Canadian Artificial Intelligence

Intelligence Artificielle au Canada

Canada's National AI Publication

La Publication Nationale en IA au Canada

Winter 2002 | No. 50 | hiver 2001

Contents

Contenu

Canadian Artificial Intelligence is published by the Canadian Society for Computational Studies of Intelligence (CSCSI). ***Intelligence Artificielle an Canada*** est publiée par la Société canadienne pour l'étude de l'intelligence par ordinateur (SCEIO). Canadian Publications Mail Product Sales Agreement No. 1497200.

ISSN 0823-9339

Web Version Web

Please **don't throw away** the **envelope** in which CAI/IAC arrived. It contains information about your userID and password, which are needed to access the members-only area of the CSCSI/SCEIO website:

http://cscsi.sfu.ca/cai.html

Sample issues and articles are accessible to non-members. The members-only area contains this issue and some past issues of *CAI/IAC*. To access the area, type your userID and password at the login window.

Your **userID** is the first letter of your first name plus up to seven letters of your last name. For example, the userID for Anne Murray is amurray.

Your **password** is based on your CSCSI/SCEIO membership number which is **printed on the envelope** in which this issue arrived. Take that number and prepend to it the first letter of your first and last name. (e.g. if Anne Murray's membership number was 876543, then her password would be am876543.)

President's Message

Robert Mercer

As my term of office as president draws to a close, I want to use this opportunity to bring to the attention of the membership a number of issues facing our society. The first item is the upcoming change to a new executive. So, to start this president's message, I would like to thank the present and past executive for (nearly) four years of commitment to our society. We have had, as vice-presidents, Russ Greiner (1998-2000) and Dekang Lin (2000-2002), as editors of the magazine, Dan Fass (1998-2000) and Ann Grbavec (2000-2002), as secretary, Guy Mineau (1998-2002), as treasurer, Howard Hamilton (1998-2002), and as past president, Fred Popowich (1998-2002). Everyone has contributed some time and effort. Because of their special duties I would like to thank Dan and Ann for maintaining the quality of the CAI magazine and for managing the website, Guy for maintaining communication with CIPS and our membership lists, Howard, for always getting the cheques out on time and for his advice on a number of matters, and of course Fred, who in addition to helping to coordinate the magazine, as past-president, helped to maintain the direction of the society.

To finish this first item: I have struck a nominating committee to present a slate for the next executive. The nominating committee is composed of Fred Popowich, Alan Mackworth, Janice Glasgow, and me. As well, I want to direct your attention to the notice for a call for nominations that is in this issue of the magazine. We are using this issue of the CAI magazine to advertise this call for nominations, even though the call is not in effect until after the nominating committee has presented its slate.Watch the website for further information: http://cscsi.sfu.ca/.

Now on to the other issues that are before us.

The New Electronic-Only Version of the CAI Magazine

Our current plan is to make this issue the final CAI magazine to be published in hardcopy form. For the 51st issue of the CAI magazine, we are moving to an electronic-only version (issues 38 through 49 are already available on the CSCSI/SCEIO website: http://cscsi.sfu.ca/ in the members-only area). The idea of replacing the paper version with an electronic version has been discussed many times in the past. The members who were present at the last AGM were quite enthusiastic about moving to an electronic-only form of magazine.

The electronic-only magazine is part of an enhanced web presence for the society. We are in the midst of launching www.cscsi.org and www.sceio.org. This launch will occur as soon as an important ontological problem (outside the control of the society) has been resolved.

The Annual Conference

Over the past decade (and maybe longer) it was oftentimes suggested that the CSCSI/SCEIO sponsored biennial conference become an annual event. Historically, the reason for having a biennial conference was to be in the years between IJCAI. These historical reasons had by the 1990s become moot with the creation of a number of other major AI conferences (AAAI, ECAI) and a number of major special interest conferences (KR, ML, UAI, KDD). In addition, we had started a conference alliance with two other Canadian societies, CHCCS and CIPPRS/ACTIRF. This common conference is known to all of us as AI/GI/VI. The other two societies have always had annual conferences. It just seemed that we were out of step.

Starting with the previous executive, discussions of holding an annual conference became more serious. Finally, at the AI/GI/VI'2000 meeting we committed CSCSI/SCEIO to holding an annual conference. This decision was made knowing that Canada had become a destination of choice for a number of AI, graphics, and vision conferences (for instance, The National Conference on Artificial Intelligence (AAAI) together with a number of allied conferences will be held in Edmonton at the end of July 2002, The International Conference on Shape Modelling and Applications will be held in Banff in May 2002, and The International Conference on Pattern Recognition will be held in Quebec City in the middle of August 2002).

We held our first annual conference in 2001. We are now in the midst of organizing the conference for 2002 in Calgary (the accepted papers have recently been announced and the local arrangements are nearly finalized). Even with AAAI being in Canada and the surprising events of September 11 (which had some of the organizers suggesting reduced conference travel), we have had a nearly normal submission of papers for AI'2002. So, this is very encouraging news for our annual conference.

We are preparing for the 2003 and 2004 conferences, as well. Yang Xiang is spearheading 2003 at Dalhousie University (Halifax) and 2004 is being looked after by Peter van Beek. The location that is currently being suggested is The University of Western Ontario (London). I invite all members to participate at AI/GI/VI'2002 in Calgary. And I hope to see you all at the Annual General Meeting which will be held during the conference.

International Organization

At the last IJCAI meeting in Seattle (2001), representatives of the national and trans-national AI societies were invited to participate in an informal discussion regarding the possibility of forming an international AI society. I was the only representative to speak in favour of an international organization. What was acceptable to all in attendance was the establishment of an international website/information repository (e.g. mailing lists, links to national organization websites) maintained by AAAI.

It was agreed that an informal meeting be held at each IJCAI to discuss issues that confront the societies. I still hold out some hope for an international organization, at least in the form of an umbrella clearinghouse, that is, some place to send dues, have those dues properly forwarded to the agencies that control publications and other things such as national memberships. Maybe a time will come for an international organization to emerge. Some group will undoubtedly try to start one.

CSCSI
SCEIO

AI 2002

The Fifteenth Canadian Conference on Artificial Intelligence

27-29 May, 2002

University of Calgary, Calgary, Alberta, Canada

Sponsored by Canadian Society for Computational Studies of Intelligence (CSCSI)
Société Canadienne pour l'étude de l'intelligence par ordinateur (SCEIO).

Call for Nominations - CSCSI Executive

Under terms of the constitution, a nominating committee has been struck. It will present a slate of candidates for these positions on 3 April 2002. The slate of candidates will be announced on the CSCSI website.

A Call for Nominations for all positions on the CSCSI Executive will open on 3 April 2002.

Positions to be filled:

- President
- Vice-President
- Secretary
- Treasurer

Each nomination must have a nominator. The nominee must agree to the nomination under separate cover. Completed nominations must arrive by 24 April 2002. Nominations for the CSCSI Executive are to be addressed to:

Robert Mercer
Department of Computer Science,
University of Western Ontario,
London, ON, N6A 5B7
email: mercer@csd.uwo.ca

Appel de Candidatures - Positions Executives du SCEIO

En vertu des conditions de la constitution, un comité de nomination a été composé. Il présentera une liste des candidats pour ces positions le 3 avril 2002. La liste des candidats sera annoncée sur le website du SCEIO.

Un appel en Nominations pour toutes les positions exécutives du SCEIO s'ouvrira 3 avril 2002.

Positions à remplir:

- Président
- Vice-Président
- Secrétaire
- Trésorier

Chaque nomination doit avoir un nominateur. Le dénommé doit approuver la nomination independemment. La documentation complète pour chaque dénomé doit arriver d'ici le 24 avril 2002. Les nominations pour les positions exécutives du SCEIO doivent être adressées à:

Robert Mercer
Department of Computer Science,
University of Western Ontario,
London, ON, N6A 5B7
email: mercer@csd.uwo.ca

CALL FOR NOMINATIONS

CSCSI Distinguised Service Award

This award is presented biannually to an individual who has made outstanding contributions to the Canadian AI community in one or more of the following areas:

- Community service
- Research
- Training of students
- Research/Industry interaction

The award, which will be presented at the CSCSI conference, provides:

- honorary lifetime membership in CSCSI
- conference fees when the award is presented
- a token gift

Recommendations for the award should be addressed to me to arrive by **30 April 2002**. They should include a brief (1 page) summary of the nominee's qualifications for receiving the award. The final decision will be made by the CSCSI executive.

It is very much hoped that the winner will be present to receive the award in person, and we ask nominators to make discreet enquiries in advance to ensure that there is a reasonable likelihood of their nominee attending the conference.

APPEL DE CANDIDATURES

Prix de Distinction de la SCEIO Pour Service Rendu

Cette récompense est présentée deux fois par année à un individu qui aura fait une contribution hors de l'ordinaire à la communauteé d'Intelligence Artificielle au Canada dans au moins un des secteurs suivants:

- Service à la communauté
- Rescherce
- Formation d'étudiants
- Intéraction Recherche/Industrie

La récompense qui sera présentée à la confèrence de la SCEIO, comprend:

- Carte de membre à vie à la SCEIO à titre de membre honoraire
- Accès gratuit à la conférénce à laquelle la récompense sera attribuée
- Cadeau d'appréciation

Les nominations devraient m'être adressé personnellement au plus tard le **30 avril 2002**. Elles devraient inclure un bref (1 page) sommaire des qualifications du candidat. La décision finale sera rendu par l'exécutive de la SCEIO.

Il serait souhaitable que le gagnant puisse venir chercher sa récompense en personne, et nous demandons aux parrains de s'informer auparavant de la disponibilité de leur candidat, quant à la possibilité que le candidat assiste à la conférence en vue de recevoir son prix.

