down

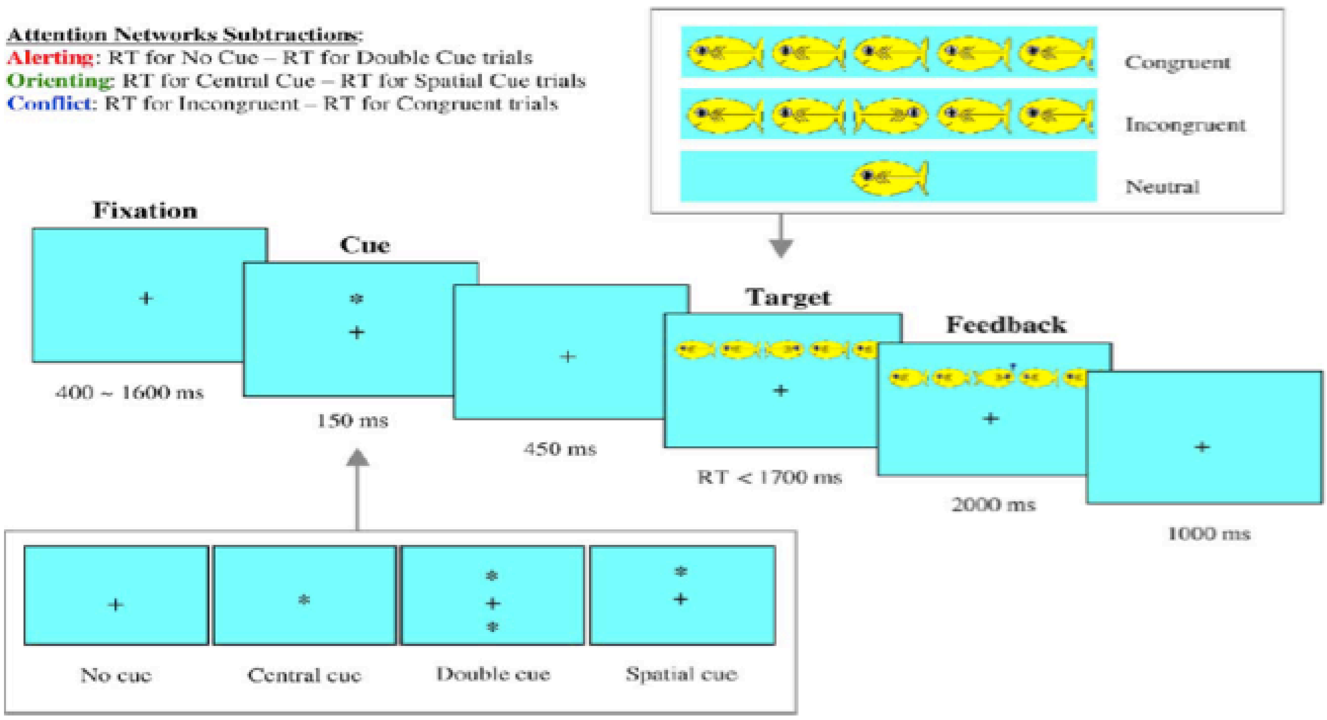


*Figure 1. Child version of the Attentional Network Test (ANT-C), in which yellow fish on a blue background replace flanker arrows in the adult version of ANT-C (Rueda et al., 2004).*

## Attentional Network Test (cont.)

the firing time of the rule which induces an element of “surprise”, so the ability to get alerted in the absence of a signal is slower in children under 10; (3) both orienting network efficiency and the ability to shift from center cue and move to the target location are at adult levels; (4) however, by simulating child performance after introducing an invalid cueing condition, a higher validity effect was found, improving up to age 10. This high validity efficiency was accounted for mainly due to slow disengaging ability, a component of orienting; (5) poor conflict resolution ability in age group 6 is due to a non-optimal refocusing ability when a distractor is selected; and (7) from the model results we conclude there is an inhibiting effect of alerting and facilitating effect of cueing on congruency in children as in adults (Callejas, et al, 2004).

## References


Anderson, J. R. & Lebiere, C. (1998). The atomic components of thought. Mahwah, NJ: Lawrence Erlbaum.

