Can Situated Robots Play Soccer?

To appear in Canadian AI-94

Michael K. Sahota and Alan K. Mackworth
Laboratory for Computational Intelligence
Department of Computer Science
University of British Columbia
Vancouver, B.C., Canada, V6T 1Z4
sahota@cs.ubc.ca, mack@cs.ubc.ca

Abstract

The goal of creating an integrated cognitive robot is still only a tantalizing dream. Current artificial intelligence and robotics research is highly divergent with little or no commonality among specialized subfields. New rich task domains are needed to pose the right challenges to extant theories and promote convergence. We propose soccer-playing as such a task since it requires situated robotics, perception, real-time decision making, planning, plan recognition, learning and multi-robot coordination and control. The technology to perform real-time vision and build autonomous robots is available; the Dynamite testbed has been built to perform experiments with multiple robots. A soccer tournament has been carried out using the testbed to evaluate aspects of the proposed reactive deliberation robot architecture. The results raise new issues and problems for research on robotic agents operating in dynamic environments.

1 Introduction

One of the dreams of Artificial Intelligence is the construction of integrated cognitive robots. Such robots must be able to integrate perception, reasoning, and action. These robots should be able to operate in the real world, which is dynamic and uncertain, not just in highly restricted environments such as factories. If building real robots is still part of the dream of AI, then we need to develop tools and theories to accomplish this goal.

Unfortunately, current research in AI is highly divergent with little or no overlap between specialized subfields such as computational vision, knowledge representation, robotics, and learning. Each group has its own conferences and journals, and when they do all meet at a single conference, they diverge in parallel sessions. The version of divide-and-conquer that we have been playing, namely, functional decomposition, is not now the best strategy.

For significant progress to be made on the AI dream, researchers must work on common tasks. But which tasks? It is clear that any science must close its eyes to most of the allures and mysteries of nature and choose a highly circumscribed fragment of reality to examine. Indeed, the key experimental task domain may well be an abstraction of the world; but we must take care to preserve the key problems and not abstract them away. For example, Galileo chose, as his blocks world, bodies sliding on a friction-free inclined plane in a vacuum; Newton considered point masses of infinite density. The danger in selecting a problem domain is that researchers must steer a course between the Scylla of enunciating a vacuous general theory of an artificial world and the Charybdis of implementing a collection of quick and dirty hacks that work, after a fashion, on an overly complex domain not properly abstracted, delimited or understood.

There have been a number of task domains that have served to focus AI research since its inception. Chess, the blocks world, video games, Tweety, the Yale Shooting Problem and many others have all served to motivate and focus the efforts of communities of researchers. We should realize that the choice of task domain is a theory-laden decision; that decision should be taken explicitly by the research community.

The Good Old Fashioned AI and Robotics (GOFAIR) [Mackworth, 1993] research paradigm has shaped the area of robotics since the time of the robot Shakey [Nilsson, 1984]. Some of the fundamental assumptions made of the world were that there is only one agent, that the environment is static unless the agent changes it, that actions are discrete and are carried out sequentially and that the world the robot inhabits can be accurately and exhaustively modeled by the robot. These assumptions proved to be overly restrictive and ultimately sterile. In the usual dynamic of the scientific dialectic, a new movement has emerged as the antithesis to GOFAIR: Situated or Nouvelle AI, which we will call the Situated Agent approach.

The Situated Agent paradigm is loosely characterized by the guiding principles set forth by Brooks: situatedness, embodiment, intelligence and emergence [Brooks, 1991].

The key idea of situatedness and embodiment is that researchers in AI should consider embodied agents that are connected to a larger world that provides the context for their activity. The essence of intelligence and emergence is that the intelligence of an agent can be judged by the quality of its interaction with its environment. The motivation for these principles is to direct research toward more realistic tasks and architectures and away from the Scylla of ungrounded theories.

A paradigmatic domain is needed to test and develop the competing GOFAIR and Situated Agent approaches. It must be suitable for testing extant theories and be sufficiently rich to bring the many threads in AI back together.