Nominations for CSCSI Distinguished Service Award are to be addressed to:

Robert Mercer
Department of Computer Science,
University of Western Ontario,
London, ON, N6A 5B7
email: mercer@csd.uwo.ca

Papers to be Presented at AI 2002

May 27-29, 2002

Calgary, Alberta, Canada

CSCSI
SCEIO

Invited Speakers

Alan Mackworth University of British Columbia, Canada
Zenon Pylyshyn Rutgers, The state university of New Jersey, USA
Len Schubert University of Rochester, USA

Papers Accepted for Oral Presentation

Scott DeLoach
Modeling Organizational Rules in the Multiagent Systems Engineering Methodology

Richard Frost and Pierre Boulos
An efficient compositional semantics for natural-language database queries with arbitrarily-nested quantification and negation

C.J. Butz and M.J. Sanscartier
On The Role of Contextual Weak Independence in Probabilistic Inference

George Lashkia
A Noise Filtering Method for Inductive Concept Learning

Terry Copeck, Nathalie Japkowicz, Stan Szpakowicz
Text Summarization as Controlled Search

C.I. Ezeife and Yu Sue
Mining Incremental Association Rules with Generalized FP-tree

Arun K Pujari, C Dhanunjaya Naidu, B C Jinaga
An Adaptive and Intelligent Character Recogniser for Telugu Scripts using Multiresolution Analysis and Associative Memory

Luc Plamondon, Guy Lapalme, Leila Kosseim
QUANTUM: A Function-Based Question Answering System

Daniel L. Silver, Robert E. Mercer
Sequential Learning through Task Rehearsal: Overcoming Impoverished Data

Peter Yap
Grid-based Pathfinding

Akihiro Kishimoto and Jonathan Schaeffer
Transposition Table Driven Work Scheduling in Distributed Game-Tree Search

Hanan Ayad and Mohamed Kamel
Topic Discovery from Text using Aggregation of Different Clustering Methods.

David Chen and Robin Cohen
AERO: An outsourced approach to exception handling in multi-agent systems

Thomas Tran and Robin Cohen
A Learning Algorithm for Buying and Selling Agents in Electronic Marketplaces

Silvia Breban and Julita Vassileva
Using Inter-Agent Trust Relationships for Efficient Coalition Formation

Kevin Kennedy and Robert E. Mercer
Using Communicative Acts to Plan the Cinematography Structure of Animations

Eric Neufeld
Clue as a testbed for automated theorem proving

Yllias Chali
Generic and Query-Based Text Summarization Using Lexical Cohesion

Michael C. Horsch, William S. Havens,Aditya K. Ghose
Generalized Arc Consistency with Application to MaxCSP and SCSP Instances

Alan Fedoruk and Ralph Deters
Using Agent Replication to Enhance Reliability and Availability of Multi-Agent Systems

Guan-Shieng Huang, Xiumei Jia, Churn-Jung Liau, Jia-Huai You
Two-Literal Logic Programs and Satisfiability Representation of Stable Models: A Comparison

S. K. M. Wong and Tao Lin
Construction of a Non-redundant Cover for Conditional Independencies

S.K.M. Wong, D. Wu
A Structural Characterization of DAG-Isomorphic Dependency Model

J. R. Parker
Genetic Algorithms for Continuous Problems

Papers Accepted for Poster Presentation

G. Grewal, T. Wilson, and C. Nell
An Enhanced Genetic Algorithm Approach to the Channel Assignment Problem in Mobile Cellular Networks

Robert St. Amant
Intelligent Data Analysis in an Interactive Planning Simulation

Yang Xiang, Chenwen Ye, and Deborah Ann Stacey
Application of Bayesian Networks to Shopping Assistance

Kamran Karimi and Howard J. Hamilton
RFCT: An Association-Based Causality Miner

Mathias Salle
Electronic Contract Framework for Contractual Agents

Tamer S. Mahdi, Robert E. Mercer
An Algorithm for Lexical Functional Mapping

Christel Kemke
A Constructive Approach to Parsing with Neural Networks - the Hybrid Connectionist Parsing Method

Jan Bakus, Mohamed Kamel
Extraction of Text Phrases using Hierarchical Grammar

F. L. Wang and J. Greer
Retrieval of Short Documents from Public Discussion Forums

Computation, Meaning and Artificial Intelligence: Some Old Problems, Some New Models

Timothy P. Racine

Résumé

Les approches conventionelles (c.-à-d., GOFAI, PDP) n'ont pas convenablement modelé l'intelligence artificielle. Je présente un sénario où le sense discernable par un système et l'erreur discernable par un système sont des critères nécessaires pour l'ascription d'intelligence à un agent. Inspiré principalement par le travail de Mark Bickhard et Ludwig Wittgenstein, j'essaye de prouver que le talon d'Achilles des approches conventionelles est la présomption de notions de correspondance par raport au sense et au problème de représentation. Je revise quelques modèles utilisés récemment (connaissance située, constructivist AI, interactivism) qui soulignent la nature située et incorporée des agents intelligents pour voir s'ils proposent une amélioration. Je conclus que l'interactivism est un modèle qui se montre prometteur pour modeler l'intelligence parce qu'il ne se fonde pas sur des correspondances pour dériver le sense et la représentation et explique également le sense et l'erreur discernables par un système.

Abstract

Standard approaches (i.e., GOFAI, PDP) have not adequately modelled artificial intelligence. I present the case that system detectable meaning and system detectable error are necessary criteria for the ascription of intelligence to an agent. Drawing primarily on the work of Mark Bickhard and Ludwig Wittgenstein, I endeavour to show that the Achilles heel of standard approaches is their reliance upon correspondence notions of meaning and representation. I review some recent models (situated cognition, constructivist AI, interactivism) that emphasize the situated and embodied nature of intelligent agents to see whether they fare any better. I conclude that interactivism is a model that shows promise in modelling intelligence because it does not rely on correspondences to derive meaning and representation and also accounts for system detectable meaning and error.

Introduction

> Some models in low-level perception have been capable of building primitive representations of the environment, but these are not yet sufficiently complex to be considered "meaningful". Much research in high-level cognitive modeling has *started* with representations...where any meaning present is already built in. There has been little work that bridges the gap between the two.
>
> Hofstadter (1995, p. 174, emphasis original)

The field of artificial intelligence materialized for the most part out of attempts to model (human) intelligence by writing programs that could serve as theories regarding its instantiation (Simon, 1996). Hooker (1996) distinguishes this symbol manipulation meta-theoretical approach (which he terms narrow AI theory, i.e., GOFAI — Good Old Fashioned AI, see Haugeland, 1985) from a broader framework (wide AI theory) Hooker constructed to incorporate the connectionist systems that eschewed formal logical programming but remained faithful to a computational metaphor of mind. Hooker created an even broader category, general naturalized intelligence theory (GNIT), to capture models that may or may not rely on computational predicates, but which necessarily frame intelligence as a basic property of natural objects.[1] Many GNIT models conceive of pattern and function as properties that emerge from interaction rather than as top-down constraints specified by higher-order cognitive structures. Accordingly, these models recast logic and language as derivative of system-environment interaction. GNIT incorporates models that rely on self-organizing nonlinear systems dynamics, and cybernetic or control principles (Hooker, 1996). As will be seen below, some of these models are representational in nature, whereas others are not.

1. In this paper, I will emphasize the situated and embodied in nature of naturalized AI models. Hooker (1996) does not touch on these themes; however, he does assume them.

The preparation of this article was supported by a Social Sciences and Humanities Research Council of Canada Doctoral Fellowship.

Part of the impetus for the construction of many naturalized models of artificial intelligence has been provided by dissatisfaction with standard AI models (whether narrow or wide).[2] This discontent may be well founded. Simply put, whereas standard models can be used to construct machines that may perform notable tasks (e.g., defeating chess masters) and may contain intriguing architectures (e.g., neural nets that can, in some sense, learn), there are compelling grounds upon which to suspect that these approaches do not and perhaps even cannot capture intelligence. To demonstrate this, I will focus on whether exemplary models working in these three meta-theoretical traditions can account for both system detectable error (error for an agent/system) and system detectable meaning (meaning for an agent/system) in a non-circular manner.[3] I will justify these criteria by investigating whether standard approaches satisfy them. In so doing, I will also question whether the standard correspondence notion of meaning (wherein meaning is said to reside in mental structures) is coherent. I will then review some recent naturalized models (situated cognition, constructivist AI, interactivism) that emphasize the situated and embodied nature of intelligent agents to see whether these fare any better. I will conclude that interactivism shows promise in modelling intelligence because it does not rely on correspondences to derive meaning and representation and also accounts for system detectable meaning and error.

The Physical Symbol System Hypothesis

Newell and Simon (1976, p. 113) claimed that "a physical symbol system has the necessary and sufficient means for general intelligent action." From this premise, they argued not only that a physical symbol system (PSS) can be programmed to behave intelligently, but also that human beings are intelligent precisely because they are physical symbol systems (see Simon's, 1969, symbol system hypothesis for earlier related ideas). Accordingly, intelligence should be explained in terms of symbols and operations on symbols. Simon (1996) reports that though there has been a great deal of empirical support for the PSS hypothesis, he is "surprised and frustrated" by how rarely this evidence is cited. He further remarks that "there appears still to be a widespread belief that the nature of human thought processes can be determined from *first principles* without examining human behaviour in painstaking detail and comparing that behaviour with claims of rigorous theories" (1996, p. 162, emphasis original). However, Simon is wrong — the plausibility of a theory is *not* an empirical issue (see Lakatos, 1970). Bickhard and Terveen (1995) have presented a series of in-principle arguments that turn on this very point. They demonstrate that there may be grounds upon which to conclude that the programmatic aims of AI mentioned above are unreachable if viewed from a narrow AI theoretical lens (or possibly even wider). A rejoinder to Simon's (1996) remark, then, is that no amount of empirical work can substantiate a theory that is fatally flawed on logical grounds (Bickhard, 1993, 1996, 1998, 1999b, 2000; Bickhard & Terveen, 1995). That is, the empirical work cannot support the claims based on it.

