SOCIAL ROBOTICS: Reality and Virtuality in Agent-Based Robotics

B.R. Duffy, R.W. Collier, G. M. P. O'Hare, C.F.B. Rooney, R.P.S. O'Donoghue
PRISM Laboratory, Dept. of Computer Science, University College Dublin (UCD), Belfield, Dublin 4, Ireland
{Brian.Duffy, Rem.Collier, Gregory.OHare, Colm.Rooney, Ruadhan.ODonoghue}@ucd.ie

Abstract

Abstract. This paper advocates the application of multi-agent techniques in the realisation of social robotic behaviour. We present the *Social Robot Architecture*, which integrates the key elements of agenthood and robotics in a coherent and systematic manner. This architecture seamlessly integrates real world robots, multi-agent development tools, and VRML visualisation tools into a coherent whole.

Keywords: *intelligent agents, robotics, agent communication languages, virtual reality*

1 INTRODUCTION

The primary concern of this paper is that of *Social Robotics* [Duffy *et al*, 1999]. We investigate the processes whereby robot communities engage in opportunistic collaborative behaviours in the solution of shared tasks.

This paper advocates the application of multi-agent techniques in the realisation of social robotic behaviour. To this end we commission *Agent Factory*, a software environment that facilitates the rapid prototyping of Multi-Agent Systems (MAS). Agent Factory offers a conduit through which robot control can be governed by a deliberative agent architecture, specifically that of a Belief-Desire-Intention (BDI) architecture. In addition, Agent Factory supports not only the creation of the social robotic community but the subsequent experimentation with and visualisation of their behaviour.

We present the *Social Robot Architecture*, which integrates the key elements of agenthood and robotics in a coherent and systematic manner. This architecture seamlessly integrates real world robots, multi-agent development tools, and VRML visualisation tools into a coherent whole. Using these elements, we deliver a development environment, which facilitates rapid prototyping of social robot communities.

Section 2 and section 3 offer respectively an introduction to BDI architectures and an overview of Agent Factory. Agent-based robotics is presented in section 4, while section 5 introduces the *Social Robot Architecture*. Section 6 describes experimental results and discussion and conclusions are presented in section 7.

2 BELIEF-DESIRE-INTENTION (BDI)

Much research work has been commissioned on Multi-Agent Systems (MAS) and Distributed Artificial Intelligence (DAI) [Bond and Gasser, 1988; Durfee *et al*, 1989; O'Hare and Jennings, 1996]. Specifically, competing agent