Callejas, A., & Lupianez, J. & Tudela, P, (2004). The three attentional networks: on their independence and interactions. Brain Cognition. 54, 225– 227.

Eriksen BA, Eriksen CW. (1974). Effects of noise letters upon the identification of a target letter in a non search task. Perception and Psychophysics, 16:143-149

Fan J., McCandliss B.D., Sommer T., Raz M. & Posner M.I. (2002).Testing the efficiency and independence of attentional networks. J. of Cog. Neurosc. 3(14):340–47.

Hussain, F. & Wood, S. (2009a). Modeling the Efficiencies and Interactions of Attentional Networks, In L. Paletta & J.K. Tsotsos, Eds. Attention in Cognitive Systems. LNAI 5395, Springer-Verlag, Berlin, Germany.

Hussain, F. & Wood, S. (2009b). Modeling the Performance of Children on the Attentional Network Test. (2009). Proceedings of 9th International Conference on Cognitive Modeling, Manchester, UK.

Posner M.I. & Fan J. (2007). Attention as an organ system. In Neurobiology of Perception and Communication: From Synapse to Society. De Lange Conference IV, Ed. J Pomerantz. London: Cambridge Univ. Press.

Posner M.I. (1980). Orienting of Attention. Quarterly Journal of Experimental Psychcology. 32, 3-25.

Posner, M.I. & Petersen, S.E. (1990). The Attention system of the Human Brain. Ann. Rev of Neuroscience, 13: 25-42.

Rueda, M.R., Fan, J., McCandliss, B.D., Halparin, J.D., Gruber, D.B., Lercari, L.P. & Posner, M.I. (2004). Development of attentional networks in childhood. Neuropsychologia, 42, 1029-1040.



Fehmida Hussain
f.hussain@sussex.ac.uk
Representation and Cognition Group,
Department of Informatics,
University of Sussex, Falmer, Brighton,
BN1 9QJ, UK


| Age | Latency data | | Accuracy data | |
|---|---|---|---|---|
| | *r* | RMSD | *r* | RMSD |
| 6 | 0.79 | 34.7 | 0.93 | 1.28 |
| 7 | 0.92 | 34.4 | 0.86 | 1.02 |
| 8 | 0.88 | 52.5 | -0.11 | 1.24 |
| 9 | 0.93 | 38.3 | 0.58 | 1.15 |
| 10 | 0.93 | 35 | 0.72 | 0.68 |

*Table 1: Correlations and RMSD are used to show statistical fit of the model to the human data for age groups 6-10 years.*

# Conference Report: International Conference on Cognitive Modelling 2009

The 9th International Conference on Cognitive Modeling in Manchester on July 24-26 2009 had the subject of computational models and computation-based theories of human behavior. This conference gathered researchers from many different subfields straddling artificial intelligence, multi-agent systems, psychology and sociology. The day before the main conference, four tutorials were given. These tutorials introduced highly exciting topics - 'agent-based simulation', 'cognitive modeling using neuronal netwoks', 'human-system integration', and 'human learning simulating'. The agent-based simulation talk given by Edmonds and Norling was about a field between sociology and multi-agent systems and considered as a new area for application of cognitive modeling.

The conference was featured by two symposia about the concept of rationality in human behavior and about comparing cognitive models. The symposia have been a very inspiring way of presenting a topic by short sequential short talks of different researchers ending with a long discussion frame. From the given talks in the main track, three interesting projects are (1) 'Comparing Human and Synthetic Group Behaviors: A Model Based on Social Psychology' of N. Fridman and G.A. Kaminka, (2) 'Fluctuations in Alertness and Sustained Attention: Predicting Driver Performance' of Glenn Gunzelmann, L. Richard Moore Jr., Lockheed Martin and Dario Salvucci as much as (3) 'The Wisdom of Crowds in Rank Ordering Problems' of Brent Miller, Pernille Hemmer, Mark Steyvers and Michael Lee.

