Achieving Human Level Reasoning and Decision-Making for Autonomous Systems: An Agent's Perspective

Sitanath Biswas[1] ,Trilok N. Pandey[2],Sarada P. Pati[3]
Institute Of Technical Education and Research
Siksha 'O' Anusandhan University,
Khandagiri Square, Bhubaneswar-751030, Orissa, India

Abstract

AI researchers have several overlapping objectives. Among these are: to build systems that aid humans in intellectual tasks; to build agents that can function autonomously in circumscribed domains; to build a general science of intelligence as manifested in animals, humans, and machines; and to build versatile agents with human-level intelligence or beyond. In this paper, we list what we think are some important considerations for those working toward building human-level autonomous AI agents.

Keywords: *Autonomous agent, AI, Human Level Reasoning, Decision Making.*

1. Introduction

In many subfields of artificial intelligence during the past several decades, there has been substantial progress that has resulted in significant near-term advances in theory and applications. However, we believe that progress towards human-level artificial intelligence and the applications it enables requires a deeper and more comprehensive understanding that cannot be achieved by studying individual areas in isolation. Two reasons, both involving integration, support this belief. First, many problems that human-level AIs must solve involve sub-problems currently addressed by different subfields, often using very different computational methods. A human-level AI must either integrate, for example, backtracking search, partially observable Markov decision processes (POMDPs), logic theorem proving algorithms, productions systems, and neural networks, or it must be based on new, heretofore undiscovered computational methods that exhibit all of the best features of these computational methods. Second, the best currently existing example of human-level intelligence is of course the human being. We believe that the history of AI demonstrate that insights into the mechanisms underlying human cognition can inform research towards human-level AI.

We propose the building of following characteristics in an Autonomous Agent for achieving Human Level Reasoning and Decision Making:

1) *Compassionate Intelligence* in an agent results in exhibition of human-level compassion or emotional intelligence.
2) *Affective Inference* in some form is essential in an agent as we approach the goal of creating human-level AI or face the challenge of applications requiring social and emotional intelligence.
3) *Belief change* in an agent is necessary because the world has been changed, and/or the agent has made a new observation of the static world.
4) *Human Reasoning* is needed to be exhibited by the agent in different "mental realms", as we discuss in this paper, for the agent to be able to reason like human-beings.
5) *Intuition* primarily capitalizes on prior knowledge acquired via a slow learning system rather than on recently encountered information kept in short-term memory.
6) *Decision making*, rationally, like a human, depends on what one believes, what one desires, and what one knows. In conventional decision models in agents, beliefs are represented by probabilities and desires are represented by utilities. Software agents are knowledgeable entities capable of managing their own set of beliefs and desires, and they can decide upon the next operation to execute autonomously.

Agent based models of human systems, such as those found in behavioural economics, implement their agents intelligence through the use of simple rules. These models have proved to be worthwhile research tools. However, improving the outputs of these models will require a greater emphasis on the development of intelligence within the model agents. A deeper understanding of the cognitive process of human decision-making is also needed. In this paper a conceptual framework for the process of achieving human level decision making and reasoning in an autonomous agent is put forward.

ISSN (Online): 1694-0814
www.IJCSI.org

2. Compassionate intelligence

The heart of what it means to be both human and intelligent includes compassion and empathy, social and emotional common sense, as well as more traditional methods of AI suitable to tasks. A program capable of empathetic decision-making or compassionate social interaction requires some meta-cognition as part of the bounded informatic situation. Namely, the cognition of such an agent includes thinking about thinking, thinking about feeling, and thinking about thoughts and feelings – its own and/or those of other agents. Our position is that human-level AI programs must not only reason with common sense about the world, but also about irrationally and with feeling, because every human being knows that to succeed in this world, logic is not enough.
An agent must have *compassionate intelligence*.

Compassionate Intelligence is the capacity of an agent to act in compassionate manner in a bounded informatic situation. Many features of human-level compassion or emotional intelligence will be wanted in an artificial agent, some will not. Thinking about designing an AI system without view of an application or without view of a particular philosophy of mind can lead to the discovery of new approaches to some of the problems of AI.