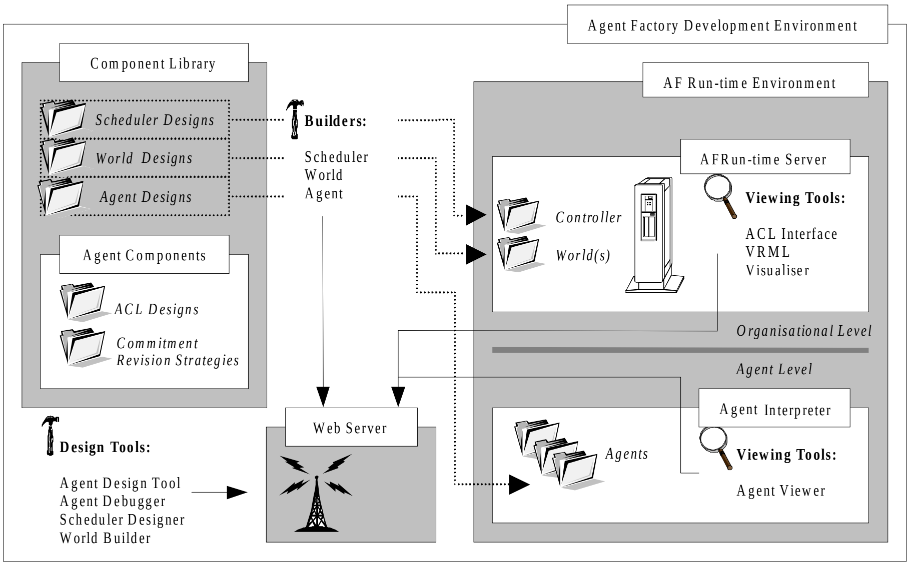

Figure 1. Agent Factory

architectures have been proposed in the literature. Two major architectural schools have emerged, namely *the reactive system school* and the *deliberative system school*. The former has predominated in the arena of autonomous mobile robot control. In this paper we synthesis reactive and deliberative reasoning.

In the delivery of computationally tractable models of deliberative reasoning, one approach that has gained wide acceptance is to represent the properties of an agent using mental attitudes such as *belief*, *desire*, and *intention*. In this terminology, an agent can be identified as having; a set of beliefs about its environment and about itself; a set of desires which are computational states which it wants to maintain, and a set of intentions which are computational states which the agent is trying to achieve. Multi-agent architectures that are based on these concepts are referred to as *BDI-architectures* (Belief-Desire-Intention) [Rao and Georgeff, 1991; Jennings, 1993; O'Hare and Jennings, 1996], and have recently been the subject of much theoretical research. Proponents of the BDI approach argue that the understanding of the dynamics of these mental attitudes and their intimate interdependencies is crucial in achieving *rational* agent behaviour. Within this research we adopt a similar standpoint. We seek to develop *social robotic systems* by embracing the BDI philosophy. To this end, we utilise Agent Factory, a multi-agent prototyping tool.

3 AGENT FACTORY

3.1 What is Agent Factory?

In essence, Agent Factory is a tool which facilitates the rapid prototyping of Intelligent Agents. The Agent Factory System has been discussed more completely elsewhere in the literature [O'Hare and Abbas, 1995; Collier, 1996; O'Hare and Jennings, 1996; O'Hare et al, 1998] [O'Hare *et al*, 1999].

The system offers an integrated toolset that supports the developer in the instantiation of generic *agent structures* that are subsequently utilised by a pre-packaged agent interpreter that delivers the BDI machinery. Other system tools support interface customisation and agent community visualisation. In creating an agent community three system components must be interwoven, those of *agents*, a *world* and a *scheduler*. The next section describes the high level architecture.

3.2 Schematic Functional Architecture

In order to provide the necessary functionality for the delivery of agent communities, the Agent Factory System has been divided into two key areas: the *Agent Factory Development Environment*, and the *Agent Factory Run-Time Environment* (see figure 1). The run-time environment provides the support necessary for the release of a completed Multi-Agent System. This environment is further sub-divided into a *Run-time Server*, and an *Agent Interpreter*. The run-time server offers two main services: access to non-agent components of the system (a controller & some worlds), and a set of tools for viewing and interacting with the agent community. The Agent Interpreter provides the functionality necessary for the execution and visualisation of agents.

The development environment represents a subset of the Agent Factory Run-time Environment. It adds a *Component Library* and a selection of tools to facilitate the rapid prototyping of agent communities. The Component Library is built from Component Design Hierarchies (CDH) that extend the standard Object Hierarchies in the OOP Paradigm. A CDH exists for each of the main system.

The *agent* is the main computational unit of Agent Factory, it combines a series of attributes that represent and support the Mental State model of the agent, a set of methods (the actuators), a set of perceptors, an Agent Communication Language (ACL), and a Commitment Revision Strategy. This design is then executed using a generic Agent Interpreter which delivers the BDI machinery.

The *scheduler* controls execution of the community, using an algorithm that exploits parallelism where possible.

Finally, the *world interface* acts as a medium between the problem domain, the community it is being developed for, and the other components of the Agent Factory System.

Access to these environments is provided both locally through Graphical User Interfaces (GUIs) and remotely through the World Wide Web (WWW) via a purpose built *Web Server*.

