Aspects of Innateness and Introspection in Artificial Agents

Chris White, David Bell, Weiru Liu
Queens University Belfast,
School of Computer Science,
Belfast, UK
{cwhite06, da.bell, w.liu} @qub.ac.uk

***Abstract* - Autonomous Computation and cognition, the set of processes that characterise intelligent behaviour are related. For example, cognition drives the autonomy of a robotic agent. In this paper, it is our position that Innateness and Introspection, sometimes referred to here as Instinct and Observation, respectively, are key areas not explicitly focused upon in many current cognitive architectures. By identifying a 'minimalistic' cognition, and incrementally adding Innateness capabilities and Introspection ability, this research defines a structure underpinning a robot explorer that can deal with uncertain environments. The flexibility offered by this structure is considered to be essential for full autonomy. This paper proposes a framework for achieving Innateness and Introspection in autonomous agents and describes briefly two experiments that have shown that a degree of Innateness and Introspection can be achieved, functionally, in robotic agents.**

Keywords - Autonomy; Innateness; Introspection; Machine Intelligence; Cognitive Architecture.

I. INTRODUCTION

Autonomous agents require cognition if they are to understand and adapt flexibly to complex worlds, and so underpin their ability to manage themselves. This is particularly important when the worlds are dynamic, uncertain and impossible to anticipate fully *ab initio.* There is a well-known definition of cognition as:

"…all processes by which the sensory input is transformed, reduced, elaborated, stored, recovered, and used..", [1].

To get the full benefits of cognition in agents, it is desirable to identify what the potential benefits are and to study cognitive processes that are manifested in humans and other higher mammals. These can be incorporated in a cognitive architecture, and they offer fundamental capabilities that are needed for practical application in AC (Autonomous Computing) systems, such as those for surveillance systems, animal behaviour modelling systems, robots, software agent systems, and other infrastructural systems for computing.

In an AC system, these capabilities should be amenable to distribution across the whole system, but, at the other extreme, there are situations where a centralised agent covering a cluster of relatively fixed low-level system elements, rather than one monolithic 'self', is appropriate. The structures and applicability of the alternative dispersion patterns is an area for further investigation. Our ideas cover the distributed scenario, but we focus on a relatively centralised situation as a starting point, and leave the distribution aspects for future elaboration.

Many functional elements are needed for autonomy in the general case. We argue that two specific capabilities, Innateness and Introspection, are needed for full cognition, and are often overlooked completely or under-cooked by researchers in AI (Artificial Intelligence) who offer cognitive architectures that could be considered for adoption in AC systems. A list of 'normative' capabilities that would be considered as 'cognitive' has been fairly well established in the literature. These are usually included in published cognitive reference models or cognitive architectures, and we believe that they should be considered by anyone seeking to produce AC systems that can deal flexibly and effectively with streams of signals and other inputs from their worlds. They include memory handling, various (somewhat subjectively selected) functions required for 'intelligent behaviour', and means of interacting with the environment. We claim that they must be supplemented.

For illustration of our approach in this introduction we use the relatively comprehensive LRMB (Layered Reference Model of the Brain) [2] for cognitive systems which, has been put forward specifically for use in AC systems. This model can be used either to help in explaining fundamental 'natural' cognitive mechanisms and processes [3 – 4], or to simply gather together specifications of capabilities that could be useful when engineering artefacts for various activities. Our focus here is on the latter use.

By common consent there are a lot of cognitive processes in 'natural intelligence'. The LRMB designers list 39 of these at six layers known as the sensation, memory, perception, action, metacognitive, and higher cognitive layers.

The designers of LRMB and other researchers taken collectively have produced a close-to-exhaustive list of features that should be possessed by any cognitive agent. There are various cognitive architectures [5] that have been mooted to capture the basis of cognition. They are intended to specify domain independent infrastructures for intelligent systems.