The following example is illustrative of the PSS hypothesis. Simon (1996, p. 164, emphasis removed) claims that: "a thought about a cat, perhaps, induced by looking at it or by remembering it, is a symbol structure in that part of the brain (probably the frontal lobe) where the 'mind's eye' resides." There are five main problems manifest in this quote and in the PSS hypothesis. (1) He argues that looking at or remembering a cat activates the 'cat' representation. However, his quote assumes the existence of 'cat'. Where did this cerebral physical symbol come from? On his model, one would presume it comes from directly encoding an actual cat, or from activation of an innate 'cat' representation. As will be shown later in the paper, both of these accounts are problematic. (2) Although Simon seems to be innocuously relying on a popular metaphor here ('mind's eye'), he inadvertently points to an initial problem with his position — a physical symbol structure is ontologically committed to an interpreter (a 'mind's eye') of said symbol in order to unpack its meaning (see Bickhard & Terveen, 1995; Heil, 1981; Müller, Sokol, & Overton, 1998).

How this is the case will also be unpacked later in this paper, but for now it suffices to point out that because Simon is in need of an interpreter for his symbol structure, he also requires an interpreter for that interpretation — which causes an infinite regress of such interpreters that seems to doom this enterprise.

(3) Simon's account also falls into the analogous infinite regress of the symbol grounding problem (see, e.g., Searle, 1981) which necessary occurs when symbol

2. Similarly, PDP came about as a result to dissatisfaction with GOFAI.

3. System detectable meaning is, of course, logically dependent upon system detectable error. If an agent cannot detect error, it is impossible for an agent to unpack meaning. I have separated these terms in this paper for illustrative purposes. Further, as will become apparent in this paper, meaning is not something that is detected by a system; rather, a system constructs meaning.

manipulation is relied upon to model intelligence. The problem arises as a result of the fact that within formal computation a given symbol is defined in respect to other purely formal symbols — thereby leaving all symbols ungrounded and meaningless.[4] (4) As a consequence of this dilemma, Simon has no account of how such representations can be assigned meaning by the agent in question (what I call the issue of system detectable meaning) whether or not one accepts that it is impossible for physical symbol structures to exist. Hofstadter (1995) noted above that this type of approach has started with extant representations where any meaning present in the system is already preordained. By definition, the meaning is not available to the machine. Depending on the aspirations of the researcher, this may or may not be a problem for designing a given system — it is problematic, however, for modelling intelligence because it is not enough that an agent perform behaviours that appear intelligent to an observer. Turing (1950), of course, proposed the opposite in the so-called Turing test.[5] An intelligent agent, however, must itself take Newell and Simon's symbols to mean what they are purported to represent. If not, in what sense should Turing or anyone else ascribe intelligence to that agent? (5) Simon (1996) has no account of how it would be possible to know whether the ostensible mental structures represent what they are taken to represent (the issue of system detectable error, Bickhard, 1999b, p. 437): "if representational error is not detectable, then neither representational goal directedness nor representational learning are possible, whether in animals or in machines. System detectable error, thus, is a fundamental criterion for an acceptable model of representation" (see also Hacker, 1990; ter Hark, 1990; Wittgenstein, 1958). I therefore operate under the assumption that system detectable meaning and system detectable error are necessary criteria for the ascription of intelligence to an agent.[6] That is, the ability to understand what the world means, to learn and pursue goals, are minimally constitutive of intelligence.

4. This is the technical consequence of point (1) above.

5. However, Turing was, in fact, cognisant of the fact that the symbol manipulation approach he advocated was limited. Brooks (1991/1999, pp. 141-142) notes that that in a paper published well after Turing's death (Turing, 1970), Turing argued that a machine capable of learning, for example, languages, would need to be 'embodied' (see below) because success would "depend rather too much on the sense organs and locomotion to be feasible." Brooks (1999) argues that Turing consciously chose to utilize a symbol manipulation (rather than embodied) method for reasons of technical convenience.

6. Whereas these criteria are necessary, they may be insufficient. See below.

The Language of Thought

Researchers and theorists of the narrow AI ilk are forced to either ignore the problem of how of it that an agent understands what symbols mean as Newell and Simon (1976) seem to have done, to assume that external objects can be directly encoded, or to embrace the nativist position that syntax is innate (Fodor, 1975, 1981).[7] Like Newell and Simon (1976), Fodor (1981) claims that (human) intelligence reduces to rule-governed operations on physical symbols. Fodor (1975) wrestles with the problem of how such a physical structure could refer to things in the world. He reasons that thought (like language) requires its own medium of representation. This medium would require a rule-governed structure — that is, a grammar. He goes on to argue that we could not know whether our thoughts have meaning without such a medium. He concludes that humans must therefore have an innate 'language of thought' onto which we graft natural language: "One cannot learn that P falls under A unless one has a language in which P and A can be represented. So one cannot learn a first language unless one already has a system capable of representing the predicates in that language and their extensions. And, on pain of circularity, that system cannot be the language being learned. Hence, at least some cognitive operations are carried out in languages other than natural languages" (Fodor, 1975, pp. 63-64, emphasis original). Fodor's language of thought, interestingly, is logically analogous to machine language. On this logic, if the human mind is analogous to a computer, then its functions must presumably require a similar structure. Fodor (1981) completes this line of argument by deducing that the language of thought must have been constructed through natural selection.

Although quick to admit that although his position may seem absurd, Fodor (1975, p. 82) argues that no better alternative to his theory exists. Nevertheless, in a more sober moment Fodor acknowledges (1981, p. 223) that he wishes "someone very nice and clever" would show him how to ground his language of thought in a foundational semantics. However, he inherits an even larger problem. Namely, if a pre-existing language delimits any conceivable thought that could arise in an agent, then there is no possible account of new ideas (Bickhard, 1993; Bickhard & Terveen, 1995). As such, evolution must have a perspicacity so awesome that it could 'foresee' everything that could ever be expressed in any natural language — a tall order indeed. It is accordingly not surprising that research has not substantiated the

7. See also Chomsky (e.g., 1964).

idea of a language of thought (e.g., Choi, 1997; Slobin, 1997).

Latent Problems and Assumptions

Chapman (1999, p. 31) notes that it seems obvious that "knowledge consists in having a mental representation that corresponds to reality as it really is". However, he shows that a problem with such a correspondence view of knowledge and meaning is that it takes for granted the very thing it is supposed to explain. This creates the monolithic problem (echoed by Bickhard's notion of system detectable error) that in order "to check the accuracy of our representations, we need some *other* way of knowing reality besides representing it. And if we cannot check, then we can never know if our representations are accurate or not " (Chapman, 1999, p. 31, emphasis original; see also Bickhard, 1993, 1996; 1998; Bickhard & Terveen, 1995; Chapman, 1987; Hacker, 1990; ter Hark, 1990; Wittgenstein, 1958). So, correspondence is not enough. Fodor was sensitive to this predicament, which is presumably why he came to the nativist position that he did. However, in light of the above his solution must be judged inadequate (see also Bickhard, 1993, 1999a; Bickhard & Terveen, 1995; Canfield, 1999; Goldberg, 1991; Heil, 1981; Shannon, 1998).

Müller, Sokol and Overton (1998) point out that the PSS hypothesis implicitly assumes the principles of empiricism wherein knowledge is passively built up from elementary sensory data. The meaning of a representation is ostensibly (and circularly) captured within the content of the object of representation. Accordingly, meaning is therefore derived from correspondence with some external referent. Empiricism assumes that "meaning resides outside the person, in pieces of the world, which the mind, in turn, reproduces in its own symbols" (Müller et al., 1998, p. 160). But if such a reproduction of symbols is logically impossible, we need a different account of how an agent can represent its world (and, relatedly, what it is taken to mean). Such an account will be developed below.

Bickhard (1993, 1996; 1998, 1999b, 2000; Bickhard & Terveen, 1995) uses *encodingism* to refer to any model that assumes all representations are encodings. I have implied one problem with this approach above — namely that encodings take for granted the very thing that they are supposed to explain (Chapman, 1999; Wittgenstein, 1958). Put another way, there is no way of verifying whether an encoding is a representation because it is assumed that *all* representations are encodings. Bickhard uses the example of Morse code to show what an encoding actually is (namely in this case, a stand-in for an alphanumeric character). He (1999b) explains that in Morse code, '...' represents (stands in for) the character 'S'. Encodings as stand-ins, then, have representational content. That is, they represent what they stand in for. Bickhard and Terveen (1995, pp. 13-14, emphasis original) point out, however, that "encodings can *carry* representational contents, and already established encodings can *provide* representational contents for the formation of some other encoding, but there is no way within encodingism per se for those representational contents to ever arise in the first place." This is the problem of emergence of representation (or more generally, the ability to learn and pursue goals, see above). Although encodingism occurs by definition when using a symbol manipulation approach (as a result of the symbol grounding problem), as will be seen below it also infects some naturalistic models.

Connectionism and PDP

Advocates of PDP (e.g., Quartz, 2001) have protested that some of their critics (e.g., Fodor, 2000) seem only familiar with connectionist architecture circa 1986. To be sure, the contemporary distributed processing approach is a force to be reckoned with (see e.g., O'Reilly & Munakata, 2000). Mixed symbolic-connectionist models (so-called 'symbolic connectionism') have also been proposed that try to avoid the weaknesses inherent in each approach (Holyoak & Hummel, 2000). For the present purposes though, I am more interested in the typical assumptions made by those working under this banner. As suggested above, Hooker (1996) constructed the wide AI meta-theoretical category in response to connectionist architectures because such systems do not rely on symbol manipulation in order to compute.