2 Why Soccer as a Task Domain?

We propose that playing soccer be a paradigmatic task domain since it breaks with nearly all of the restrictive assumptions on which GOFAIR is based and meets the standards proposed in the Situated Agent approach. The soccer domain can be characterized by the following:

- ☐ Neutral, friendly, and hostile agents
- ☐ Interagent cooperation
- ☐ Real-time interaction
- ☐ Dynamic environment
- ☐ Real and unpredictable world
- ☐ Objective performance criteria
- ☐ Repeatable experiments

The GOFAIR assumptions do not hold in the soccer world. The one agent assumption is violated: there are cooperating agents on the robot's team, competing agents on the other team, and neutral agents such as the referee and the weather. The world is not completely predictable: it is not possible to predict precisely where the ball will go when it is kicked, even if all the relevant factors are known. The simplifying assumption of discrete sequential actions is violated: continuous events such as a player running to a position and the ball moving through the air occur concurrently.

Soccer meets the standards of the Situated Agent approach. In soccer, robot agents are embodied and are situated in an unfolding game. Although it is still true that the intelligence of an agent can be judged from the dynamics of interaction with the environment, soccer also provides *objective performance criteria*.

The ability to score and prevent goals and the overall score of the game are objective measures of success. These measures allow explicit comparisons of alternative controller designs. The effects of chance can be factored out by carrying out repeated experiments. With objective criteria and repeatability, short-term and long-term learning strategies, as well as experiments in automatic evolution of controllers, become feasible. The availability of objective criteria is a critical feature of soccer that distinguishes it, along with the aspect of a real and unpredictable environment, from many of the other task domains proposed for driving the new research paradigm.

Soccer as a task domain is sufficiently rich to support research integrated from many branches of AI. In addition to the obvious potential of the soccer domain for research in perception and robotics, there are many other areas of AI that are applicable: reasoning under uncertainty, on-line reasoning, resource-bounded reasoning, planning, decision theory, qualitative physics, plan recognition, learning, and multi-agent theory.

Soccer is not the real world, but a suitably circumscribed fragment of it. Soccer is an appropriate abstraction of the world to challenge research in AI to focus on achievable tasks, and to drive the development of relevant theories.

3 Dynamite: A Testbed for Multiple Mobile Robots

The Dynamite testbed provides a practical platform for testing theories in the soccer domain using multiple mobile robots. The testbed consists of a fleet of radio controlled vehicles that perceive the world through a shared perceptual system [Barman *et al.*, 1993]. In an integrated environment with dataflow and MIMD computers, vision programs can monitor the position and orientation of each robot while planning and control programs can generate and send out motor commands. This approach allows umbilical-free behaviour and very rapid, lightweight fully autonomous robots.

The mobile robot bases are commercially available radio controlled vehicles. We have two controllable 1/24 scale racing-cars, each 22 cm long, 8 cm wide, and 4 cm high excluding the antenna. The testbed (244 cm by 122 cm in size) with two cars and a ball is shown in Figure 1. The cars have each been fitted with two circular colour markers to allow the vision system to identify their position and orientation. The ball is the small object between the cars.

The hardware used in this system is shown in Figure 2. There is a single colour camera mounted in a fixed position above the soccer field. The video output of the camera is transmitted to special-purpose video processing DataCube hardware in Figure 2. The DataCube is a dataflow computer which has been programmed to classify image pixels into different colour classes at video rate (60 Hz). This information is transmitted to a network of transputers which form a MIMD computer. Additional vision processing is performed on the transputers to find the position, in screen coordinates, of the centroid of each coloured blob and to transform these positions from screen to world coordinates. The vision subsystem is called the Vision Engine [Little *et al.*, 1991]. The Vision Engine produces the absolute position of all the objects on the soccer field; the orientation of each car is also reported. This is done at 60 Hz with an accuracy in position of approximately 1 mm.