A. *Mind Training From India and AI*

Recently ancient mind-training practices involving meta-cognition have become very popular in western cultures. Recently ancient mind-training practices involving meta-cognition have become very popular in western cultures. FMRI and other diagnostic tools have shown that when humans engage in persistent mind training involving meta-processes there is a positive effect on mental and physical health as well as permanent changes in brain structure and function [Begley 2007] [Lutz et. al. 2004]. A dramatic example of this idea is the practice of TUMMO [Crommie 2002], [Benson, 1982]. But how do human mind training practices like Vipassana(meaning "insight meditation" in Sanskrit) relate to AI systems? Natural systems of cognition have always been inspirational to AI researchers (e.g. vision, memory, locomotion.) Cultures where human mind training has evolved for hundreds and thousands of years present an untapped resource of ideas for researchers working towards human-level AI or compassionate intelligence. Many of these special mind training methods use an architecture of mind that is based on meta-cognition and meta-processes similar in structure and function to the diagram developed by Cox and others [Cox and Raja 2008] for this workshop.

The systems engage practitioners in an active mental process of observation or visualization of mental objects, often with a representation of self, along with meta-processes that effect transformation in behaviour, state of mind, affect, and or body function. One advantage in choosing "insight meditation" as a philosophy of mind is that it focuses on aspects of consciousness important for kind behaviour. The mental processes of Vipassana have been documented to create compassion and empathy in even the most hardened criminals [Ariel and Menahemi 1997]. If human history is any prediction of the future, it looks like we will need and want some robots to exhibit kind behaviour. In the least, agents will need to effect kind behaviour in applications involving communication, health monitoring, or when working with hardware to measure emotional state of patients (users).

3. Affective inference

Definition: Affective Inference is a method of inferencing whereby emotion can be the antecedent to a consequent thought, and vice versa. There is no question that in the natural course of thinking sometimes an emotion can give rise to a thought, or that thoughts may also give rise to an emotional state, which in turn gives rise to more thoughts or further emotional states, and so on. The meta-cognition process of insight meditation highlights the interdependency of feelings and thoughts in human cognition. Many 17th century "common sense" philosophers such as Hume, Locke, and others spent considerable time on the issue of affect and rational thought. As we approach the goal of creating human-level AI or face the challenge of applications requiring social and emotional intelligence, it is essential to have some form of *affective inference*.

This style of inferencing presupposes a programming language where an agent's mental state contains explicitly named objects of emotional concept such as mood, emotional state, disposition, attitude, and so on in addition to traditional non-affective concepts and objects, and that these emotional objects require a separate but interdependent computational apparatus. Affective inference differs from logical inference in some important ways. First, by its very nature affective inference is volatile – moods and emotions change over time and so will the inferences that depend on them. An essential component of an affective inference machine or agent is a truth maintenance mechanism.

A. ***Love is Blind****: An Affective Inference Example*

We demonstrate the idea of Affective Inference in an example called *Love is Blind*. The example is interesting because it gives a different outcome depending on the personality of the agent.

R1: If FEELS(In-Love-With(x)) then Assert(Handsome(x))
R2: IF Obese(x) then NOT(Handsome(x))
R3: If Proposes(x) and Handsome(x) Then
Accept-Proposal(x)
P1: FEELS(In-Love-With(Peppy))
A1: Handsome(Peppy) {{A1}}

P2: Proposes(Peppy) Premise {{}}
D1: Accept-Proposal(Peppy) {{A1}}

Now suppose we learn
P3: Obese (Peppy) Premise {{}} Then
D2: NOT (Handsome)

ISSN (Online): 1694-0814
www.IJCSI.org

4. Practical reasoning agent

There are at least two kinds of reasoning methods applied in constructing an agent, namely practical reasoning and theoretical reasoning. Practical reasoning is directed towards actions – the process of figuring out what to do by weighing different acting options against with agent desires and believes. For example, to choose a specific transportation by a bus instead by a train is a matter of weighting alternatives against a value or care – time, cost, safety, etc. While theoretical reasoning is directed towards beliefs. For example, if I believe theory that ‘all living systems are open for environment’ and I believe that ‘human is a living system’, then the theoretical reasoning will draw a conclusion that ‘human must be open for environment’. This paper focuses on the practical reasoning agent in the following.