Narrow AI assumed that the brain operates on physical symbols, whereas wide AI left behind physical symbols by using brain connections as the inspiration for modelling intelligence (Rumelhart, 1989). This is progress in that connectionists have tried to look at how the brain actually might function as opposed to assuming that it functioned in a particular way. It notable though that McDonough (1999, p. 173) points out that "almost everyone, it seems, holds that human intelligence is produced by an internal intelligence machine (the brain)." He argues that cognitive science has traditionally either relied on Mental Representation Theory (narrow AI), or Neural Net Theory (wide AI) to which I shall presently turn. Theories of mental representation assume that the brain performs computations on discrete (encoded) representations, whereas neural net theories assume knowledge can be distributed over all the elements in the brain (but see Bickhard & Terveen, 1995). Does this get us any closer to system detectable error and meaning?

Connectionism and PDP[8] are an improvement over narrow AI in many respects: (a) they operate at a 'subsymbolic' level (Smolensky, 1988) and accordingly can be said to avoid the symbol grounding problem; (b) neural nets need not be programmed but 'only' trained; (c) they can respond to novel inputs that coalesce into attractor states,[9] whereas standard symbol manipulation approaches are limited to combinations of extant symbols; (d) the operations of neural nets, on the face of it, seem consistent with brain function;[10] and, (e) these systems are also less sensitive to noise than their classical counterparts. However, they lack system detectable meaning because it is clear that any meaning present is available only to an observer, user, or designer of the system, not to the system itself (Bickhard, 1999b; Bickhard & Terveen, 1995). Bickhard adds (1999b, p. 451): "Neither transducers nor nets...are capable of system detectable error concerning what they take to be on the other end of their input correspondences — neither one takes anything to be on the other end of its inputs. Connectionism, in other words, does not address this basic problematic of representation for the system."[11] Therefore, at present, distributed processing models fail on both criteria.

Situatedness and Embodiment

In Bickhard's (1999b) Morse code example above, '...' can stand-in for 'S' because we have an independent way of verifying what 'S' represents. Bickhard and Terveen (1995, p. 15, emphasis original) argue that "a supposed mental encoding of a cup, for example, does not represent the same thing a cup represents — the cup is not a representation at all, and therefore the cup *cannot* be representationally stood-in-for. The cup might be representationally stood-for, but it cannot be representationally stood-in-for." I will offer a kindred perspective on the same problem by considering two ideas from Wittgenstein. The first is roughly derived from his private language argument(s) (Chapman, 1987; Hacker, 1990; ter Hark, 1990; Wittgenstein, 1958), the second from Wittgenstein's notion of a language-game and a corresponding form of life (Hacker, 1990; ter Hark, 1990; Wittgenstein, 1958).

Wittgenstein (1958) argues that symbols take on meanings by the role they play in their respective language-games. A language-game is a set of conventions that governs how linguistic terms and the actions they describe are interwoven (Chapman, 1987). Wittgenstein (1958, para. 23, emphasis removed) noted that "the term 'language-game' is meant to bring into prominence the fact that speaking a language is part of an activity, or a form of life." Forms of life, then, are the practices within which language-games are embedded. Müller (1999) puts it thusly: "[a] form of life is the common ground which we must share with other people in order to understand their actions and words and in order to come to an understanding of them in our judgments." Wittgenstein demonstrated that an agent needs public (outer rather than inner) criteria to come to understand what a given object, utterance, gesture, picture, and the like might signify. Chapman (1987, p. 105) notes that "criteria are those publicly observable circumstances which might be used in teaching the correct use of [an] expression to a child or someone else learning our language." If there were no public criteria for the correct use of words, then agents could never agree that were using terms correctly. Communication would therefore become impossible. Wittgenstein (1958) showed that introspection (hence, private language) does not allow for system detectable meaning or error — in both cases there is no way of verifying whether a symbol represents what it is taken to represent. Agents must accordingly come to know what things and words mean by looking at how they are used within a language community.

Language-games show that understanding meaning is not 'just' a matter of acting upon symbols or attractor regions[12] (if indeed such systems even took themselves to be so doing, which they do not); rather, language is a

8. I will assume that readers are sufficiently familiar with connectionist models for me to not have to describe them here. I direct interested readers to, for example, Rumelhart (1989).

9. Bickhard and Terveen (1995) in fact note that PDP systems allow for input pattern differentiators that are emergent in nature (see below).

10. Bickhard (2000; Bickhard & Terveen, 1995) points out though that the brain is infinitely more complex than these models. However, this may prove to be a matter of constructing richer models rather than a substantive problem.

11. To date, connectionist models also operate on the problematic basis of correspondence notions of meaning and representation (see below).

12. Dartnall (1999, p. 222) argues that "connectionism typically treats...inner states as representational, whether local or distributed."

means within which those involved in a particular language-game can follow the conventions of what is expected regarding the use of a term. Chapman (1987, p. 106, emphasis removed) notes that "in contrast to the common defining characteristics of formal concepts, the criteria of many natural language usages are neither necessary nor sufficient in themselves for justifying the use of the corresponding expressions. A given criterion is not necessary for a given usage, because there might be alternative criteria as well. But neither is a given criterion sufficient in itself for justifying the expression, since it provides justification only under normal conditions." Put another way, the same referent can have different criteria, and different referents can have the same criteria because they may be embedded within different language-games (see Heil, 1981; Goldberg, 1981; ter Hark, 1990). This argument should not to taken trivially as problematic only for constructing systems that can use language; it shows that a correspondence notion of meaning is untenable. As Heil observes, "a gap remains between the thing itself (whatever it is) and its application to states of affairs" (Heil, 1981, p. 333). Both narrow and wide AI models are on shaky foundations to the extent that they rely on correspondence notions of representation because meaning is a fundamentally *situated construct.*[13]

In addition to situatedness, the glue that holds many of the naturalized models together is their reliance on an embodied (rather than disembodied) approach to the problem of modelling intelligence. Merleau-Ponty (1945/1962) claimed it is impossible to speak about the body and of life in general, but only, for example, the animal body and animal life, or the human body and of human life. Lakoff and Johnson (1999) have developed to this idea to argue that the human capability to understand or abstractly reason is attributable to particularly human concrete bodily experiences rather than the other way around. As such, 'high-level' intelligence is parasitic on embodiment.

They argue that the form of our bodies is critical to the representations that we develop and use for both our thought and language. This has spurred some AI workers to use robotic agents to test their theories about intelligence.

Although in this paper I will use 'embodiment' in the same manner as the theorists reviewed below, there may be reason to be wary of the working definition implicit in Bickhard, Brooks and Drescher. The crux of it is that embodiment in AI is assumed to be tantamount to building machines that, in some sense, 'have bodies' and possess little in the way of top-down programming constraints. One of the world's foremost Wittgensteinian scholars has cautioned, however, "machines, unlike living creatures, do not have a body, although they are bodies" (Hacker, 1990, p. 169). Why? "Thinking is a capacity of the animate, manifest in the behaviour and action characteristic of its form of life" (Hacker, 1990, p. 170). The form of life in question, of course, includes acting and behaving as we do, but also includes the behavioural manifestation of perception, desire, emotion, pleasure and pain (in addition to thought). Clearly, machines to this point cannot fully participate in this form of life, and perhaps should not be considered embodied agents. Another way of putting this is that robotic bodies are only very rough approximations of human forms at this point. This may limit the extent to which they can participate in 'human' activities. As Wittgenstein put it (1980, para. 563), "What a lot of things a man must do in order for us to say that he thinks." In this sense then, I will gloss over a weighty issue in this paper — namely, whether artificial intelligence seems programmatically possible on any grounds. There are reasons to suspect that it may not be. Two additional and roughly convergent lines of argument that are distinct from (yet complement) the point I made above follow from Heidegger (see Agre & D. Chapman, 1987; Dreyfus, 1992) and Wittgenstein (see Hacker, 1990, ter Hark, 1990; Shanker, 1998).

Notwithstanding the above, a research programme that operates from a premise of embodiment necessary models intelligence by constructing systems possessing bodies (of sorts) that are situated in environments. For the descriptive purposes of this paper, these bodies and environments can be simulated or actual. However, Brooks (1999) argues that modelling intelligence

13.Bickhard (1987) argues that Wittgenstein does not take his (implicit) critique of correspondence theories far enough — and that, in effect, Wittgenstein ends up with a *correspondences* model of meaning. Bickhard's basic point is that Wittgenstein implicitly relies on an encodingist (see below) conception of representation by using 'criteria' to link a representational function of language to the world. Bickhard sees this as derivative of what he calls the 'linguistic idealism' in Wittgenstein's "theory" ('theory' is in scare quotes because Wittgenstein did not advance a theory). I do not have the space to consider Bickhard's comprehensive argument, but I will note that a Wittgensteinian could also argue that a naturalistic programme such as Bickhard's may have fallen prey to the 'biologism fallacy' (see ter Hark, 1990). At the very least, language games and forms of life do establish a *situatedness*, an intractable context dependency, to meaning that will suffice for the present argument.

through robotic agents in real-world environments is the only possible way for truly intelligent behaviour to emerge (Brooks, 1999; see below) — and I think he is right on this point. I also believe that a reliance on embodied agents to model intelligence would come as no surprise to Wittgenstein who appreciated that "only of a living human being and what resembles (behaves like) a living human being can one say: it has sensations; it sees; is blind; hears; is deaf; is conscious or unconscious" (1958, para. 281). However, Hacker (1990, pp. 160-170) uses a line of argument derived more or less from the same quote to argue that Wittgenstein may have therefore concluded that AI is not possible in principle. Although Hacker may be correct, so-called embodied AI models (which Hacker seems not to have been aware of when his volume was published) often rely on the very ideas that Wittgenstein championed. Further, Brooks and Drescher (see below) have constructed agents that satisfy some of the criteria that would need to be met in order to construct a embodied agent of which Wittgenstein probably would have approved — if one suspects, that is, that embodiment is more than building machines that 'have bodies' and little in the way of top-down constraints.