Fridman and Kaminka (1) presented a way how to connect theories from social psychology with agent-based simulation domain in the concrete case of group/crowd behavior. The work concentrates on implementation and subjective (12 subjects) evaluation of Festinger's Social Comparison Theory (SCT). The implementation is done in the agent-based way of simulation with a visual output. A heuristic based on SCT is compared to an egoistic and a centrally coordinated heuristics. There are similarities to the works from the Agent-Based Simulation Workshop of KI 2008. This topic is also mentioned in the tutorial given by Edmonds and Norling. The result was that agents using the SCT heuristic appear to behave more coordinated way than in the egoistic case but not so coordinated as in the centrally coordinated case. The second work (2) models human performance in driving under condition of sleep deprivation. Fatigue causes a big part of vehicle accidents. The authors constructed an ACT-R model of driver's steering behavior, where they could successfully apply their previous work on mechanisms of fatigue. The driver is simulated by a control update cycle defined in ACT-R. This cycle is responsible for avoiding deviation from the lane. The so called microlapses used for modelling fatigue are breakdowns in the cognitive processing. A higher probability and length of these breakdowns cause a high error rate. The third work (3) is about the phenomenon also known as the 'wisdom of crowds'. The talk was quite interesting - one mentioned that in 91% of cases in the game 'Who Wants to be A Millionaire', the audience was correct. The investigated task was to guess the correct order of items, where the correct one or the ground truth exists. Only 1% of subject could find out the comletely correct answer. The work proved that a Markov chain Monte Carlo model based on the Thurstonian approach for aggregating different answers to a single matches best the ground truth.

One of the new topics emerging is the production of decision sequences in repeated zero-sum matrix games. There are two independent works about this issue - 'Applying Occam's razor to paper (and rock and scissors, too): Why simpler models are sometimes better' by Antonio Napoli and Danilo Fum as much as 'On modelling typical general human behavior in games' by Rustam Tagiew. The first considers the production of decision sequences as independent to the decisions of the opponent but dependent on rewards and models it using ACT-R. The second searches for causalities in decision making using data mining techniques.

Rustam Tagiew
Institute for Computer Science
TU Bergakademie Freiberg
Germany

# The Mining Game
## A Brief Introduction to the Stochastic Diffusion Search Metaheuristic

In recent years studies of social agents have suggested several new metaheuristics for use in search and optimisation; Stochastic Diffusion Search (SDS) [1] is one such 'Swarm Intelligence' algorithm. SDS is a distributed population-based search algorithm utilising interaction between simple agents to locate a global optimum; such 'communicating agents' have recently been suggested as a potential metaphor for some cognitive processes [6].

SDS is most easily applied to discrete search and optimisation problems where the task is to identify the hypothesis, h, which maximises the value of a *decomposable objective function* (A decomposable objective function $F$ is one which can be independently evaluated via a set of partial functions, $f_i$, such that $F = \sum_i f_i$). Unlike many nature inspired search algorithms SDS has a solid mathematical framework which fully describes the behaviour of the algorithm, investigating its: resource allocation [4], convergence to global minima [5], robustness and minimal convergence criteria [2] and time complexity [3]. In the following brief summary we deploy a simple new metaphor—the mining game—to introduce SDS to readers of the AISBQ.

The mining game provides a high-level description of a search to identify the best hill, $H_{best}$, in a large mountain range on which a group of miners should prospect for gold; each hill is quantised into a fixed set of regions, R, where each region yields a specific 'rate of return' $R_j$ of gold (concentration). Thus the 'best' hill for the miners is the hill, $H_i$, which maximises the value of the (decomposable) objective function $F = H_i \sum_j R_j$

**The mining game:** a group of miners learn that there is gold to be found on the hills of a large mountain range but have no information regarding its distribution. To maximize their collective wealth, the maximum number of miners should dig where the concentration of gold is highest; but this information is not available a-priori. Thus the goal of the mining game (i.e. the resource allocation process) is to allocate the most miners to the hill which has the richest seams of gold. In order to solve this problem, the miners employ a simple Stochastic Diffusion Search.