Practical reasoning consists of at least two kinds of activities. One is deliberation – to decide *what to do* from various desires based on the beliefs about the world. And another is means-end reasoning – to decide *how to do* it. Means-end reasoning goes ahead to generate a feasible plan to realize the intention. Among many endeavours to find out a good architecture to implement practical reasoning agent, the Believe-Desire-Intention (BDI) model has been mostly discussed in the literatures and widely accepted by the agent society.

A. Belief-Desire-Intention model

The Belief-Desire-Intention (BDI) model has been explicitly embodied in the Procedure Reasoning System (Georgeff and Lansky, 1987). The Procedure Reasoning System architecture consists of five components: plans, beliefs, desires, intentions, and interpreter. (See figure 1)The *plans* is a library of pre-compiled action plans, which are manually input by the agent programmer when the agent was built up. *Data* is input from the environment to update agent’s beliefs. And the agent will generate a plan based on his *desires* and *intentions*.

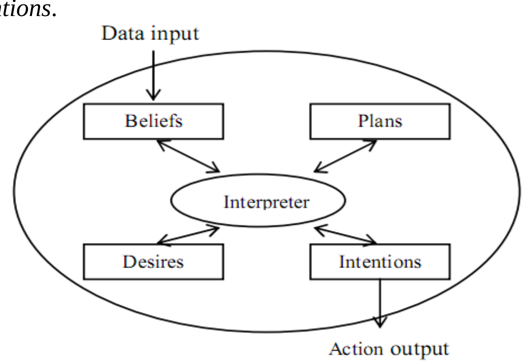

Fig. 1 Procedure reasoning system

The *interpreter* transforms, interprets, and integrates all the *beliefs, desires, plans, and intentions.*

The Procedure Reasoning System architecture was the first architecture covered all components in BDI-model, and it provided with a comprehensive template for implement an agent with BDI-model. However, the relationship among those five components cannot be easily programmed. Some researchers pointed out that implementation based on this framework (BDI) typically have no explicit representation of either desires or goals, and consequently no mechanisms for checking consistency of desires (Thangarajah and Padgham et al., 2001).

5. Belief Change

As a very young field, belief change has not been recognized as a subject of its own until the middle of the 1980’s [Hansson, 1997]. Since it is so new, it does not even have a well-established name. Belief change is just one name of the field among others such as: database updating, theory change, belief dynamics and belief revision. In general, belief change is about changing the beliefs of minds and the data of databases to accommodate new information. As already have been pointed out by [Keller and Winslett, 1985] that there are usually two types of reasons why an agent should change its beliefs. One is because the world has been changed, and the other is that the agent has made a new observation of the static world. The first type, change-recording incorporation of new information, is often called belief update. The term belief revision is reserved for the second type, knowledge-adding incorporation of new information.

The research subject which we now call belief change has mainly two origins [Hansson, 1997]. In philosophy, belief change has been studied to investigate the revision of scientific theories and logical theory. The first milestone of philosophical researches on belief change is the series studies of Levi [1977; 1980] in the 1970’s, which have under pinned the major concerns of the field and provided the basic formal framework. The next milestone is the AGM theory (named after its originators Alchourr´on, G¨ardenfors and Makinson) which has provided a more general and versatile formal framework for studying belief change [Alchourr´on and Makinson, 1982; Alchourr´on *et al*., 1985; G¨ardenfors, 1988]. In a nutshell, the AGM theory assumes the beliefs of an agent are represented by a deductively closed set of sentences (or, a belief set) of some logical language, and mainly studies how to incorporate (remove) a new sentence into (from) a belief set in a rational way. The AGM trio have studied belief change mainly in two ways. They first have introduced the so-called rationality postulates, which they claimed should be respected by any rational belief change operator. The guiding criterion of the AGM postulates is the so-called minimal change principle, that is to change the belief set as little as possible. Also, they have proposed models of constructing concrete rational belief change operators. The advent of the AGM theory finally helped the field to grow up as an important subject of its own. Since then belief change becomes a flourishing and interdisciplinary field of researches. Many researchers from different fields find the value of belief change in their own fields and thus get involved in the development of belief change. The second origin of belief change is computer science. Specifically, database theorists are interested in models of database update which are more sophisticated than those of the usual relational database [Winslett, 1990]. One important development in this direction is Doyle's "Truth Maintenance Systems" [Doyle, 1979]. Also, the problem of updating a belief set (base) is an important topic in AI [Herzig and Rifi, 1999; Dalal, 1988]. Parallel to the AGM theory, Katsuno and Mendelzon (KM) have proposed a general framework for belief update [Katsuno and Mendelzon, 1991a].