Starting from such suggestions, we desire to identify and understand some particular mechanisms and interactions for use in complex processes such as those requiring AC. However we find that they do not adequately cover the features on which we focus, i.e., on the Innateness relationships and interactions between what are called in LRMB the 'inherited and the acquired' cognitive functions, as well as Introspective processes that support deep reasoning. These features tend to be understudied in other reference models. For example according to Wang, in LRMB, these capabilities are relatively simple: the analogue of the operating system and applications in a computing system, particularly a real-time system and there is little mention of isolating a way of 'keeping an eye on' what is transpiring during cognitive activity, as a significant subsystem.

In this paper, our hypothesis is that consideration of cognitive architectures can greatly enhance AC systems, and we outline two novel aspects of a cognitive architecture that builds on previous work in this area, and make it particularly well suited to AC. This paper is therefore laid out as follows: we discuss the useful reference point or baseline of minimal cognition briefly, identifying the most basic mechanisms needed if a system is to be called cognitive. After this we define a general cognitive structure where we look at representative architectures in the literature and discuss the Innateness and Introspection functionality. We argue that these features are essential for true cognition. Basic learning methods to be applied in this research are considered and briefly demonstrated in an example of an exercise involving instincts. Finally, the paper discusses implications for measuring autonomy.

II. Minimal Cognition

Simple devices, which do not have functions that are commonly associated with cognition, such as reasoning and learning, can produce seemingly complex behaviour – using a stimulus response mechanism. Such mechanisms have been investigated by Konrad Lorenz [6], considered to be the father of Ethology. They have been termed 'fixed action patterns' or alternatively 'innate release mechanisms' where specific stimuli will trigger a fixed response. These events may appear intelligent but in fact are simple pre-programmed codes that do not deviate from pattern. An engaging example is given by Sharkey [7] where apparently conscious and smart behaviour can be observed with simple reaction mechanisms. The observed behaviour can be interpreted in many anthropomorphic ways, but the operation of the device is very simply explained by these limited, somewhat 'brainless' responses to sense data.

There are many important questions arising from thought experiments like this and related studies. We focus on one: When does a system become intelligent? However the *gedanken* above suggests a more fundamental preliminary question – one that we are actively studying in our labs. What is the minimal structure that is needed if we are to have a cognitive system? An alternative to the definition mooted earlier is to say that agents 'that reason act, perceive, and learn in changing, incompletely known, and unpredictable environments' are cognitive. The device in Sharkey's thought experiment clearly does not qualify.

Now, a two-component signal transduction (TCST) system [8] a molecular sensorimotor system in bacteria, has been discovered relatively recently, and this is seen by some to mark a boundary for cognition. The TCST system is important because it elucidates a molecular mechanism for adaptation and memory. Its sensorimotor organization is still dependent on metabolic activity, but it is 'organizationally autonomous', and functionality similar to that of the nervous system is claimed for it.

We claim that a minimally cognitive system must have two particular features before it can be called cognitive – Innateness and a degree of self-awareness, or Introspection. These features are discussed later.

III. Generalised Cognitive Architecture

As stated above, in research in this area at least two 'flavours' are identifiable…. 'Modelling invariant aspects of human cognition' – explaining/matching psychological phenomena; and 'an effective path toward building intelligent agents' – generating intelligent behaviour

Baars with the GWT (Global Workplace Theory) [9] and Moreno et al with CERA (Consciousness and Emotional Reasoning Architecture) [10] give examples of the types of architectural frameworks that are needed for a computational model of consciousness.

The first of these mainly deals with what Moreno calls A-Consciousness – accessibility of contents of memory for reasoning volition and speech. It seems to aim primarily at understanding consciousness in organisms. For CERA, Moreno includes inner perception or introspection (M-consciousness) and self-recognition and reasoning about the self (S-Consciousness) in his Reasoning consciousness. The functioning of artefacts based on CERA is important; however a goal of that model that we do not aim for is to get close to computational correlates of biological neural structures. Another architecture that is very well developed is ACT-R [11], and it presents a comprehensive list of functionality that would largely be agreed by any researcher working in this domain: Sense functions for visual and other sense processing; Motor functions for action; Memory functions for, e.g., short-term buffers and a long-term memory.