At the start of the mining process each miner is randomly allocated a hill to mine (his hypothesis, $h$). Every day each miner digs a randomly chosen region on the hill, $R_j$. At the end of each day the probability that a miner is 'successful' is proportional to the amount of gold he has mined (the test phase). Each evening the miners congregate and each miner who is not successful selects another miner at random for communication. If the chosen miner has been successful then they share the location of the gold and subsequently both maintain it as their hypothesis, $h$; if not, the first (unsuccessful) miner selects a new region to mine at random (the diffusion phase).

By iterating through test and diffusion phases miners stochastically explore the whole solution space. However, since tests succeed more often on rich seams than on poor, an individual miner will spend more time examining good regions, at the same time recruiting other miners, which in turn recruit even more miners. Candidate solutions are thus identified by concentrations of a substantial population of miners.

Central to the power of SDS is its ability to escape local minima. This is achieved by the probabilistic outcome of the partial hypothesis evaluation in combination with reallocation of resources—miners—via stochastic recruitment mechanisms. Partial hypothesis evalution allows a miner to quickly formulate an 'opinion' on the quality of the investigated solution without exhaustive testing (e.g. a miner forms an opinion on the best hill in the range without exhaustively digging for gold in each region on every hill).

At each iteration of the algorithm the 'success' of the miners in their search for gold can be evaluated probabilistically or deterministically; in the former case each region has a probability of yielding gold, in the latter gold is either present or absent at each region in discrete parcels. In both cases at the end of the test phase miners are either successful or unsuccessful (cf. standard SDS); thus the mining game can be further refined through either of the following two assumptions:

1. Finite resources: the probability of finding gold is reduced each time a miner successfully mines a region;
2. Infinite resources: the probability of finding gold in a region does not vary as gold is extracted from it.

In the first case—finite resources—the task is dynamic and analagous to the robot search tasks described by Steels [8] and Krieger [7]; the second case is analogous to conventional discrete optimisation problems. Ongoing research seeks to apply SDS to the wider gamut of continuous optimisations problems—revisiting the Mining Game metaphor, this is analagous to locating the (real-valued) point(s) on the mountain range where the concentration of gold is highest. In Cognitive Science other recent work offers 'stochastic diffusion' as a potential new metaphor for neuronal operation [6].

### References

[1] J. M. Bishop, Stochastic Searching Networks, Proc. 1st IEE Int. Conf. ANNs, London UK, pp. 329-331, 1989.

[2] D. M. Myatt, J. M. Bishop and S. J. Nasuto, Minimum stable convergence criteria for Stochastic Diffusion Search, Electronics Letters 22(40), pp. 112-113, 2004.

[3] S.J. Nasuto, J.M. Bishop & S. Lauria, Time Complexity of Stochastic Diffusion Search.

## The Mining Game (cont.)

Neural Computation '98, pp. 260-266, Vienna Austria, 1998.

[4] S. J. Nasuto, Resource Allocation Analysis of the Stochastic Diffusion

Search, PhD Thesis, The University of Reading, Reading UK, 1999.

[5] S.J. Nasuto and J. M. Bishop, Convergence Analysis of Stochastic Diffusion

Search, Journal of Parallel Algorithms and Applications 14(2), pp: 89-107,

1999.

[6] S.J. Nasuto, J. M. Bishop, and K. De Meyer, Communicating neurons: a

connectionist spiking neuron implementation of stochastic diffusion search,

Neurocomputing (72), pp. 704-712, 2009.

[7] M. J. Krieger, J. B. Billeter, and L. Keller, Ant-like task allocation and

recruitment in cooperative robots, Nature, vol. 406 (6799), pp. 9925, 2000.

[8] L. Steels, Cooperation between distributed agents through selforganization,

Proc. IEEE Int Workshop on Intelligent Robots and System (7), pp. 814,

1990.

Mohammad Majid al-Rifaie and
Mark Bishop
Department of Computing
Goldsmiths College
University of London
m.majid@gold.ac.uk
mas02mb@gold.ac.uk

# Conference Report:
## 3rd International KES Symposium on Agents and Multi-agent Systems – Technologies and Applications (KES-AMSTA 2009)

Yet another conference on Multiagent Systems? That might be the first reaction when hearing about KES-AMSTA 2009; however this symposium, organized by KES International (the third in its series) and hosted by Uppsala University during June 2009, helped draw interesting conclusions about the state of the art and current research trends of the field.