4 AGENT BASED ROBOTICS

Initial research focused on the behaviour [Brooks, 1986; Steels, 1994] and navigation problems associated with single robots, more recently the area of systems of multiple robots has demanded considerable attention. There are numerous advantages in the use of multiple robots, these include *inter alia* distributed capabilities; parallelism; task and load distribution; increased functionality with minimal complexity. However, achieving coherent behaviour presents considerable challenges, none more acute than overcoming problems of co-ordination and interference. It seems clear that multi-agent techniques are amenable to transference to systems of multiple autonomous robots, in particular to addressing the problems of co-ordination and interference.

Initial work on agent based robotics emerged from cellular robotics where the robots had limited functionality and relied on swarm like intelligence to achieve their desired task, typically exhibiting *emergent* capabilities [Cai *et al*, 1995; Fukuda *et al*, 1989]. More information can be found elsewhere: [Dudek *et al*, 1996; Arkin and Balch, 1998; RoboCup].

5. TOWARDS SOCIAL ROBOTICS

The motivation behind this research is to demonstrate, in a tangible form, the correspondence between systems of multiple agents and systems of multiple robots. Research in the multi-agent community has concentrated primarily on software agents with little or no physical manifestation of their behaviour. Research in the robot community has focused on non-collaborative robotic behaviour with relatively simplistic cognitive models. Neither community has suitably addressed the opportunity of applying BDI-based agent technologies in the control of social robots.

5.1 Terminology

It is our conjecture that a distinction exists between *societal robotics and social robotics*. The former represents the integration of robotic entities into the human environment or society, while the latter deals specifically with the social empowerment of robots permitting opportunistic goal solution with fellow robots.

5.2 Social Robot Architecture

The computational machinery needed to facilitate team building and collaborative behaviour is non-trivial. We describe the *Social Robot Architecture*, which goes some way toward achieving this through the judicious synthesis of the reactive model with that of the deliberative model. The layered architecture (figure 2) has four fundamental elements: the physical level, the reactive level, the deliberative level provided via Agent Factory, and the social level, which we consider in more detail in the subsequent sections. The control architecture offers basic reactive behaviours for reflex robot responses to unexpected or dangerous events. These constitute a set of primary *survival behaviours* for the robot. Goal oriented behaviours are delivered through the intentional agent structures of Agent Factory.

We now consider the architecture's functional components.

5.2.1 The Physical Level

Robots in our terminology may take the form of either that of a physical entity, specifically the Nomad Scout II from Nomadic Technologies or a simulated entity, provided by Nserver, a robotic simulator for the Nomad Scout II robots. Robots are situated in the world and exhibit behaviour based upon degrees of perception, action and interaction. Three such physical robots Aengus, Bodan and Bui[1] are used for experimentation.

5.2.2 The Reactive Level

A series of fundamental reflex behaviours are implemented at this level. They provide the basic *survival kit* of the robot necessary for dynamic and unpredictable environments. These *Survival Behaviours* include such behaviours as `avoid_obstacle`, `stop`, and `retreat`.

The sensory information received from the 16 ultrasound sensors, the bumper ring, and the odometry is processed at this level, resulting in clear *agent events* being generated and communicated to the agent's deliberative level (i.e. `door_found(X,Y,Xi,Yi)`).

5.2.3 The Deliberative Level: Agent Factory

The deliberative layer provides the deliberative machinery. This is achieved through the BDI architecture described in section 3. An agent's beliefs are generated based on its belief set, and are updated with the receipt of new agent states or events from the robots sensors and communication with other robots

An agent starts by gathering *perceptions* (raw sensor data). Low level robot processing filters these and thus achieves attention focusing, generating a set of agent events. If emergency action is required, a reflex behaviour is activated to overcome the immediate problem and this action is communicated as an agent event to the deliberative level. The perception process at this level deals with converting these agent events into beliefs and adding them to the belief set providing the agent with an up to date model of its current

[1] Aengus, Bodan and Bui are ancient characters in Irish legend and their graves can be found at Bru na Boinne, Ireland.