In such architectures, 'soft computing' functions are also needed, as are intentional functions for goals etc., along with a coordinator. The main components of 'Soft Computing' [12] (in Zadeh's SC Institute UC Berkeley) are:

- Fuzzy logic (FL)
- Neural network theory (NN)
- Probabilistic reasoning (PR)

Our emphasis is on exploration. We seek integrated systems for intelligent 'agental' behaviour, rather than piece-wise improvement of individual functions/modules. We want to accommodate the following generalised functionality: Perception and Action (motor) - outer stratum of Fig 1; Reasoning/Predicting/Deciding/Learning – second stratum, Soft (Computing) Functions', of Fig 1; Remembering/Learning - short or long memories STM (Short Term Memory) and LTM (Long Term Memory), in Fig 1. The general architectural framework for cognition presented in outline in Fig 1 gives the 'shape' of a sort of consensus of many contributors to this topic. In addition to showing the 'common consent' framework, we use it to distinguish our own architectural focus.

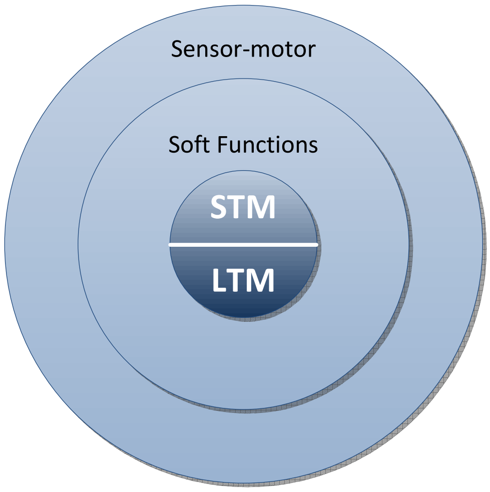

Figure 1. Standard Schema for a generalised Cognitive Architecture.

IV. Learning, Reasoning, and Uncertainty and Inconsistency Handling Mechanisms

Our research is approached from the direction of data and knowledge engineering. Our particular interest is in the areas of inconsistency handling and to an extent, machine learning. Both of these areas should arguably be key focuses for AC. A variety of projects have led to development of and usefulness of artefacts with this sort of functionality e.g., Medical, Transportation, Education, General Engineering, Telecoms, Music, Manufacturing, and Biology. Important topics include transitional (esp. time series) mining, causality, categorisation, reinforcement learning, and harmonisation issues, especially those connected with the obtaining of new knowledge.

The latter topic is of particular interest here – it is the primary catalyst which triggered the present proposals. As indicated above, it is all treated very much from an engineering viewpoint. It has led quite naturally to the contemplation of the possibility and practicability of adding a self-conscious aspect of agents to support 'deep' reasoning and thereby facilitate autonomous behaviour and AC.

To manage the changes of agent's beliefs, we need to consider ways of revising or updating an agent's current beliefs when new knowledge/evidence is obtained. To achieve this, a success principle must be maintained which states that new knowledge should be retained, and the minimal change principle is crucial. It argues that the agent's prior knowledge should also be retained as much as possible while maintaining consistency. When the new knowledge is not guaranteed to be kept, merging has to take place to determine a new belief set based on the strengths of the prior beliefs and new evidence. There is much research [13 – 20] on various aspects of 'Soft Computing' relevant to this work on revising and measuring the amount of conflict and agreement between prioritized knowledge bases, and resolving the conflicts in such knowledge bases e.g., by lexicographic aggregation, combining them by negotiation, or merging them under various constraints.

Revision in numerical related theories, such as probability theory is handled differently. In Probability theory, revision is done by Bayes' updating rule or Jeffery's rule. In other theories, such as the Dempster-Shafer theory and possibility theory, the counterparts of Jeffery's rule or Bayesian updating rule have been developed. The key idea of belief revision and merging in an agent environment is to accommodate new knowledge and to reach a consistent set of ne beliefs for the agent. In machine learning, methods are available for 'Rough' computations, discovering causal patterns in data through mining, reinforcement learning and feature subset selection based on relevance.

V. Specific Cognitive Architecture

The standard model of Fig 1 is supplemented by two important ingredients in our scheme [21] – functions and memory modules for each of two capabilities – viz: Innateness and Introspection ("Observer").