Disembodiment is derivative (or at least reflective) of the empiricist epistemology that was alluded to above wherein meaning is ostensibly independent of agents, waiting, as it were, to be digested whole. Müller et al. (1998, p. 161) argue that "understanding and meaning are 'embodied', that is, contingent upon subjects being embodied agents who are engaged in, or acting upon, the world. Through these embodied interactions, [agents] construct mental structures, which, in turn, organize experience and make further understanding possible." Wittgenstein (1958, 1969) construed actions as primary — that is, prior to words. From this point of view, meanings "emerge at the level of the whole person's context-embedded deeds" (McDonough, 1999, p. 179). Meaning therefore is not only intractably situated; it is also parasitic upon embodiment. In fact, one could easily turn this around to argue that if meaning is inherently situated and embodied, a "disembodied mind" is logically impossible (see Dartnall, 1999).

Situated Cognition

In the introduction of Clancey's (1997) *Situated Cognition*, he states that "the theory of situated cognition...claims that every human thought and action is adapted to the environment, that is, *situated*, because what people *perceive*, how they *conceive of their activity*, and what they *physically do* develop together" (emphasis original). A decade earlier, Agre and D. Chapman (1987) emphasised the need to construct artificial agents capable of *leaning on the world* (versus consulting an explicit internal model of the world) in order to model intelligence. For the sake of economy, I will limit myself to a consideration of Brooks' (1986/1999, 1991/1999) model[14] (but see also, e.g., Agre & D. Chapman, 1987; Beer, 1995; Clancey, 1997). Brooks (1986/1999, 1991/1999) has built robotic agents by eschewing explicit representation altogether (see also Keijzer, 1998). As I am only interested in the assumptions in Brooks's account, I have not described his more recent Cog project in any detail. Those interested should consult, for example, Brooks, Breazeal, Irie, Kemp, Marjanovic, Scassellati and Williamson (1998).

Nilsson (1984) describes the mobile robot, Shakey, that was constructed in the late 1960s. Shakey was able to navigate aspects of its environment. However, this occurred as a result of precise control over the environment in which this machine moved and the instructions that were fed into it. Brooks (1991/1999) points out that no research team had been able to replicate Shakey's 'feats' in an unstructured environment. Brooks, however, has managed to do so with his *subsumption architecture* (wherein a simpler structural-functional layer is subsumed into a more complex one, and so on). He uses primitive, but embodied and complete, systems with the goal of developing complexity incrementally wherein each layer of the architecture tightly couples perception to action.

Brooks (1991/1999, p. 138-139) describes four ideas (the first two of which should be familiar by now) that demarcate his work (and the work of those whom I shall describe below) from that of standard AI and robotics: (a) *Situatedness* — the behaviour of his robots results from a confrontation with "the here and now of the environment" (not abstract descriptions of it); (b) *Embodiment* — his robots experience the world directly through their bodies wherein the agents' actions immediately feedback on their sensations; (c) *Intelligence* — the source of intelligence is not limited to the computational engine, but comes jointly from the world, the signal transformations within sensors, and the physical coupling of the robot to the world (cf. McDonough, 1999); and, (d) *Emergence* — systemic intelligence emerges

14.Brooks might object to being included under the category of 'situated cognition', but this is a reasonable characterization for what I am trying to accomplish here.

from interactions with the world and occasionally from indirect interactions amongst the agents' components.

With respect to the last point, Brooks (1991/1999, p. 139) states that "it is sometimes hard to point to one event or place within the system and say that is why some external action was manifested." It is clear that "this new approach to robotics makes claims about how intelligence should be organized that are radically different from the approach assumed by traditional AI" (Brooks, 1991/1999, p. 61). What is not as clear is how Brooks' model and architectures relate to the necessary criteria for the ascription of intelligence that I have delineated. I want to first point out though that the title of Brooks' (1999) volume, *Cambrian Intelligence,* tips his hand with respect to his emphasis. His robotic agents are designed to instantiate what has occurred in our evolutionary history. Intelligence in the primitive (*pre-vertebrate* Cambrian era) sense is what Brooks has tried to model — with the assumption that more complex higher-order operations emerge out of such lower forms by subsuming them (hence, his term subsumption architecture).

Although Brooks rejects traditional AI representation schemes (including, perhaps problematically, goal representation), he allows for representations that are "partial models of the world" (1991/1999, p.177).[15] These are not standard explicit correspondence representations, and as such Brooks avoids the incoherence of correspondence models. These representations can be in error (see e.g., Brooks, 1986/1999) and these agents accordingly have system detectable error. Due to their prehistoric status, they seem, however, to lack system detectable meaning (with respect to understanding, for example, what a symbol might mean). Hofstadter (1995) claimed above that any representations in such a low-level perceptual system would be too primitive to be considered meaningful — presumably to Hofstadter, that is, not to the agent — and I agree with the following caveat. Although such agents cannot 'detect' higher-order meaning, they may be able to classify environments in a way that has lower-order 'meaning' (i.e., significance) to the system (with respect, for example, to indicating future interactive outcomes, see below). In fact, these lower-level perceptual experiences seem phylogenetically required before higher-order (culturally-derived) meanings could emerge. This suggests that Hofstadter's (1995) dichotomy is misleading one when applied to situated and embodied agents.

15.Bickhard (Bickhard and Terveen, 1995) rightly point out that Brooks' robots would be implicitly representational on Bickhard's terms. See below.

Constructivist AI

Drescher takes his cue from Piagetian theory[16] by designing a model to roughly instantiate Piaget's sensorimotor stage of development. He argues that investigating infants is a better way to study the situated emergence of intelligence than "Brooks' insect-robots" (Drescher, 1991, p. 192). Brooks has of course since moved on to do 'humanoid robotics' in his Cog project (see e.g., Brooks et al., 1998). Although Drescher's (1991) model is situated and embodied, he uses a simulated agent that makes contact with a simulated (micro) world. As noted above, Brooks (1991/1999) has argued that the real world poses problems for an agent that a simulated one cannot (due the constraints of the designer). Drescher (1991) argues however that real-world domains can have similar problems, and (1991, p. 192-193) that the utility of either approach is an empirical matter. Drescher, however, relies on a connectionist architecture to instantiate his model (and has accordingly inherited the strengths and weaknesses of this approach). Although satisfying the joint criteria I allude to above would therefore be impossible given the constraints of current connectionist modelling (including Drescher's), the assumptions guiding his model are notable. Bickhard's interactivism model, which is sketched next, is also roughly inspired by Piagetian principles. Dresher's (1991) constructivist AI model makes for both a good introduction to Piaget and a fitting counterpoint to interactivism (see below).

Piaget (1954) holds that an infant begins life with only basic reflexes. The infant *constructs* knowledge of the world somewhat like Brooks' robots (because such knowledge cannot be present *a priori* or built up *passively* with experience, see Müller et al., 1998). Like Brooks' principle of subsumption, this construction builds upon and incorporates previous structures, leading ultimately (in the case of Drescher's simulated infant) to the construction of a permanent object. Piaget's schemas[17] roughly indicate that a certain action in a specific context will have a particular outcome. To make this more transparent, Drescher (1991, p. 4,

16.As opposed to the punctuated equilibria theory of Gould and Eldredge (1977) wherein Brooks (1991/1999, p. 139) claimed to find his inspiration.

17.The term 'schema' is a translation error in the source that Drescher relied upon. Piaget actually intended 'scheme' not 'schema', but I have left it in Drescher's words.

emphasis removed) notes that "the [schema] mechanism...follows what we might call a prediction-value paradigm, in contrast with a situation-action paradigm: the mechanism does not directly learn what action to take in a given situation, but rather what would happen next for each of several possible actions. It may then select what action to take based in part on the value of an achievable result."

By incorporating Piagetian principles into his model, Drescher enables his machine to construct concepts that are qualitatively different from extant ones (despite the fact that the novel concept emerged in some sense out of them). Drescher's (1991) machine constructs a new concept by creating an element called a *synthetic item.* This is analogous to Piaget's (1954) notion of conservation wherein an agent hypothesizes an invariant mental structure in order to anchor changing (or absent) elements of experience. In contrast to symbol manipulation accounts, this allows new concepts to emerge: "In the interests of being able to transcend *a priori* domains (and in the interests of modelling Piaget's theory), the schema mechanism starts without...knowledge. This way, from the outset, the mechanism demonstrates the ability to learn in unprecedented domains — since, to the mechanism, all domains are unprecedented" (Drescher, 1991, p. 5). Despite the limits of his connectionist architecture, Drescher's work may again make it possible for epistemic systems to have low-level 'meaning' (in the sense of schematic significance, Piaget, 1954) and high-level meaning in a way that can emerge *ontogenetically* in principle, as opposed to assuming that it must phylogenetically (cf. Brooks, 1999; Hofstadter, 1995).[18]

Interactivism

Although there are many points and subtleties to interactivism, I want to merely demonstrate how this model avoids the problems associated with the above. I direct interested readers to the works I have cited. Those interested in constructing agents consistent with interactivist principles should consult Bickhard and Terveen (1995) in particular. The mathematical theories upon which interactivism is based and the implementations of design criteria are addressed in this work in detail.

Bickhard's (1993, 1996, 1998, 1999b; Bickhard & Terveen, 1995) interactivism model is consistent with the situationist account outlined above — except for the rather crucial difference that Bickhard affords a central role to representation in his model (that is based on a radically different set of assumptions than those of standard AI accounts). Bickhard and Terveen (1995) argue, in fact, that the situationist critique of representation is only valid to the extent that it is aimed at an encodingist notion of representation. Interactivism in also consistent to a certain extent with constructivist AI, but Bickhard and Terveen (1995, p. 278-279, emphasis removed) note that "Drescher's items always 'represent' environmental conditions." This is not a surprising concern coming from Bickhard because he levels the same criticism at Piagetian theory in general (see Bickhard & Campbell, 1989). Despite the similarities between situated cognition and constructivist AI, Bickhard concludes that constructivist AI, therefore, falls into the very encodingist circularity and incoherence of which situationism is a critique. Although Bickhard critiques encoding*ism,* he is not critical of encodings or symbolic representation in and of itself. Rather, he argues that symbolic representation emerges from the *originally non-representational* actions of an agent.