The reasoning and control components of a vehicle can be implemented on any number of transputers out of the available pool. Currently, each vehicle is controlled

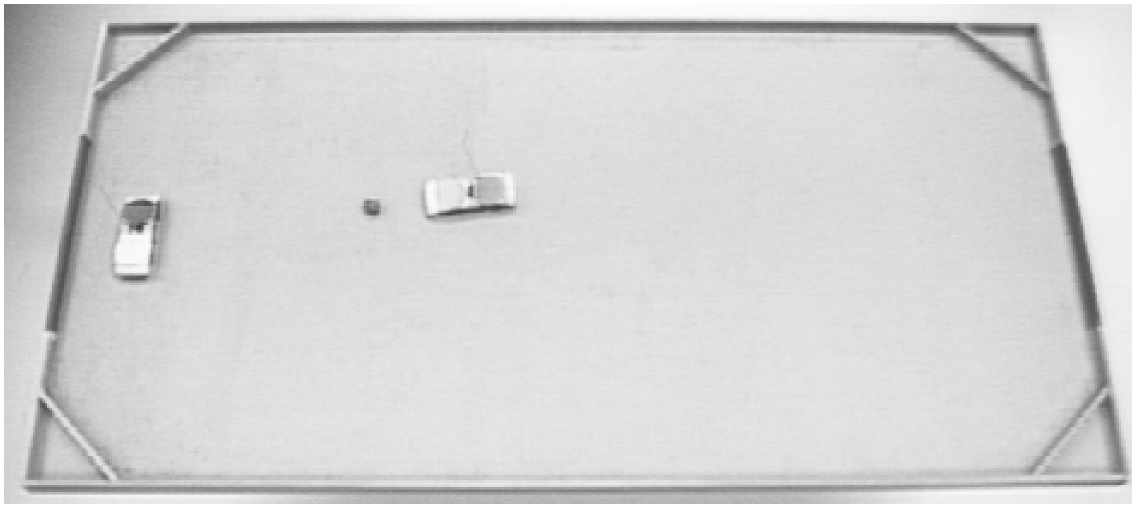

Figure 1 Robot Players on the Soccer Field

by a distributed user program running on two transputer nodes. An arbitrary number of nodes, labeled 1 to n in Figure 2, can be used in parallel to control independent vehicles. The movement of all vehicles is controlled through radio transmitters attached to a single shared transputer node. Commands are transmitted to the vehicles at a rate of 60 Hz.

A physics-based real-time graphics simulator for the Dynamite world is also available for testing and developing reasoning and control programs.

A feature of the Dynamite testbed is that it is based on the “remote brain” approach to robotics. The testbed avoids the technical complexity of configuring and updating on-board hardware and makes fundamental problems in robotics and artificial intelligence more accessible. We have elected not to get on-board the on-board computation bandwagon, since the remote (but untethered) brain approach allows us to focus on scientific research without devoting resources to engineering compact electronics.

4 A Robot Architecture for Dynamic Domains

Most extant theories of robot architectures do not directly address the problems posed by dynamic environments. In a changing world, an agent must be able to generate intelligent behaviour in real-time. The soccer domain is a good testing ground for theories that address these issues since it is a highly dynamic environment. In this section, an architecture targeted towards dynamic environments, reactive deliberation, is described.

Much of the previous work on architectures for dynamic environments has been addressed by two distinct schools. Architectures in the situated behaviour school [Brooks, 1986; Agre and Chapman, 1987; Kaelbling and Rosenschein, 1990] typically allow frequent changes in the actions of the robot, yet restrict the allowable computational models. The planning school [Nilsson, 1984; Firby, 1992; Gat, 1992] allows unrestricted computational models, yet the commitment to arbitrary length plans hinders the ability