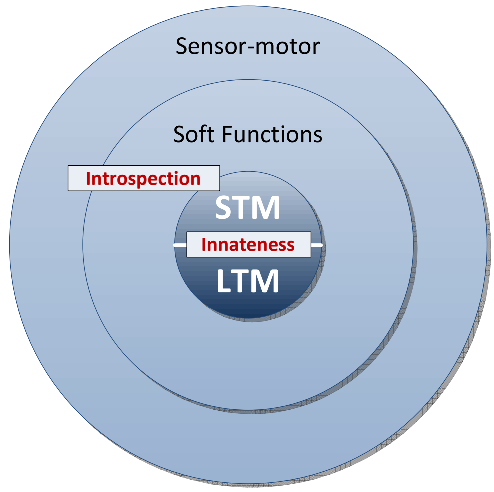

Figure 2. Cognitive Architecture Schema.

A. *Innateness*

A key factor in many programmes where flexibility is required of agents is what is built-in *ab initio*. Instinct is innate ability of agent to detect/react to/associate stimuli from the environment or from internal urges. By it an agent can begin, for example, to form de facto categories [22] ("*We are sensorimotor systems who learn to sort and manipulate the world according to the kinds of things in it, and based on what sensorimotor features our brains can detect and use to do so.*"), and thus learning. It can then use the categories and other learning outcomes to plan and predict, and to some extent modify the built-in innate reactions. Some behaviour can still be explained as innate, but the agent can learn and solve problems related to 'goals' – maybe, in some organisms, even 'let its hypotheses die in its place'.

 ISBN: 978-1-61208-134-2

In an attack-flee scenario in one of our projects which we conduct using Khepera platform [23], two innate instincts are: Investigate (E) / Beware (B). Both of these trigger the collection of sense data (S) from which the robot learns. When a new object (e.g., a light) appears – both E and B lead to S and dominate when enough data is available, learning is possible, and this takes place systematically.

The decision resulting from a particular episode of exploration indicates which of E or B results – essentially whether the robot investigates further (E) or flees (B). This result depends on the evidence acquired according to 2 rules (illustrations are given in Rules 1 and 2 below) – from sense data.

Rule 1 - If light then B and S (strength .90) (an 'innate' rule) ….this says that if a light is detected, the robot has an instinct to be careful, but to passively collect data on the light's behaviour.

Rule 2 - If test is positive then E and S (a learned/acquired rule) ….this says that if the robot tries some test, such as: approach the light and look to see if object reacts aggressively, in which case the robot is put off (score -0.5 if yes); alternatively if the light does not react, score 1, and the robot is positively prompted to investigate further. Nothing is added otherwise.

There are two phases in episodes captured as behavioural traces. The first is illustrated in the table below – the result depends simply on the rules (instincts) and the sense data. At the end of this phase, when some termination criterion is reached, the balance of evidence lies in some particular direction e.g., Beware (as here) with certain strength (such as 0.90), probably indicating that the object is dangerous. The termination criterion for the tabled data is when the accumulated 'score' exceeds 5.

TABLE I. TRACE OF ACTION: STEPS IN A BEHAVIOURAL EPISODE

Agent Actions	**Object Actions**		
Approach	***Attack***	***React***	***Score Added***
1	0	0	1
1	0	0	1
1	0	1	-0.5
1	0	0	1
1	0	0	1
1	0	1	-0.5
1	0	0	1
1	0	1	-0.5
1	0	0	1

Just one piece of evidence is considered here – there could be others. Evidential reasoning can then be used to come to a decision. It is important to distinguish situations where this does not necessitate a particular decision - e.g., strength is only 0.90 - especially one which is against the purpose of the robot. This is in contrast to a more general purpose – such as to 'scientifically' reflect 'the real world'. Such a decision should be taken as a suggestion, but not a necessitation [24]. In the second phase, not illustrated here, further episodes could be conducted to get more persistent knowledge.

At the end of the episode in Table 1, we are at the point in the trace where the accumulated score exceeds 5 (so B dominates here). This pattern of robot plasticity or flexibility of behaviour could be applied in many other AC application areas such as surveillance, software agent systems and other complex computational systems. There is no requirement for self-awareness or Introspection in this example. We add the observer functionality/module as below.