In effect, Bickhard's model is grounded on a non-circular situationist constructivism. Recall that Piagetian schemes indicate that a certain action in a specific context will have a particular outcome. To the extent that Piaget relies on correspondent representations as the *basis* for further encodings[19], he encounters circularity and incoherence. Interactivism is instead predicated on the type of implicit (non-encodable) representation that can be seen, for example, in Brooks (1991/1999). Bickhard explains that "if [a] subsystem ends up in final [internal] state 'A', then it has encountered the type of environment that yields that final state when engaged in interaction, and similarly for final state 'B'...the possible final states implicitly define the classes of environments that would yield them, and actual interactions classify environments among those implicitly defined classes" (Bickhard, 1996, p. 179).

The point is that these environmental classifications (which Bickhard terms 'differentiations') do not lead to representational content regarding the environment — other than indicating, for example, final state 'A' instead of 'B'. Because there is no representational content about the differentiated environments caused by interactive differentiation, there can be no encodings. The logi-

18.Drescher, or course, has not demonstrated that it did or will emerge. However, his *Connection Machine* has behaved in ways that are strikingly consistent with Piaget's (1954) 1st through 6th sensorimotor stage. His architecture managed to accomplish the requisite construction — the possibility of a permanent object.

19.This is a matter open to debate.

cal basis for the non-circularity in Bickhard's (1993, 1996, 1998, 1999b; Bickhard & Terveen, 1995) model, then, is provided by the fact that implicit definition is not a representational phenomenon. It can therefore provide a foundation that allows for system detectable error (and meaning) in a way that avoids encodingism: "interactive differentiators...model what is standardly construed as encodings. But the interactive model makes no claims that these correspondences somehow magically constitute representations of what has been differentiated" (Bickhard, 1996, p. 180). As such, Bickhard's model resolves the problems associated with encodingist/correspondence notions of representation and meaning.[20]

Conclusion

Standard AI models run afoul of designing intelligent systems as a result of correspondence assumptions that are derivative of an empiricist epistemology. In particular, no attempts are made within these frameworks to account for system detectable error, and derivatively, system detectable meaning. Connectionist architectures show promise, but at present still fall into the same dilemma. Naturalized AI models may be able to avoid these problems because they are inherently situated and embodied (to some degree). Situationists share some basic epistemological assumptions with constructivists regarding the modelling of intelligence. For example, both deny that knowledge can be defined independently of a mind. However, constructivists may be more likely to argue that an individual mind creates its own particular knowledge structures in response to the world (i.e., as 'memories') — whereas situationists stress that constructed knowledge does not exist in memory *per se*, but instead emerges from interaction with the environment (Self, 1995). (Narrow AI has, of course, assumed that representations have some psychological reality in that they correspond (typically not in a literal, but a functional sense) to memory structures.) Situationists hold that representations created by practical activities of an agent are not themselves knowledge. Rather, knowledge is more a capacity to interact (Self, 1995).

Interactivism refines the situationist position by framing a central role for representation in a way that avoids the circularity associated with both standard and constructivist AI. Bickhard subsumes situationist models into his framework by pointing out that all of the aforementioned are representational on his notion of (implicit) representation. He points out that some of these models (e.g., Agre & D. Chapman, 1987) are incoherent because they rely on the same correspondence notions as does, Bickhard argues, constructivist AI. As such, interactivism shows promise in modelling intelligence because it does not rely on correspondences to derive meaning and representation and also accounts for system detectable meaning and error. Further, this is a model that explains in a non-circular and parsimonious manner how low-level perception can, in some sense, be meaningful to a system, *and* how these lower-level perceptions are necessary for higher-order meaning (cf. Hofstadter, 1995). Interactivism may offer a viable alternative within which AI can proceed.

Epilogue

I'd like to acknowledge at this point though that outlining Bickhard's model above in some detail was intended to crystallize the problems with correspondence-empiricist-rationalist-nativist approaches to AI. His model also handily fills in some of the gaps in logic that exist in the other two approaches I outlined above.

Rod Brook's humanoid robotics programme, however, basically squares with what Bickhard would want from an AI. That is, Bickhard's (Bickhard & Terveen, 1995) critiques of Brooks are more logical than methodological. Brooks and his collaborators have been working with? on? a robot they named Cog for about the last 5 years at the MIT AI labs (http://www.ai.mit.edu/projects/humanoid-robotics-group/cog/cog.html).

To date, Cog can point, imitate head nods, detect faces, perform reflexive arm withdrawal (like an infant), reach for a target and perform orienting behaviours — with no ontological programmatic commitments on the part of Brooks and his collaborators to incoherent models. I have a few quarrels with some of their moves, but if what I have outlined above is accurate then we may see some progress — probably very slow progress — in the coming years with this and related projects. As intimated above though, the only obvious problem with this approach to AI is that it may be impossible in a technical sense to truly construct an embodied machine that is isomorphic to a human body. Wittgenstein and others imply a close approximation will not be enough. It is conceivable that scientific advances in constructing humanoid bodies (rather than minds) will solve this problem.

20. It should also be apparent that interactivism is another model that suggests a way out of the observation made above by Hofstadter (1995).

References

Agre, P. E. & Chapman, D. (1987). Pengi: An implementation of a theory of activity. *Proceedings of the Sixth National Conference on Artificial Intelligence* (pp. 196-201). Seattle, WA.

Beer, E. (1995). A dynamical systems perspective on environment agent interactions. *Artificial Intelligence, 72,* 173-215.

Bickhard, M. H. (1987). The social nature of the functional nature of language. In M. Hickmann (Ed.), *Social and functional approaches to language and thought* (pp. 39-65). New York: Academic.

Bickhard, M. H. (1993). Representational content in humans and machines. *Journal of Experimental and Theoretical Artificial Intelligence, 5,* 285-333.

Bickhard, M. H. (1996). Troubles with computationalism. In W. O'Donohue & R. F. Kitchener (Eds.), *The philosophy of psychology* (pp. 173-183). London: Sage.

Bickhard, M. H. (1998). Levels of representationality. *Journal of Experimental and Theoretical Artificial Intelligence, 10,* 179-215.

Bickhard, M. H. (1999a). Concepts: Where Fodor went wrong. *Philosophical Psychology, 12,* 5-24.

Bickhard, M. H. (1999b). Interaction and representation. *Theory and Psychology, 9,* 435-458.

Bickhard, M. H. (2000). Dynamic representing and representational dynamics. In E. Dietrich & A. B. Markman (Eds.), *Cognitive dynamics: Conceptual and representational change in humans and machines* (pp. 31-50). Mahwah, NJ: Erlbaum.

Bickhard, M. H. & Campbell, D. T. (1989). Interactivism and genetic epistemology. *Archives de Psychologie, 57,* 99-121.

Bickhard, M. H. & Terveen, L. (1995). F*oundation issues in artificial intelligence and cognitive science: Impasse and solution.* Amsterdam: Elsevier.

Brooks, R. A. (1999). *Cambrian intelligence: The early history of the new AI.* Cambridge, MA: MIT. (Original articles published 1986-1991.)

Brooks, R. A., Breazeal, C., Irie, R., Kemp, C. C., Marjanovic, M., Scassellati, B. & Williamson, M. M. (1998). Alternate essences of intelligence. In *Proceedings of the American Academy of Artificial Intelligence* (AAAI-98).

Canfield, J. V. (1999). Folk psychology versus philosophical anthropology. *Idealistic Studies, 29,* 153-172.

Chapman, M. (1987). Inner processes and outward criteria: Wittgenstein's importance for psychology. In Chapman & R. A. Dixon (Eds.), *Meaning and the growth of understanding* (pp. 103-127). Berlin: Springer-Verlag.

Chapman, M. (1999). Constructivism and the problem of reality. *Journal of Applied Developmental Psychology, 20,* 31-43.

Choi, S. (1997). Language-specific input and early semantic development: Evidence from children learning Korean. In D. I. Slobin (Ed.), *The crosslinguistic study of language acquisition: Expanding the contexts* (Vol. 5, pp. 41-133). Mahwah, NJ: Erlbaum.

Chomsky, N. (1964). On the nature, use and acquisition of language. In A. Goldman (Ed.), *Readings in philosophy and cognitive science* (pp. 511-534). Cambridge, MA: MIT Press.

Clancey, W. (1997). *Situated cognition: On human knowledge and computer representations.* Cambridge: Cambridge University Press.

Dartnall, T. (1999). Normative engines: Stalking the wild Wittgensteinian machine. *Idealistic Studies, 29,* 215-230.

Drescher, G. L. (1991). *Made-up minds: A constructivist approach to artificial intelligence.* Cambridge, MA: MIT Press.

Dreyfus, H. L. (1992). *What computers still can't do: A critique of artificial reason.* Cambridge, MA: MIT Press.

Fodor, J. A. (1975). *The language of thought.* Cambridge, MA: MIT Press

Fodor, J. A. (1981). *RePresentations: Philosophical essays on the foundations of cognitive science.* Cambridge, MA: MIT Press.

Fodor, J. A. (1983). *The modularity of mind.* Cambridge, MA: MIT Press

Fodor, J. A. (2000). *The mind doesn't work that way: The scope and limits of computational psychology.* Cambridge, MA: MIT Press.

Goldberg, B. (1991). Mechanism and meaning. In J. Hyman (Ed.), *Investigating Psychology* (pp. 48-66). London: Routledge.

Gould, S. & Eldredge, N. (1977). Punctuated equilibria: The tempo and mode of evolution reconsidered. *Paleobiology, 3,* 115-151.