Later on, various extensions of the classical belief change have been proposed and extensively investigated. In particular, many belief revisionists are interested in iterated belief revision [Darwiche and Pearl, 1997; Nayak, 1994b], that is how should an agent revise its beliefs in response to a sequence of new information. Recently, non-prioritized belief revision, in which the new information is not always accepted in the revise belief set, has also drawn considerable attentions from the community of belief change [Hansson, 1999; Booth, 2001]. Non-prioritized revision can handle more realistic domains where there is no strict correlation between the chronology of the information and the credibility of their contents. The classical belief change is concerned with the beliefs of a single agent. There are also extended models of belief change designed for multi-agent scenarios, e.g., both belief merging [Konieczny and P´ erez, 1998; Gauwin et al., 2005] and belief arbitration [Revesz, 1997; Liberatore and Schaerf, 1998] are about how to extract the coherent common beliefs out of a set of (possibly mutually inconsistent) belief sets.

6. Human reasoning

Consider the situation of two children playing with blocks. Even in this simple situation, the children may have concerns that span many "mental realms":

Physical: What if I pulled out that bottom block?
Bodily: Can I reach that green block from here?
Social: Should I help him with his tower or knock it down?
Psychological: I forgot where I left the blue block.
Visual: Is the blue block hidden behind that stack?
Spatial: Can I arrange those blocks into the shape of a table?
Tactile: What would it feel like to grab five blocks at once?
Self-Reflective: I'm getting bored with this—what else is there to do?

We argue that no present-day AI system demonstrates such a broad range of common-sense skills. Any architecture we design should aim to achieve some competence within each of these and other important mental realms. We propose that to do this we work within the simplest possible domain requiring reasoning in each of these realms. We suggest that we develop our architectures within a physically realistic model world resembling the classic Blocks World, but where the world was populated by several simulated beings, and thus emphasizing social problems in addition to physical ones. These beings would manipulate simple objects like blocks, balls, and cylinders, and would participate in the kinds of scenarios depicted in figure 2, which include jointly building structures of various kinds, competing to solve puzzles, teaching each other skills through examples and through conversation, and verbally reflecting on their own successes and failures.

The apparent simplicity of this world is deceptive, for many of the kinds of problems that show up in this world have not yet been tackled in AI, for they require combining elements of the following:

Spatial reasoning about the spatial arrangements of objects in one's environment and how the parts of objects are oriented and situated in relation to one another. (Which of those blocks is closest to me?)

Physical reasoning about the dynamic behaviour of physical objects with masses and colliding/supporting surfaces. (What would happen if I removed that middle block from the tower?)

Bodily reasoning about the capabilities of one's physical body. (Can I reach that block without having to get up?)