B. Introspection

In previous work [25 – 26], the main theme of the simulation was to show grounding of real world experiences in an agent's sense data and hence some aspects of the common intuition of "understanding". The mapping of this sense data to internal mental constructs, categorising experience and patterns, and most importantly to justify actions, is considered to be a basis for understanding.

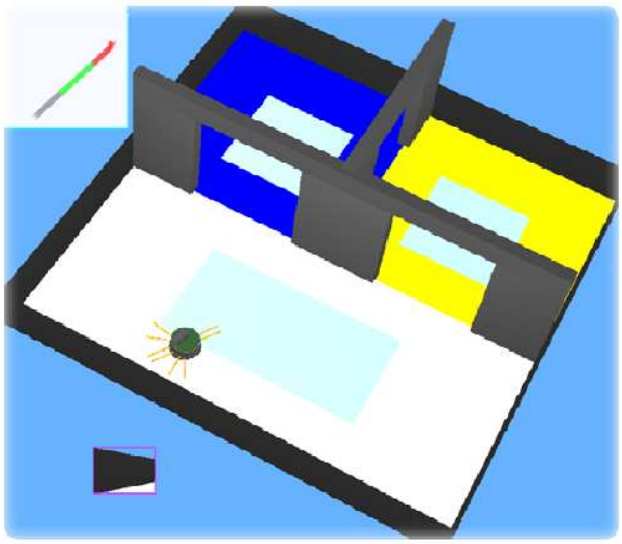

Figure 3. A simulated experiment in Webots, 3 rooms representing different activities an agent can engage and a colour trace top left that can be linked by an 'Observer' to a mood.

It is important to consider what understanding is. For present purposes it is the ability to rationalise ideas through abstractions in order to form a concept, and later to justify any resulting actions, in a given environment. A key consideration here is to ground concepts in sense data. In 'playback' systems the agent carries out actions which are pre-programmed, with no initial understanding of what it is doing. In our scenarios this is their situation initially and actions can be considered here as primitive "instincts". Concurrently there is a "mind" (Observer) linking these actions to an internal construct by abstraction. By doing this it can be said that body and mind have been separated function-wise to achieve this understanding.

Here we suggest that a level of understanding can be achieved by means of symbolic grounding. Two cognitive sub-agents called actor and observer are implemented to allow deep reasoning/query response in simple environments. The actor subagent pursues a programme involving engaging in activities and changing mood in a manner depending on the activities and their ordering. The observer sub-agent records the trace of activities and moods and links the activities and moods to its own sense data using a simple symbolic grounding function. For example, the actor from time to time enters rooms while leaving a colour trail. This colour trail changes depending on the room. The actor can be interrupted from time to time and this is recorded to see what the relationship it has with mood and activity pattern. Colours are linked to moods and rooms to activities. As the actor 'experiences' a particular path, the observer is able to classify, reason and explain what the actor is doing and experiencing in terms of sense data and summaries thereof. See Fig 3.

The environment is a working space which can be varied – e.g., from very simple rooms up to rather complex labyrinths, with various complexities of event mixes. The agent can pick up attributes of the environment form sensors and with the help of a built-in observer, answer factual (related to direct sense data values) and explanatory (related to summaries, etc. of sense data) questions about what it experiences. An example of patterns /characteristics of an agent are : while the activity is reading the mood is initially good, musing turns it to fair and discussion has a bad effect on mood.

This grounding can again extend the flexibility of artefacts such as the common AC applications, this time by relating 'understanding' to the basic inputs from the agent's environment.

VI. TOWARDS THE MEASUREMENT OF AUTONOMY

From an engineering point of view, it is important to be able to compare AC systems with respect to their degrees of autonomy. Like the closely related topic of intelligence tests, this is as yet not well understood and part of our work is to look at this issue. We do so from various angles here.