Hacker, P. M. S. (1990). *Wittgenstein, meaning and mind* (An analytic commentary on the Philosophical investigations; vol. 3). Cambridge: Blackwell.

Hark, M. ter (1990). *Beyond the outer and the inner: Wittgenstein's philosophy of psychology.* Dordrecht, The Netherlands: Kluwer.

Haugeland, J. (1985). *Artificial intelligence.* Cambridge, MA: MIT Press.

Heil, J. (1981). Does cognitive psychology rest on a mistake? *Mind, 90,* 321-342.

Hofstadter, D. & The Fluid Analogies Research Group (1995). *Fluid concepts and creative analogies: Computer models of the fundamental mechanisms of thought.* New York: Basic.

Holyoak, K. J. & Hummel, J. E. (2000). The proper treatment of symbols in a connectionist architecture. In E. Dietrich & A. B. Markman (Eds.), *Cognitive dynamics: Conceptual and representational change in humans and machines* (pp. 229-264). Mahwah, NJ: Erlbaum.

Hooker, C. A. (1996). Towards a naturalized cognitive science. In W. O'Donohue & R. F. Kitchener (Eds.), *The philosophy of psychology* (pp. 184-208).

Jordan, M. J. & Rumelhart, D. E. (1992). Forward models: Supervised learning with a distal teacher. *Cognitive Science, 16,* 307-354.

Keijzer, F. A. (1998). Doing without representations which specify what to do. *Philosophical Psychology, 11,* 269-302.

Lakatos, I. (1970). Falsification and the methodology of scientific research programmes. In I. Lakatos and A. Musgrave (Eds.), *Criticism and the growth of knowledge* (pp. 231-278). Cambridge: Cambridge University Press.

Lakoff, G. & Johnson, M. (1999). *Philosophy in the flesh: The embodied mind and its challenge to Western thought.* New York: Basic.

McDonough, R. (1999). Bringing cognitive science back to life. *Idealistic Studies, 29,* 173-214.

Merleau-Ponty, M. (1962). *The phenomenology of perception* (trans C. Smith). New York: Humanities Press. (Original work published 1945)

Müller, U. (1999). *Wittgenstein and Piaget: A critique of mechanistic theories of meaning.* Paper presented at the 29th Annual Symposium of the Jean Piaget Society, Mexico City, Mexico.

Müller, U., Sokol, B. & Overton, W. F. (1998). Reframing a constructivist model of the development of mental representation: The role of higher-order operations. *Developmental Review, 18,* 155-201.

Newell, A. & Simon, H.A. (1976). Computer science as empirical inquiry: Symbols and search. *Communications of the ACM, 19,* 113-126.

Nilsson, N. J. (Ed.) (1984). *Shakey the robot.* Stanford Research Institute AI Center, Technical note 323.

Piaget, J. (1954). *The construction of reality in the child.* New York: Ballentine.

Quartz, S. (2001). Anything to celebrate? *Trends in Cognitive Science, 5,* 223-224.

Rumelhart, D. E. (1989). The architecture of mind: A connectionist approach. In M. I. Posner (Ed.), *Foundations of cognitive science* (pp. 133-160). Cambridge, MA: MIT Press.

Russell, B. (1921). *The analysis of mind.* London: Allen & Unwin.

Ryle, G. (1949). *The concept of mind.* London: Hutchinson.

Searle, J.R. (1981). Minds, brains, and programs. In J Haugeland (Ed.), *Mind design* (pp. 282-306). Cambridge, MA: MIT Press.

Self, J. (1995). *Computational mathetics: Towards a science of learning systems design.* Unpublished manuscript. University of Leeds, Leeds, UK.

Shanker, S. (1998). *Wittgenstein's remarks on the foundations of AI.* New York: Routledge.

Shannon, B. (1998). Fodor's impasse — A converse perspective and a way out. *Journal for the Theory of Social Behaviour, 28,* 130-145.

Simon, H. A. (1969). *The sciences of the artificial.* Cambridge, MA: MIT Press.

Simon, H. A. (1996). Computational theories of cognition. In W. O'Donohue & R. F. Kitchener (Eds.), *The philosophy of psychology* (pp. 160-172). London: Sage.

Slobin, D. I. (1997). The origins of grammaticizable notions: Beyond the individual mind. In D. I. Slobin (Ed.), *The crosslinguistic study of language acquisition: Expanding the contexts* (Vol. 5, pp. 265-323). Mahwah, NJ: Erlbaum.

Smolensky, P. (1988). On the proper treatment of connectionism. *Behavioral and Brain Sciences, 11,* 1-74.

Turing, A. (1950). Computing machinery and intelligence. *Mind, 59,* 433-460.

Turing, A. (1970). Intelligent machinery. In B. Meltzer & D. Michie (Eds.), *Machine intelligence 5* (pp. 2-23). New York: Elsevier.

Wittgenstein, L. (1958). *Philosophical investigations* (3rd ed., trans G. E. M. Anscombe). Englewood Cliffs, NJ: Prentice-Hall.

Wittgenstein, L. (1969). *On certainty* (Eds. G. E. M. Anscombe and G. H. Von Wright, trans D. Paul and G. E. M. Anscombe). New York: Harper & Row.

Wittgenstein, L. (1980). *Remarks on the philosophy of psychology* (vol. 1, Eds. G. E. M. Anscombe and G. H. Von Wright, trans by G. E. M. Anscombe). Oxford: Blackwell.

Acknowledgements

The preparation of this article was supported by a Social Sciences and Humanities Research Council of Canada Doctoral Fellowship. I would like to thank Jeremy Carpendale and Bill Turnbull for their helpful comments on earlier drafts of this article.

About the Author

Timothy Racine is a doctoral candidate in the Department of Psychology, Simon Fraser University, Burnaby, BC V5A 1S6, Canada. His interests include the theoretical and empirical issues related to the construction of mind in human infants and robotic agents. Correspondence concerning this article should be sent via electronic mail to <timothy_racine@sfu.ca>, or addressed to Timothy P. Racine, Department of Psychology, Simon Fraser University, Burnaby BC V5A 1S6.

VISION INTERFACE (VI 2002)

May 27-29, 2002
Calgary, Alberta, Canada

http://www.visioninterface.org/vi2002

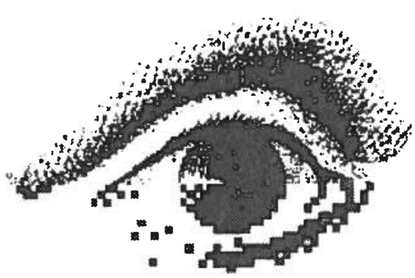

Graphics Interface 2002

May 27-29, 2002
Calgary, Alberta, Canada

For the latest information on the 2002 conference, visit: www.graphicsinterface.org

CSCSI
SCEIO

AI 2002

The Fifteenth Canadian Conference on Artificial Intelligence

Sponsored by Canadian Society for Computational Studies of Intelligence (CSCSI)
Société Canadienne pour l'étude de l'intelligence par ordinateur (SCEIO).

27-29 May, 2002

University of Calgary,
Calgary, Alberta, Canada

http://www.cs.unb.ca/ai2002/

AI'2002, the Fifteenth Canadian Conference on Artificial Intelligence, is a forum that showcases current research in artificial intelligence from researchers around the world. The conference features research paper presentations and invited speakers. AI'2002, the (now annual) meeting of the Canadian Artificial Intelligence community, will be held in conjunction with GI'2002 (Graphics Interface) and VI'2002 (Vision Interface).

Conference Chair
Ali Ghorbani
Faculty of Computer Science
University of New Brunswick
Fredericton, NB, E3B 5A3
Canada

Program Co-Chairs
Robin Cohen (University of Waterloo)
Bruce Spencer (NRC and University of New Brunswick)

The Eighteenth National Conference on Artificial Intelligence

Fourteenth Innovative Applications of AI Conference

July 28 - August 1, 2002
Edmonton, Alberta, Canada

AAAI-2002

The Fourteenth Innovation Applications of Artificial Intelligence Conference and the Eighteenth National Conference on Artificial Intelligence will be held 28 July - 1 August, 2002, at the Shaw Conference Centre and the Westin Edmonton Hotel, in Edmonton, Alberta, Canada. Further information on the two conferences will be made available at:

http://www.aaai.org/Conferences/National/2002/aaai02.html

CG 2002

COMPUTERS AND GAMES 2002

Third International Conference on
Computers and Games
Edmonton, Canada, July 25-27 2002
Email: cg2002@cs.ualberta.ca
URL: http://www.cs.ualberta.ca/cg2002

CALL FOR PAPERS

The biennial *Computers and Games* conference series is a major international forum for researchers and developers interested in all aspects of artificial intelligence in computer game playing. After two terms in Japan, the conference will be held in North America for the first time. Starting this year, we are extending the scope to include commercial games such as sports games and role-playing games.

TOPICS OF INTEREST:

Relevant topics include, but are not limited to:

- the current state of game-playing programs,
- new theoretical developments in game-related research,
- general scientific contributions produced by the study of games,
- AI techniques applied to games, such as machine learning, heuristic search, knowledge representation, data-mining, and path finding,
- social aspects of computer games,
- cognitive research on how humans play games, and
- issues related to networked games.

PAPER SUBMISSION:

The conference proceedings will be published by Springer in the *Lecture Notes in Computer Science Series* (http://www.springer.de/comp/lncs/index.html). All submitted papers are refereed. Accepted papers will be presented at the conference and printed in the conference proceedings. Please visit the CG'02 web page for detailed instructions about the submission procedure and the layout of the papers.

IMPORTANT DATES:

April 1, 2002	Deadline for paper submissions.
May 15, 2002	Accept/Reject notifications sent to authors.
June 15, 2002	Early registration deadline.