The most prominent one among these trends was the explicit consideration of business aspects, serving either as an input or an output of agent-based technologies. This fact was encouraged by the symposium organizers, who devoted a good number of sessions on the intersection of Multiagent Systems and Business: E-Commerce, Management and eBusinesss, as well as Digital Economy were a few of the related main tracks and invited sessions. Papers within these sessions but also other generic conference tracks covered a wide range of business life facets, such as marketing, supply chain management, sales, trading, logistics, human resource management, decision making and knowledge management. It is also worth mentioning that numerous presenters came from business universities or management-related research centres, a fact that lead to interesting discussions between experts from the two disciplines.

Two papers within the Digital Economy session dealt with the procurement process, i.e. the process of buying raw materials or product components required for the manufacturing activity. The first paper by Konrad Fuks and Arkadiusz Kawa, titled "Simulation of Resource Acquisition by e-Sourcing Clusters Using NetLogo Environment", utilised agent-based simulation to verify an e-sourcing strategy that recognizes the benefits of coopetition. According to this strategy, companies are encouraged to join e-sourcing clusters where several enterprises, which are normally competing against each other, cooperate in order to achieve lower prices for common resources. Using NetLogo as the simulation environment, the authors showed that e-sourcing clusters can boost up the resource acquisition process. The second paper by Laor Boongasame, Ho-Fung Leung, Veera Boonjing, and Dickson Chiu, titled "Forming Buyer Coalitions with Bundles of Items", illustrated the benefits of buying bundles of items within a buyer coalition. An algorithm for forming such a coalition was introduced, providing total discounts close to the optimal case (according to simulation results).

The area of Supply Chain Management was dealt with in two papers. The paper "Agent Based Decision Support in Manufacturing Supply Chain" by Per Hilletofth, Lauri Lätilä and Olli-Pekka Hilmolaand suggested the use of agent based modelling as an information fusion method to enable decision making with high visibility, thus indicating tasks and information sources for different supply chain agents. The paper I presented with Jessica Chen-Burger on "Analysing Supply Chain Strategies using Knowledge-Based Techniques" aimed to help understand different supply chain scenarios and make business rationale more transparent through a declarative approach. A knowledge based framework for abstracting, analysing and improving supply chain models and a workflow engine for simulating business processes were thus introduced.

A Human Resource Management policy was investigated in the paper "Intelligent Agent Based Workforce Empowerment" by Nazaraf Shah, Edward Tsang, Yossi Borenstein, Raphael Dorne, Anne Liret and Chris Voudouris. The authors assessed the effects of empowering filed engineers in BT using an agent-based model and simulator. Furthermore, Florian Messerschmidt, Andreas Lattner and Ingo Timm in "Customer Assistance Services for Simulated Shopping Scenarios" illustrated how agents representing customers and exhibiting different shopping behaviours may be affected by customer assistance services. Last, Mats Apelkrans and Anne Häkansson presented work on "Applying multi-agent system technique to production planning in order to automate decisions". A multiagent architecture was presented with the vision to automate the information logistics process, and thus the decision making process during production planning. The authors highlighted the importance of security in an actual business application.

After so much interesting discussion about business applications of intelligent agents and multiagent systems, one should ask himself why these technologies are not particularly popular in the business world. That could be one of the questions to be handled in future KES-AMSTA symposia.

Areti Manataki
Centre for Intelligent Systems and their Applications
The University of Edinburgh

# Society News

## AISB Convention 2011

The 2011 AISB Convention will be held at the University of York, from the 4-7th April 2011, organised by Dimitar Kazakov and George Tsoulas. Details will be available on the AISB web site soon.

## Broadening the Scope...

The AISB committee has become concerned that some areas of research in Artificial Intelligence and the Simulation of Behaviour are not well represented by the society. In particular, we are concerned that the society has a tendency to narrow its view to topics related closely to Computer Science at the expense of topics related to the Simulation of Behaviour and areas covered by Cognitive Science, Psychology and other related disciplines. The committee has always played an important role in championing areas of research by raising awareness in committee meetings but, perhaps more importantly, by suggesting articles for AISBQ and symposia for the Convention.