Keedwell [27] presents an initial proposal of staged testing for Intelligence, known as the Staged Developmental Machine, based on staged testing methods for developing children, per Piaget's theory. An analogous staged test device can be envisaged for AC systems. This proposed test offers a scalar value measure rather than the Yes/No of the Turing test, and there is no requirement for Natural Language Processing. The idea is that 'machines could be judged by their effective stage'. For use in testing for autonomy levels here, we highlight the following stages/part stages, per Keedwell, that we expect our artefacts to attain...

Stage 1 Sensory Perception

- Reacts to Basic Stimuli
- Understands Cause and Effect – predicts next step…
- Understands concept of objects (those controllable and those not) and what to expect from them.
- Uses trial and error to learn about the World (experimentation).

Stage 2 Pre-Operational

- Responding to (NOT Language understanding) relating to self (e.g., answering questions on state).
- Relating objects (though NOT via language) though not in current perceptual field (memory).

Stage 3 Concrete Operation

- Conservation of volume etc. of objects(e.g., estimating quantity of objects in visual field)
- Classification of objects logical rather than on attribute basis (animals/shapes…)
- Sorting objects (e.g., size or colour)
- Effect of Reverse of action (undoing) predictable

Stage 4 Formal Operations

- Ability to create hypotheses/experiments
- Abstract thought – prediction of interactions of objects in novel ways

We also propose [21] a staged approach to the measurement of degree of intelligence or autonomy via a similar scale based on complexity of the environment similar to the Sphex test [28], as well as a degree of Innateness and a degree of Introspection.

VII. SUMMARY AND CONCLUSION

This research is part of a project which focuses on developing a robot explorer. A more generally applicable, unique cognitive architecture has been proposed that can underpin wider AC functionality. It is expected that most autonomous agents at some point will encounter uncertain 'territory'.

Two key and somewhat distinctive features – Innateness and Introspection - are included in the architecture. These are considered to be indispensible if the flexibility and adaptability required for applications requiring autonomy is to be supported. We will use a probabilistic/possibilistic approach combined with classification methods to establish (inexact) rules and tools for AC. We will develop novel methods of evaluation for the added functionality. In particular, as an exemplar of dynamic application systems where the additional capabilities described here are targeted, a robotic agent will be able to explore and describe environments effectively.

REFERENCES

[1] Neisser, U. *Cognitive Psychology*. New York: Appleton-Century-Crofts, 1967. Print.

[2] Wang, Y. 2007, "Toward Theoretical Foundations of Autonomic Computing", *International Journal of Cognitive Informatics and Natural Intelligence*, vol. 1, no. 3, pp. 1-16.

[3] Langley, P., Laird, J.E., and Rogers, S. 2009, "Cognitive architectures: Research issues and challenges", *Cognitive Systems Research,* vol. 10, no. 2, pp. 141-160.

[4] Langley, P. and Choi, D. 2006, "A unified cognitive architecture for physical agents", *proceedings of the 21st national conference on Artificial intelligence - Volume 2,* AAAI Press, pp. 1469-1474.

[5] Newell, A. *Unified Theories of Cognition*. Cambridge, MA: Harvard UP, 1990. Print.