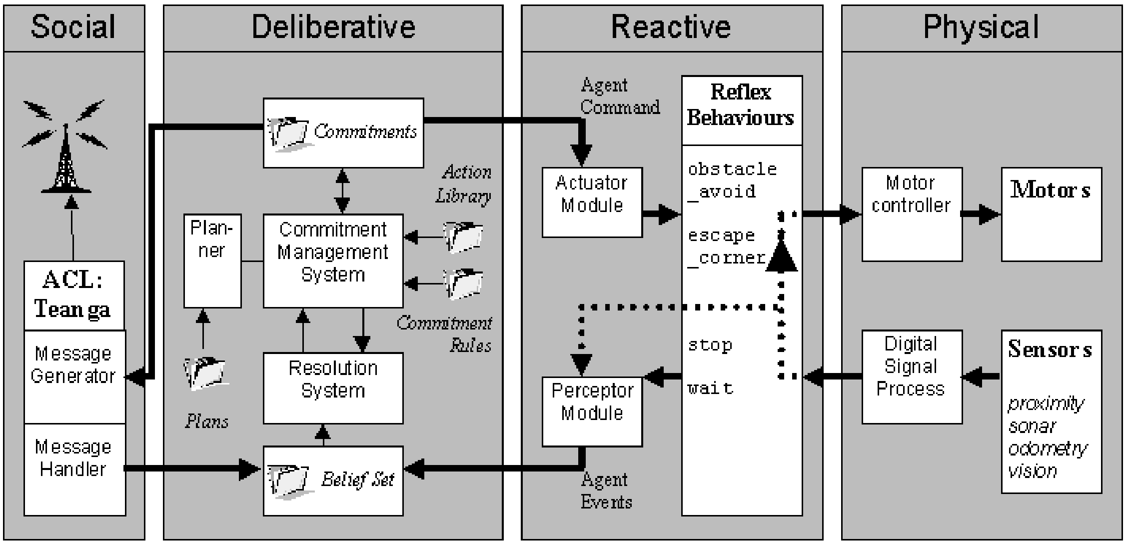

Figure 2. The Robot Agent

perceived situation. This situation is then analysed and the agents' commitments are updated. Pre-existing commitments are analysed and those pertaining to the current time frame honoured resulting in either a communicative act being sent to the social level or a physical act being passed to the actuators via the reactive level.

By way of example, a team of robots may have a goal of mapping their environment. Decomposing this problem, the individual robots might have to map constituent rooms, which in turn are made up of walls and a door. By analysing the data gathered from the sonar's these components might be identified. If a wall is found the robot then generates an agent event of the form, `wall_from(X,Y,Xi,Yi)`. Upon receipt the robot agent will revise it's belief set and accordingly add the belief `Bel(wall_from(X,Y,Xi,Yi))`.

5.2.4 The Social Level

Our agents interact via an Agent Communication Language (ACL), Teanga, which is described in more detail in section 5.3. The social level comprises of a message generator and a message handler. When a message is received the message handler decodes, checks and passes it on to the deliberative layer where it is dealt with accordingly, e.g. the belief set is updated or a commitment adopted. When a commitment to send a message is acted upon, the commitment management system utilises the services of the message generator to correctly generate and send the appropriate message. The hardware required to achieve this inter-robot communication is the Mercury-EN wireless unit from Nomadic Communications Inc, and the RangeLAN from Proxim Inc.

5.3 Teanga

Inter-agent communication demands a language with enough expressive power to represent and express the concepts of beliefs, intentions, information and service requests and replies to such. Agent communication languages (ACLs) generally have a syntax powerful enough to generate a wide range of communicative actions. Our agents interact via *Teanga*[2] [Rooney *et al*, 1999], an Agent Communication Language based upon Speech Act Theory [Austin 1969].