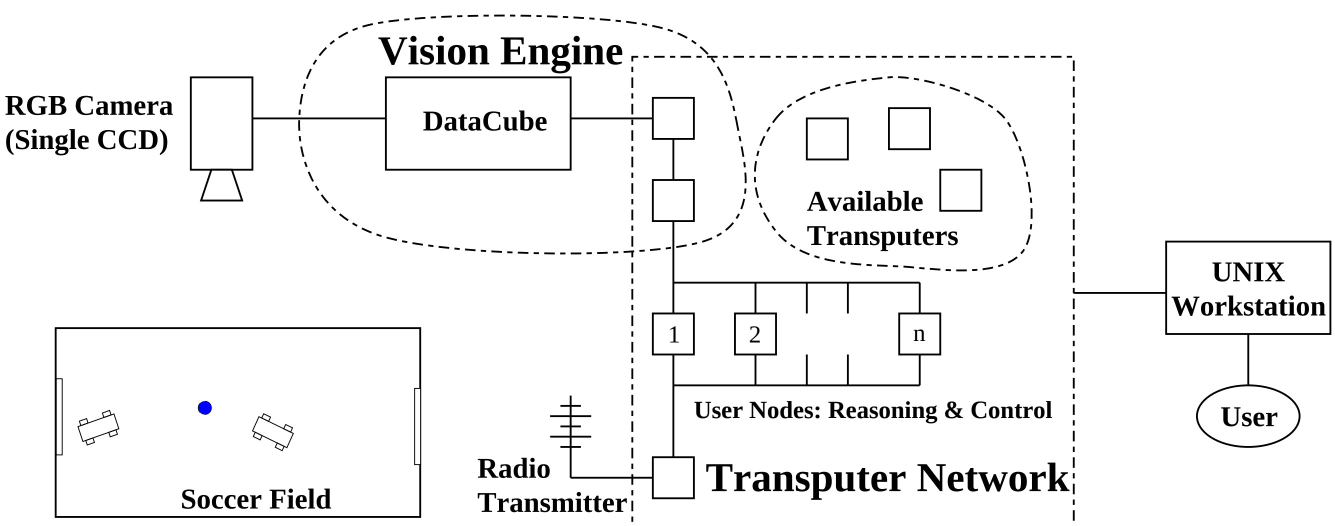

Figure 2 The Dynamite Hardware Setup

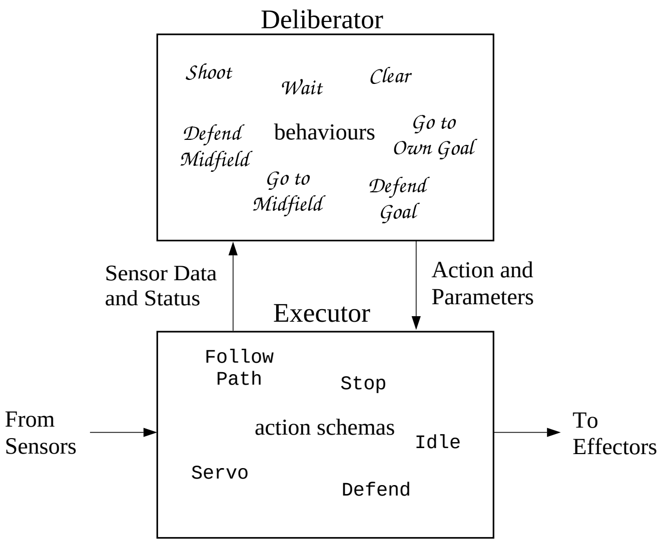

Figure 3 The Reactive Deliberation Controller

of the agent to change its goals and actions in response to unanticipated changes in the environment.

The problem of deciding what to do next has also been addressed in decision theory [Kanazawa and Dean, 1989], Maes' dynamics of action selection [Maes, 1990], and Minsky's mental proto-specialists [Minsky, 1986]. Decision theoretic tools are limited in their ability to handle continuous variables and perform sophisticated spatial reasoning. The dynamics of action selection performed poorly in simulations [Tyrrell, 1993] in part due to a reliance of the model on predicate inputs. Minsky's arguments against mental proto-specialists (that bid against one another for control of the agent) neglect to include the external state of the world as a valid basis for decisions.

Reactive deliberation is a robot architecture that combines responsiveness to the environment with intelligent decision making [Sahota, 1993; Sahota, 1994]. Even deliberation must be to some extent be reactive to respond to changes in the environment. Although the name is apparently an oxymoron, it is consistent with Artificial Intelligence nomenclature (cf. Reactive Planning).

Under reactive deliberation, the robot controller is partitioned into a deliberator and an executor; the distinction is primarily based on the different time scales of interaction. Informally, the deliberator decides what to do and how to do it, while the executor interacts with the environment in real-time. These components run asynchronously to allow the executor to interact continuously with the world and the deliberator to perform time consuming computations. A structural model illustrating the partition with examples can be seen in Figure 3. The deliberator is responsible for generating a single action, whereas other planning-based architectures generate a complete plan (i.e. sequences of actions). This distinction helps focus the deliberative activities on the immediate situation.