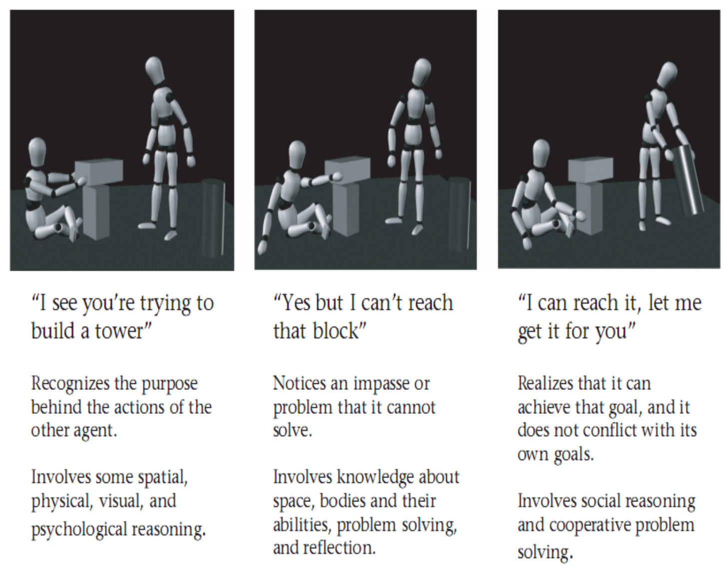

Fig. 2 Reasoning in Multiple Mental Realms to solve a problem in the Model World

Visual reasoning about the world that underlies what can be seen. (Is that a cylinder-shaped block or part of a person's leg?)

Psychological reasoning about the goals and beliefs oneself and of others. (What is the other person trying to do?)

Social reasoning about the relationships, shared goals and histories that exist between people. (How can I accomplish my goal without the other person interfering?)

Reflective reasoning about one's own recent deliberations. (What was I trying to do a moment ago?)

ISSN (Online): 1694-0814
www.IJCSI.org

Conversational reasoning about how to express one's ideas to others. (How can I explain my problem to the other person?)

Educational reasoning about how to best learn about some subject, or to teach it to someone else. (How can I generalize useful rules about the world from experiences?)

So a system should be made to face a substantial library of graded sequences of mini-scenarios that require it both to learn new skills, to improve its abilities to reflect on them, and (with practice) to become much more fluent and quick at achieving these tasks. These orderings should be based on such factors as the required complexity of objects, processes, and knowledge involved, the linguistic competence required, and the understanding of how others think and feel. That library could include all sorts of things children learn to do in such various contexts as dressing and undressing dolls, colouring in a picture book, taking a bath (or washing a dog), making toys out of Meccano and other construction kits, eating a meal, feeding a baby, cleaning a mess made by spilling some powder or liquid, reading a story and answering questions about it, making up stories, discussing behaviour of a naughty person, and learning to think and talk about the past, the future, and about distant places, etc.

7. Intuition

An obvious experience of everyday life is that people frequently make judgments and decisions without explicit use of all the relevant information that is available from the environment and from their memory. Moreover, even if people are aware of all the particular details, they do not necessarily analyse every piece of information on a deeper level and weigh it in an explicit way before making decisions. On the contrary, people often go with the very first response that enters their mind, which is usually an immediate feeling, a spontaneous idea, or a sudden appearance of "I know what to do" or "this is the best choice." This typically happens without any apparent effort, and if asked, people cannot say why they came up with a certain response. Nevertheless, people tend to trust their intuitions so frequently simply because they are quite successful with them, and the intuitions seem to "satisfice" their needs in many situations (Simon, 1955). In addition, there is plenty of evidence that people's intuitions can outperform deliberate thinking processes under specific conditions (e.g., Wilson, 2002).

Intuition is a process of thinking. The input to this process is mostly provided by knowledge stored in long-term memory that has been primarily acquired via associative learning. The input is processed automatically and without conscious awareness. The output of the process is a feeling that can serve as a basis for judgments and decisions.

The definition specifies the input to intuitive processing: knowledge acquired through experience and stored in long-term memory. As such, intuition primarily capitalizes on prior knowledge acquired via a slow learning system rather than on recently encountered information kept in short-term memory. The slow learning system is guided by the principles of associative learning that have been studied by using, for instance, Pavlovian, evaluative, and operant conditioning procedures. The intimate relation between prior experience and intuition is highlighted by most of the theories on intuition.