[2] Teanga is Irish for "language" and interestingly is an anagram for a agent.

There are a number of reasons why we chose to develop our own ACL rather than work with an established language such as KQML or FIPA's ACL [FIPA 97]. One of the main reasons, however, is that our language should be compositional. We wish to be able to support nested speech acts (and speech acts contained inside composite actions, e.g. plans). To implement a nested communicative act in KQML requires that the content be a KQML message. However, in KQML the content language is independent and there is therefore no content checking. *"A disadvantage of content independence is that it prevents the content from being checked for compatibility with the speech act type"* [Cohen & Levesque, 1995]. Our language shall have the constraint that the content language considers speech acts as actions like any other. So, an agent may, for example, request, or commit to, a speech act which will in turn be represented in our syntax and which may in turn contain nested speech acts.

Examples of the communication structure follow:

SPEECH ACT(Aengus Bodan) (**INFORM BELIEF** Bel(wallfrom(x1y1;x2y2)) now)
SPEECH ACT(Bodan Bui) (**REQUEST ACHIEVE** followwall(x1y1;x2y2) now)
SPEECH ACT(Bui Bodan) (**COMMIT ACHIEVE** followwall(x1y1;x2y2) now)

5.4 The Worlds

One of the key tenants of our research has been the provision of multiple views of the operation of the same multiple robot system. We introduce the concept of the *virtual robotic workbench*.

The system provides two distinct views. The first provides a physical view of the Nomad Scout II's navigating the *robot world*. We supplement this view with a virtual view to support behaviour generation and testing. This provides a virtual reality via the Agent Factory Visualiser which renders a 3-D VRML world (see figure 3) Internet observable through any standard browser.

Herein we harness the advantages of using virtual environments, by directly relating simulation and reality in a seamless manner. Such alternate views permit multiple views, information hiding and abstraction, system interaction, and behaviour scrutiny via snapshots and recordings.

The *virtual robotic workbench* supports the specification and invocation of robot experiments via a Java built Internet interface. This allows the selection of a virtual world (which may represent an existing real world) and the embodiment of either one or more robots within this world. This is achieved via the Agent Factory Web Interface and therefore allows the use of virtual reality as a real-time web visualisation tool for remote behavioural analysis of real world robot experiments.

6 EXPERIMENTS

A virtual world of the PRISM Research Laboratory has been built to enable comprehensive and controlled testing of basic and complex behavioural models. This world includes four offices connected principally by a corridor, as seen in figure 3, and replicates to scale the reality of the robots' experimental environment.

This therefore provides two environments: the real world and its Virtual Reality (VR) representation. A series of systematic tests have been conducted and examples of these are presented in the following subsections.

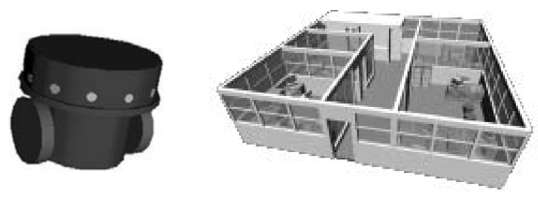

Figure 3: Virtual reality model of a robot & the PRISM Research Laboratory

6.1 Linking Reality and Virtuality

The initial task was to establish a correspondence between the physical robots and their virtual counterparts. An agent was specified within Agent Factory and associated with a physical robot via the Agent Factory Robot Interface. A virtual robot representation of the robot was then situated in its virtual world. The virtual robot should act as a 3-dimensional view on the behaviour of the real robot, mirroring its movement. To achieve this the Agent Factory Visualiser monitors the telemetry information from each of the agents' *agent event* queue and updates the virtual robots' positions accordingly.

We demonstrated through a variety of robot tasks how both views corresponded to the anticipated robot behaviour. Figure 4 shows the dual perspective of this experiment showing the link between the virtual and real world.