The executor is composed of a collection of action schemas. An *action schema* is a robot program that interacts with the environment in real-time to accomplish specific actions. Only one action schema is enabled at a time and it interacts with the environment through a tight feedback loop. The active schema receives run-time parameters from the deliberator that fully define its activity.

The focus of the deliberator is on an effective mechanism for selecting actions or goals in a timely manner. A central feature of reactive deliberation is that the deliberator is composed of concurrently active modules called *behaviours* that represent the goals of the robot. The notion of a behaviour is used in the sense of Minsky's mental proto-specialists [Minsky, 1986] with some important distinctions. In reactive deliberation, each behaviour computes an action and generates a bid reflecting how suitable it is in the current situation. The most appropriate behaviour, and hence action, is determined in a distributed manner through inter-behaviour bidding. Some examples of behaviours are: shoot ball, defend goal, go to midfield, clean floor, and deliver mail.

A *behaviour* is a robot program that computes an action that may, if executed, bring about a specific goal. Behaviours propose actions whereas action schemas perform actions. Each behaviour must perform the following: 1) select an action schema, 2) compute run-time parameters for the schema (plan the action), and 3) generate a bid describing how appropriate the action is.

Behaviours in reactive deliberation have a number of features. Different computational models can be used within behaviours to provide flexibility in the design of robot controllers. Inter-behaviour bidding is an effective mechanism for goal arbitration [Tyrrell, 1993] and can also be accomplished in a distributed computing environment. Another important property is that behaviours can be used as a mechanism for distributing computational resources.

Reactive deliberation is not a panacea for robotic architectural woes. A further disclaimer is that it is an incomplete robot architecture since it focuses on the issues related to dynamic domains and ignores a number of issues such as perceptual processing and the development of world models. The proposal is orthogonal to those issues. However, it makes explicit the need to evaluate the actions and goals of the robot at a rate commensurate with changes in the environment.

5 Some Experimental Results

Several controllers based on reactive deliberation have been implemented to allow robots to compete in complete one-on-one games of soccer [Sahota, 1993]. Current functionality includes various simple offensive and defensive strategies, motion planning, ball shooting and playing goal. The robots can drive under accurate control at speeds up to 1 m/s, while simultaneously considering alternate actions. We have produced a 10 minute video that documents these features.

As documented in [Sahota, 1993], a series of experiments, soccer games, called the Laboratory for Computational Intelligence (LCI) Cup were performed using the Dynamite testbed. The most elaborated reactive deliberation controller competed with subsets of itself to provide, through the scores of the games, an objective utility measure for some of the architectural features of reactive deliberation and the behaviour themselves. Through the results of the LCI Cup the importance of modifying goals in response to changes in the environment has been shown. Further, the results demonstrate that the architectural elements in reactive deliberation are sufficient for real-time intelligent control in dynamic environments.

The reactive deliberation architecture provides a first step towards an integrated intelligent agent for dynamic environments. The current version of the controller can only play adequately in one-on-one soccer. Even in this restricted task domain, there are many unresolved problems. There are several important issues that need to be further addressed in building robot agents, such as:

- ☐ Real-time decision making — Reasoning about the world and selecting appropriate actions in real-time.
- ☐ Planning — Efficiently computing motion plans, predicting future world states, and reasoning about actions in an uncertain world.
- ☐ Plan recognition — Identifying the goals, actions and plans of other agents.
- ☐ Modeling — Acquiring implicit or explicit models of the robot and the environment.
- ☐ Learning — Changing behaviour at many levels through tuning models and refining actions using objective performance criteria.
- ☐ Multi-agent theory — Determining how agents can cooperate to accomplish group tasks.
- ☐ Robot architectures — Integrating all of the above components in new organizational forms.

We have shown that the Dynamite testbed is a useful abstraction of the soccer domain that can be used to test and develop many theories. However, it has a significant limitation. Off-board perception through an overhead camera leads to the pervasive use of world coordinates. The convenience of using a world model bypasses many important issues in robot vision and sensory robotics. For soccer experiments to address these issues in situated perception, a new testbed with on-board sensing will have to be developed.