According to our definition, the output of intuition is a feeling, for instance, the feeling of liking an entity or a feeling of risk. Feelings are a powerful means of communication, not only between individuals (e.g., via facial expressions) but also within the organism. Feelings arise involuntarily and immediately break into consciousness (Wundt, 1907; Zajonc, 1968, 1980). Thus, they can serve as an interruption device changing subsequent motivations. In a nutshell, intuition is assumed to exploit the capability of one's mind to process information in parallel. Parallel processes can handle a huge amount of information. In judgment and decision making, long-term memory provides the main database for this process. Therefore, intuition in judgment and decision making rests to a great extent on knowledge that is well consolidated and deeply ingrained in memory.

Several approaches to intuitive judgment and decision making converge in assuming that automatic and deliberate processes can (and mostly do) operate simultaneously and thus jointly shape thought and action. Presumably, the ideal case of pure intuition or pure deliberation does not exist in reality. Accordingly, when distinguishing different strategies of thinking on the empirical level, one should seek to specify the relative contribution of the two processes. If automatic processes dominate, one should expect that the judgments or decisions are reached very quickly by virtue of parallel processing. At the same time, judgments and decisions should reflect a strong sensitivity for entire samples of prior experiences even if the sample is huge because parallel processes can handle a large amount of information. Moreover, intuitive judgments and decisions should be prone to undervalue the weight of new evidence because they strongly rely on consolidated knowledge. Figure 3 shows a coordinate system containing styles of processing, amount of information, and its degree of consolidation as dimensions.

ISSN (Online): 1694-0814
www.IJCSI.org

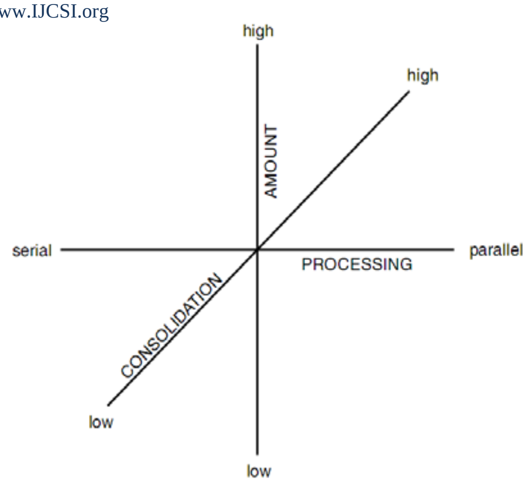

Fig. 3 Coordinate System containing styles of processing, amount of information, and its degree of consolidation as dimensions

8. Decision Making

The human decision-making process can be regarded as a complex information processing activity. According to (Rasmussen, 1983), the process is divided into three broad categories that correspond to activities at three different levels of complexity. At the lowest level is skill-based sensorimotor behaviour, representing the most automated, largely unconscious level of skill-based performance such as deciding to brake upon suddenly seeing a car ahead. At the next level is rule-based behaviour exemplified by simple procedural skills for well-practiced, simple tasks, such as inferring the condition of a game-playing field based on the current rainy weather. Knowledge-based behaviour represents the most complex cognitive processing. It is used to solve difficult and sometimes unfamiliar problems, and for making decisions that require dealing with various factors and uncertain data. Examples of this type of processing include determining the status of a game given the observation of transport disruption.

Human decision-makers often weigh the available alternatives and select the most promising one based on the associated pros and cons. The P3 model will represent these pros and cons as logical sentences with embedded probabilities as follows:

<bel>Heavy Rain?<sup0.7 >Cancelled

The above sentence can be interpreted as follows: if the agent believes that it rained heavily then it asserts that there is a 70% chance (equivalently, generates an amount of support 0.7) that the game will be cancelled. An agent may obtain evidence from different sources as support for the cancellation, as well as support against the cancellation, as follows:

<bel>Club FinancialCrisis ?<sup0.6 >¬Cancelled

The above sentence states that if the club is in financial crisis then there is a 60% chance that the cancellation will be avoided. These types of P3 sentences that provide support both for and against various decision options constitute arguments used by the agent to solve a decision problem. Such an argumentation-based decision-making framework has been developed in (Das *et al*. 1997; Fox and Das, 2000).

9. Conclusions

A key aspect of modeling human behavior is capturing the effect of meta-information on information processing, situation assessment, and decision-making. Our experience performing analyses of human cognition and action in different decision-making domains has shown that humans use (and/or fail to use) this meta-information when making decisions.