[6] R.W. Burkhardt Jr., Konrad Lorenz, In: Michael D. Breed and Janice Moore, Editor(s)-in-Chief, Encyclopaedia of Animal Behavior, Academic Press, Oxford, 2010, pp. 298-303.
[7] Sharkey, N.E. and Heemskerk, J.N.H. 1996, "The Neural Mind and the Robot", *Neural Network Perspectives on Cognition and Adaptive Robotics* IOP Press, pp. 169-194.
[8] Van Duijn, M., Keijzer, F., and Franken, D. 2006, "Principles of Minimal Cognition: Casting Cognition as Sensorimotor Coordination", *Adaptive Behavior - Animals, Animats, Software Agents, Robots, Adaptive Systems*, vol. 14, no. 2, pp. 157-170.
[9] Baars, B.J. In the Theatre of Consciousness; Global Workspace Theory, A Rigorous Scientific Theory of Consciousness, *Journal of Consciousness Studies*, **4** (4), 1997, pp. 292-309.
[10] Arrabales, M. R. and Sanchis de Miguel, A. 2008, "Applying machine consciousness models in autonomous situated agents", *Pattern Recognition Letters*, vol. 29, no. 8, pp. 1033-1038.
[11] "About ACT-R". Act-R: Theory and Architecture of Cognition. 2011. March 24, 2011, http://act-r.psy.cmu.edu/about/
[12] "BISC Program; Soft Computing". The Berkeley Initiative in Soft Computing. 2009. March 24, 2011, http://www.eecs.berkeley.edu/~zadeh/
[13] Qi, G., Liu, W., Glass, D.H., and Bell, D.A. 2006, "A split-combination approach to merging knowledge bases in possibilistic logic", *Annals of Mathematics and Artificial Intelligence*, vol. 48, no. 1-2, pp. 45-84.
[14] Qi, G., Liu, W., and Bell, D.A. 2006, Merging stratified knowledge bases under constraints, *Proceedings of the 21st American National Conference on Artificial Intelligence*, (AAAI06): pp. 281-286.
[15] Qi, G., Liu, W. and Bell, D. 2010, "Measuring conflict and agreement between two prioritized knowledge bases in possibilistic logic", *Fuzzy Sets and Systems*, vol. 161, no. 14, pp. 1906-1925.
[16] Bell, D.A., Guan, J.W., and Lee, S.K. 1996, "Generalized union and project operations for pooling uncertain and imprecise information", *Data & Knowledge Engineering*, vol. 18, no. 2, pp. 89-117.
[17] Guan, J.W., Guan, Z., and Bell, D.A. 1997, "Bayesian probability on boolean algebras and applications to decision theory", *Information Sciences*, vol. 97, no. 3-4, pp. 267-292.
[18] Bell, D.A. and Guan, J.W. 1998, "Computational methods for rough classification and discovery", *Journal of the American Society for Information Science*, vol. 49, no. 5, pp. 403-414.
[19] Wang, H., Bell, D., and Murtagh, F. 1999, "Axiomatic Approach to Feature Subset Selection Based on Relevance", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 21, no. 3, pp. 271-277.
[20] Guan, J.W. and Bell, D.A. 1993, "A generalization of the dempster-shafer theory", *Proceedings of the 13th international joint conference on Artificial intelligence - Volume 1* Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, pp. 592-597.
[21] Bell, D. 2011, "The CogArch Framework", Unpublished paper.
[22] Harnad, S. "To Cognize is to Categorize: Cognition is Categorization". Categorization in Psychology. 2003. June 22, 2010, Proc UQAM Summer Institute in Cognitive Categorisation, http://www.ecs.soton.ac.uk/~harnad/Temp/catconf.html
[23] "Khepera Platform". K-Team Corporation | Mobile Robotics. 2011. March 24, 2011, http://www.k-team.com/
[24] An, Z., McLeish, M., Bell, D.A., and Hughes J, (1993) "How Did the Tiger Rumble the Donkey?" Proc 6th FLAIRS Symposium, Florida, pp136-141.
[25] Jin, Z. and Bell, D.A. 2003, "An experiment for showing some kind of artificial understanding", *Expert Systems*, vol. 20, no. 2, pp. 100-107.
[26] White, C. and Bell, D. 2010, "Actions and observations: Demonstrating aspects of understanding in a simple world", *Towards a Comprehensive Intelligence Test (TCIT): Reconsidering the Turing Test for the 21st Century Symposium*, ed. Ayesh, Bishop, Floridi, Warwick, April 2010, pp. 24-27.
[27] Keedwell, E. 2010, "Towards a Staged Developmental Intelligence Test for Machines", *Towards a Comprehensive Intelligence Test (TCIT): Reconsidering the Turing Test for the 21st Century Symposium*, ed. Ayesh, Bishop, Floridi, Warwick, April 2010, pp. 28-32.
[28] Esperjo-Serna J.C. 2010, "Connecting the Dots my own Way: Sphex-test and Flexibility in Artificial Cognitive Agents", *Towards a Comprehensive Intelligence Test (TCIT): Reconsidering the Turing Test for the 21st Century Symposium*, ed. Ayesh, Bishop, Floridi, Warwick, April 2010, pp. 1-6.