PROGRAM CHAIRS:

Jonathan Schaeffer (University of Alberta, Canada)
Martin Müller (University of Alberta, Canada)

General Chair:
Moonis Ali, USA

Program Chair:
Paul Chung, UK

Program Co-Chair:
Chris Hinde

Local Organization Chair:
Christian Dawson

Web Site Chair:
Shuanghua Yang

Program Committee:
Chimay Annumba, UK
Ruth Aylett, UK
Senen Barro, Spain
Fevzi Belli, Germany
Ivan Bratko, Slovenia
Edmund Burke, UK
CW Chan, HK
Kai Chang, USA
Roberto Desimone, UK
Gion Dini, Italy
Tadeusz Dobrowiecki, Hungary
Graham Forsyth, Australia
Christian Freksa, Germany
Perttu Heino, Finland
Tim Hendtlass, Australia
Mohamed Jmaiel, Tunisia
Catholijn Jonker, Netherland
Mark Lee, UK
Andras Markus, Hungary
Steve McCoy, UK
Laszlo Monostori, Hungary
Derek Partridge, UK
Don Potter, USA
Henri Prade, France
Chris Price, UK
David Robertson, UK
Padro Sanz, Spain
Michael Schroeder, UK
Fabrizio Sebastiani, Italy
Flavio Soares Correa da Silva, Brazil
Kazuhiko Suzuki, Japan
Liansheng Tan, China
Takushi Tanaka, Japan
Spyros Tzafestas, Greece
Sunil Vadera, UK
Xuezhong Wang, UK


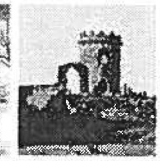


IEA/AIE-2003
CALL FOR PAPERS

The Sixteenth International Conference on Industrial & Engineering Applications of Artificial Intelligence & Expert Systems

(http://gradients.lboro.ac.uk/iea2003aie)

June 23-26, 2003
Loughborough, UK
Sponsored by:
International Society of Applied Intelligence
Organized in Cooperation with:
AAAI, ACM/SIGART, CSCSI/SCEIO, ECCAI, ENNS, IEE, INNS, JSAI, LU, SWT

IEA/AIE continues the tradition of emphasizing applications of artificial intelligence and knowledge-based systems to engineering and industrial problems as well as to other areas. Topics of interest include, but are not limited to:

- Adaptive Control
- Applications to Manufacturing
- Autonomous Agents
- Applications to Design
- Case-based Reasoning
- Computer Vision
- Data Mining
- Distributed Problem Solving
- Expert Systems
- Fuzzy Logic
- Genetic Algorithms
- Genetic Programming
- Heuristic Search
- Internet Applications
- Intelligent Interfaces
- Intelligent Systems
- KBS Methodology
- Knowledge Processing
- Knowledge Management
- Machine Learning
- Model-based Reasoning
- Natural Language Processing
- Neural Networksz
- Planning and Scheduling
- Reasoning Under Uncertainty
- Robotics
- Spatial Reasoning
- System Integration
- Speech Recognition
- Temporal Reasoning

Authors are invited to electronically submit their paper, written in English, of up to 10 single spaced pages, presenting the results of original research or innovative practical applications relevant to the conference. Practical experiences with state-of the-art AI methodologies are also acceptable when they reflect lessons of unique value to the conference attendees. Shorter works, up to 6 pages, to be presented in 10 minutes, may be submitted as SHORT PAPERS representing work in progress or suggesting possible research directions. (Please, indicate "short paper" in the submission letter in this case.) **Paper submission should be sent by email to iea2003aie@lboro.ac.uk either in MS Word or PDF format by November 12, 2002. All papers should include a key word list containing one or more of the topics listed above.** ALL papers will be formally peer reviewed. Notification of the results of the review process will be made by the end of January 2003, and final copies of papers for inclusion in the conference proceedings will be due by February 21, 2003. The conference proceedings, published in a bound volume by Springer-Verlag in their 'Lecture Notes in Artificial Intelligence' series, will be available at the conference. Referees will be asked to nominate papers for a Best Paper award to be announced at the conference. All papers, but particularly those nominated for the Best Paper competition, will be automatically considered for publication in an expanded form in the International Journal of Applied Intelligence.

Workshop:
Workshop proposals are invited. For further details, consult the Program Chair or see the conference web site.
Conference web site:
http://gradients.lboro.ac.uk/iea2003aie

Professor Moonis Ali
General Chair of IEA/AIE-2002
Southwest Texas State University
Department of Computer Science
601 University DriveSan Marcos TX 78666-4616 USA
Fax: (+1) 512-245-8750
Email: ma04@swt.edu

Professor Paul Chung
Program Chair of IEA/AIE-2003
Department of Computer Science
Loughborough University
Loughborough
Leicestershire LE11 3TU
England, UK
Phone: +44 (0) 1509 222543
Fax: +44 (0) 1509 211586
Email: p.w.h.chung@lboro.ac.uk

PRECARN UPDATE
NOUVELLES DE PRECARN

The Canadian intelligent systems (IS) industry is at an interesting point in its evolution. Research has been going on in Canada for decades and while public interest has always been evident, at no other time moreso than right now, people are starting to take a concentrated interest in this field of technological development – researchers, governments, even our friends and neighbours.

Applications of IS technologies are flourishing – from nickel mining to data mining, from factory robots to robotic surgery, from biosciences to biometrics and from telerobotics to telemedicine. Brought into sharp focus by the events of September 11, safety and security are application areas that are receiving more attention these days.

As Canada's leader in collaborative intelligent systems R&D, Precarn Incorporated has a unique perspective on these trends. Universities, government research organizations and private companies are working together under the program to respond to the needs of a changing environment. Already some projects are underway:

Research in predictive maintenance for looking at transients in a power system is one such example. A project led by Powertech Labs of Surrey, BC, is developing a system to recognize that as certain variables occur in sequence or pattern, there is need for action or preventative measures to be taken.

Development of a diagnostic and monitoring system for nuclear reactors is another example. The research, led by Ontario Hydro in Toronto, ON, looked into the accurate detection of failures in the cooling system and control systems to prevent a dangerous situation from progressing.

Extrapolation technologies, taking hundreds of variables to determine the best approach, are also being developed in a project led by Hatch Associates Ltd., Advanced Systems Group, Mississauga, ON, with the prototype being designed for the scheduling of a truck fleet at a surface mine. Further, several companies and research teams have approached Precarn with innovative project ideas that may offer solutions to pressing security issues.

Precarn is a national consortium of corporations, research institutes and government partners working within the intelligent systems industry. Precarn also manages IRIS (the Institute for Robotics and Intelligent Systems), one of the federal government's Networks of Centres of Excellence. Expanding the program, IRIS has just received an additional $13.1 million to fund new and continuing projects.

Researchers within the IRIS network are developing intelligent software agents capable of complex reasoning. They are also looking at data mining – analyzing large amounts of data from structured and unstructured sources. Applications for these technologies are far-reaching, from providing remote database access to information at a job site, to suggesting the best possible road on which to travel, given variables, such as perhaps weather, traffic and even taking into consideration unsafe war zones in a ground battle. Further, IRIS researchers are looking at image matching and analysis, another area where intelligent systems are developing at a rapid pace.

There are hundreds of applications for the technologies being developed. Precarn and IRIS and other research organizations are working to ensure that the Canadian development in these areas reaches its potential, both within our borders and in the world markets.

Precarn and IRIS will hold their 12th Annual Canadian Conference on Intelligent Systems in Calgary on May 29-31, 2002. For more information on becoming involved in the network or for conference information, visit www.precarn.ca.

Intelligent Systems. Thinking Technology.
Systèmes Intelligents. Penser technologie.

CSCSI/SCEIO Membership Application

I wish to join CSCSI/SCEIO and receive *Canadian Artificial Intelligence /Intelligence Artificielle au Canada*

Yes ___ / No ___ Web Access Only ($30.00* Cdn./yr.)

Yes ___ / No ___ Printed Copy and Web Access ($40.00* Cdn./yr.)

I am a student

Yes ___ / No ___ Web Access Only ($15.00* Cdn./yr.)

Yes ___ / No ___ Printed Copy and Web Access ($15.00* Cdn./yr.)

I am a member of CIPS

Yes ___ / No ___ Web Access Only ($25.00* Cdn./yr.)

Yes ___ / No ___ Printed Copy and Web Access ($30.00* Cdn./yr.)

*Includes applicable G.S.T.

Name ....

Job Title ....

Company ....

Address ....

....

....

....

E-mail ....

Phone ....

Fax ....

Visa or Amex Number ....

Expiry Date ....

Signature ....

Please send your membership application to:

Mail: CIPS National Office, One Yonge St., Suite 2401, Toronto, Ontario M5E 1E5

Phone: (416) 861-2477

Canadian Artificial Intelligence
Intelligence Artificielle au Canada

The editor welcomes contributions from members and non-members. Possible formats include, but are not limited to:

- commentary, in the form of letters or short articles, on any AI issue of concern to you (e.g. research priorities, funding, educational trends, cultural, political, or ethical aspects of AI)
- interdisciplinary work
- survey papers
- reports on new work by students
- abstracts
- web site reviews
- news
- humour
- employment opportunities
- conference announcements or reviews.

Submission information • Guide de soumission

Canadian Artificial Intelligence / Intelligence Artificielle au Canada welcomes submissions on any matter related to artificial intelligence. Please send your contribution, preferably in electronic form, with an abstract and biography to:

Ann Grbavec
Editor, *Canadian AI / IA au Canada*
School of Computing Science
Simon Fraser University
Burnaby, British Columbia, Canada
V5A 1S6 or — agrbavec@cs.sfu.ca
Phone (604) 291-3208
Fax (604) 291-3045

Please contact the editor as early as possible to obtain complete submission guidelines. Sample issues and articles can be viewed at: http://cscsi.sfu.ca/cai.html

Advertising rates are available upon request from the editor.

Recycled / Recyclable