The work reported in this paper demonstrates that researchers can make substantial progress towards achieving human-level artificial intelligence by integrating the approaches of multiple sub-disciplines in artificial intelligence and by drawing inspiration from the study of human intelligence.

It is obvious that developing human-level intelligence is a huge challenge. However, important parts of that scientific and engineering enterprise are the methods and practices for evaluating the systems as they are developed. In this paper, we present some of the primary challenges that arise in evaluation that distinguish it from research on more specialized aspects of artificial intelligence. We also attempt to characterize the types of scientific claims that arise in research on HUMAN LEVEL RESONING AND DECISION-MAKING, distinguishing different classes of claims that can be made at the system level, and then further analyzing the independent and dependent variables of those claims. However, much work remains to identify methods for measuring and evaluating these capabilities.

References

[1] *A Survey of Agent-Based Intelligent Decision Support Systems to Support Clinical Management and Research:* Darren Foster Carolyn McGregor Samir El-Masri Health Informatics Research School of Computing and IT University of Western Sydney

[2] *Influence-Based Autonomy Levels in Agent Decision-Making*:B. van der Vecht 12 and A.P. Meyer1 and R.M. Neef 1 and F. Dignum2 and J-J.Ch. Meyer2.

[3] *The St. Thomas Common Sense Symposium: Designing Architectures for Human-Level Intelligence* :Marvin Minsky, Push Singh, and Aaron Sloman.

[4] *Belief change in reasoning agents:* Dipl. Inf. Jin.

[5] (2002) The IEEE website. [Online]. Available: http://www.ieee.org/

[6] *Considerations Regarding Human-Level Artificial Intelligence:* Nils J. Nilsson (nilsson@cs.stanford.edu) Robotics Laboratory Department of Computer Science Stanford University

[7] Brooks, R. A. 1997. From Earwigs to Humans. Robotics and Autonomous Systems 20(2–4): 291–304, June.

(Available at people.csail.mit.edu/u/b/brooks/public_html/papers/ascona.pdf.)

[8] Drescher, G. 1991. Made-up Minds: A Constructivist Approach to Artificial Intelligence.Cambridge, MA: MIT Press.

[9] Dzeroski, S.; De Raedt, L.; and Driessens, K. 2001. Relational reinforcement Learning. Machine Learning 43(1–2): 7–52.

[10] Education Direct 2005. Education Direct Website. Stamford, CT: Thomson Learning, Inc. (Available at www.educationdirect.com/programs_cat.html.)

[11] Engelmore, R., and Morgan, T., eds. 1988.Blackboard Systems. Menlo Park, CA: Addison-Wesley Publishing Company.

[12] Feigenbaum, E. A. 2003. Some Challenges and Grand Challenges for Computational Intelligence. Journal of the ACM, 50(1): 32–40, January.

[13] Friedland, N. S., et al. 2004. Project Halo: Towards a Digital Aristotle. AI Magazine 25(4): 29–47, Winter.

[14] Hawkins, J., and Blakeslee, S. 2004. On Intelligence. New York: Times Books.

[15] Koza, John R., et al. 2003. Genetic Programming IV: Routine Human-Competitive Machine Intelligence. Norwell, MA: Kluwer Academic Publishers.

[16] Lakoff, G., and Johnson, M. 1980. Metaphors We Live By. Chicago: University of Chicago Press.

[17] Lenat, D. B. 1995. CYC: A Large-Scale Investment in Knowledge Infrastructure. Communications of the ACM 38(11): 33–38, November. (Available at web.media.mit.edu/~lieber/Teaching/Common-Sense Course-02/Lenat-CACM.pdf.)

[18] McCarthy, J. 1996. From Here to Human Level Intelligence. Unpublished memo, Department of Computer Science, StanfordUniversity, Stanford, CA. (Available at www-formal.stanford.edu/jmc/human.html.)

[19] McCarthy, J. 1999. The Well-Designed Child. Unpublished memo, Department of Computer Science, Stanford University, Stanford, CA. (Available at http://www-formal.stanford.e du/jmc/child.pdf.)