As a step towards widening the breath of material covered by the society, we have decided to publicise the research interests of committee members. These will shortly start appearing alongside the committee names in AISBQ and on the website. There will shortly be a new round of committee elections and we would like to encourage all members of the society to review the current list of committee interests and, if they feel their research area is not represented, to either nominate themselves or a suitable colleague to the committee.

## New Fellows

The committee is pleased to announce that Prof. Luciano Floridi from the University of Hertfordshire and Prof. Mike Wooldridge from the University of Liverpool have been appointed as Fellows of the Society.

## Prof. Richard Gregory

The committee was sorry to hear of the death of Prof. Richard Gregory, Fellow of the Society, on the 17th May 2010.

*University of York*

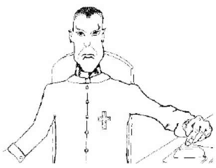

# Dear Aloysius...

Cognitive Divinity
Programme
Institute of Applied
Epistemology

**About the Society**
The Society for the Study of Artificial Intelligence and Simulation of Behaviour (AISB) is the UK's largest and foremost Artificial Intelligence society. It is also one of the oldest-established such organisations in the world.

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

AISB membership includes the following benefits:

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

You can join the AISB online via:

http://www.aisb.org.uk

## Fr. Aloysius Hacker answers your questions

*Dear Aloysius,*

*Both EPSRC and the REF are now demanding that I write accounts of the economic impact of my research. But my speciality is the cognitive modelling of consciousness. While there are obvious long-term commercial applications, I'm unlikely to have anything in the short term to impress these agencies. How can I defend myself against this philistine attack on my curiosity-driven science?*

*Yours, Curious*

Dear Curious,

Your work has immediate commercial application, even before you achieve your ambition to simulate consciousness. I assume your cognitive models are quite large computer programs containing many procedures and variables. Were you aware that there are more than 10 independent companies charging people large sums of money to name a star. How much would these people pay to name the key procedure or main variable in a potentially conscious artificial agent? Here at my Institute, we have developed a service tailor-made to your needs. For a small consideration, **PNOMIC™** (**P**rocedure **N**ame **O**rismology **M**atched to **I**nvesting **C**lient) will act as an intermediary between you and a lucrative market of people seeking scientific immortality.

Yours, Aloysius

Dear Aloysius,

*My last two grant applications have been ranked in the bottom half of the priority list. I've been told that my future applications have been curtailed and that I have to undertake a process of 're-education'. My Head of Department has warned me that I will be classified as non-research-active and given a double teaching load. Can you help me?*

*Yours, Doomed*

Dear Doomed,

Your first step should be to change your name. This will help you evade the funding agency's penalties. You might be able to use the excuse of change of religion or marital status. If not, just do it anyway. If you're lucky, you might even find yourself eligible for a First Grant. Secondly, take advantage of my Institute's **PIMP-UP**™ (**P**roposal **I**mprovement using **M**odish **P**hraseology and **U**nqualified **P**romises). Adventure, risk and transformative technology are the key to unlocking funding treasure; **PIMP-UP**™ will ensure that your future proposals tick all these boxes.

Yours, Aloysius

Dear Aloysius,

I've been a loyal researcher at a top AI lab for 15 years but, in the current economic climate, my lab is being closed down and I'm going to be thrown onto the scrapheap. To make matters worse, my retirement nest egg has been badly eroded by the stock market crash. Can you help?

Yours, Sacked

Dear Sacked,

My research into your employers has revealed that its business is predicated on a huge collection of commercially priceless information. This data can all be stored on a small, portable, terabyte hard drive. If the company's original records were then all deleted, they would be keen to buy your backup for enough money to keep you in luxury for the rest of your life. To assist you, my Institute's **SHAKEDOWN**™ (**S**uccessful **H**eist to **A**cquire **K**nowledge then **E**liminate **D**ata and **O**btain **W**ealth **N**ow) system will automate the whole process for you, including the anonymous emails and the money/data exchange.

Yours, Aloysius

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