6 Conclusions

Soccer has been proposed as a task for the development and unification of divergent theories in Artificial Intelligence. Soccer captures a number of essential properties of the real world including dynamics, real-time requirements, and cognitive functions. To perform experiments with soccer, the Dynamite testbed has been constructed with support for multiple mobile robots. A theory of robot architecture, reactive deliberation, has been applied to the soccer domain using the Dynamite testbed with demonstrated success. The results suggest that a wide range of theories from decision theory to robot control need further development to be successful in domains like this. This paper can be viewed as a challenge to researchers to apply their theories to the soccer domain to determine whose team of agents will win the Robot Soccer World Cup.

The question posed in the title, “Can Situated Robots Play Soccer?” has at least four possible answers: “Yes”, “No”, “Don’t Know”, and “Don’t Care”. We claim to have provided evidence for “Yes”. But, one could argue for “No” based on the limitations of our experiments or our theories. “Don’t Know” now seems inappropriate. “Don’t Care” is a response that ignores the current theoretical and experimental needs of the field. Not only *can* situated robots play soccer but they also *should*!

Acknowledgments

We are grateful to Rod Barman, Keiji Kanazawa, Stewart Kingdon, Jim Little, Dinesh Pai, Heath Wilkinson and Ying Zhang for help with this. This work is supported, in part, by the Canadian Institute for Advanced Research, the Natural Sciences and Engineering Research Council of Canada and the Institute for Robotics and Intelligent Systems Network of Centres of Excellence.

References

[Agre and Chapman, 1987] Philip Agre and David Chapman. Pengi: An implementation of a theory of activity. In *AAAI-87*, pages 268–272, 1987.

[Barman *et al.*, 1993] R. Barman, S. Kingdon, J. Little, A. K. Mackworth, D.K. Pai, M. Sahota, H. Wilkinson, and Y. Zhang. Dynamo: real-time experiments with multiple mobile robots. In *Proceedings of Intelligent Vehicles Symposium*, pages 261–266, 1993.

[Brooks, 1986] Rodney A. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, RA-2:14–23, 1986.

[Brooks, 1991] Rodney A. Brooks. Intelligence without reason. In *IJCAI-91*, pages 569–595, 1991.

[Firby, 1992] R. James Firby. Building symbolic primitives with continuous control routines. In *First International Conference on Artificial Intelligence Planning Systems*, pages 62–69, 1992.

[Gat, 1992] Erann Gat. Integrating planning and reacting in a heterogeneous asynchronous architecture for controlling real-world mobile robots. In *AAAI-92*, pages 809–815, 1992.

[Kaelbling and Rosenschein, 1990] Leslie Pack Kaelbling and Stanley J Rosenschein. Action and planning in embedded agents. In Pattie Maes, editor, *Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back*, pages 35–48. M.I.T. Press, 1990.

[Kanazawa and Dean, 1989] Keiji Kanazawa and Thomas Dean. A model for projection and action. In *IJCAI-89*, pages 49–54, 1989.

[Little *et al.*, 1991] J. Little, R. Barman, S. Kingdon, and J. Lu. Computational architectures for responsive vision: the vision engine. In *Proceedings of Computer Architectures for Machine Perception*, pages 233–240, 1991. Paris.

[Mackworth, 1993] Alan Mackworth. On seeing robots. In A. Basu and X. Li, editors, *Computer Vision: Systems, Theory, and Applications*, pages 1–13. World Scientific Press, 1993.

[Maes, 1990] Pattie Maes. Situated agents can have goals. In Pattie Maes, editor, *Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back*, pages 49–70. M.I.T. Press, 1990.

[Minsky, 1986] Marvin Minsky. *The Society of Mind*. Simon & Schuster Inc., 1986.

[Nilsson, 1984] Nils Nilsson. Shakey the robot. Technical Report 323, SRI International, 1984. Collection of Earlier Technical Reports.

[Sahota, 1993] Michael K. Sahota. Real-time intelligent behaviour in dynamic environments: Soccer-playing robots. Master's thesis, University of British Columbia, 1993.

[Sahota, 1994] Michael K. Sahota. Reactive deliberation: An architecture for real-time intelligent control in dynamic environments. In *Proceedings of AAAI-94*, 1994. Forthcoming.

[Tyrrell, 1993] Toby Tyrrell. *Computational Mechanisms for Action Selection*. PhD thesis, Edinburgh University, 1993.