Figure 4: VR and real view of same robot performing the above behaviour

In addition the WWW observable VR representation of reality could be used to remotely observe the physical robots behaviour.

6.2 Co-ordination between Robots

Having established a correspondence between the virtual and real robots, simple experiments were then undertaken to investigate inter-robot communication.

A robot (Aengus) was situated in a complex office environment with a second (Bodan) situated in a featureless environment without obstacles of any form. Aengus performed a *wander* behaviour whereby the robot moved around its environment avoiding obstacles for a period of time with Bodan replicating the actions of Aengus.

As can be seen in figure 5, the robot Bodan was able to replicate the motions of Aengus, demonstrating inter-robot communication of abstract agent events within Agent Factory.

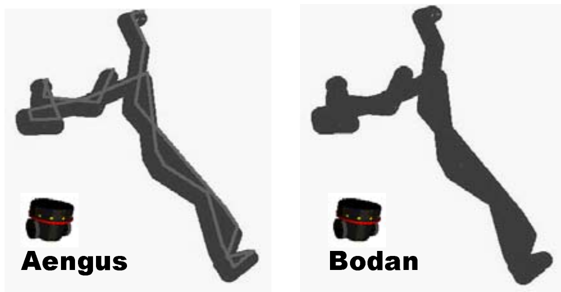

Figure 5: Bodan replicating the motion trajectories of Aengus

6.3 A Social Task

A suitable collaborative task was identified which demanded Aengus and Bui to waltz. The specific waltz was that of the "Box Waltz" which consists of a simple square shape repetitive motion.

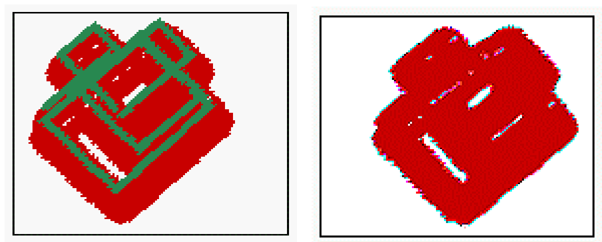

Figure 6: The robot duo's performed behaviour in a square room

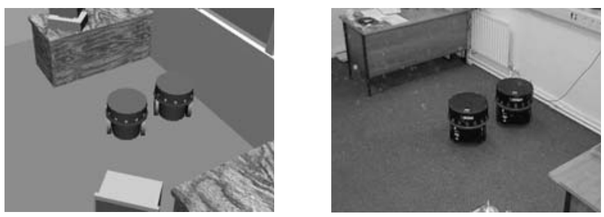

Figure 7: The Aengus-Bodan duo performing a waltz, shown in VR and reality

The objective was to perform this with two partner robots (Aengus and Bui) *waltzing* side by side in a space-restricted environment. When one robot encounters an obstacle or wall, a reflex behaviour is invoked resulting in communication with the second robot and their paths modified by the invocation of a high level reflective behaviour resulting in the next step of the dance. The objective was for a duo of robots to perform the waltz utilising the maximum space available in the room. Figure 6 illustrates the robot motion traces while figure 7 shows the two corresponding views of our dancing duo.

This experiment demonstrated co-ordinated high-level behaviours subjected to reflex actions when necessary, controlled by the Agent Factory system, and viewed remotely through virtual reality via the Internet.

7 DISCUSSION AND CONCLUSIONS

This paper introduces *Social Robots* and the accompanying *Social Robot Architecture*. The architecture provides firstly, a powerful robot control mechanism, which integrates reactive and deliberative features and secondarily, a rich visualisation medium, which offers seamless, contrasting, yet consistent views of social robotic behaviour in both reality and virtual reality.

This work represents one of the first concerted efforts to control multiple robot systems with BDI machinery. To this end we have embedded a pre-existing BDI system, namely Agent Factory, into a *social robot* architecture. Agent Factory delivers complex robot behaviour control and enables the rapid prototyping of such behaviours. The individual robots are equipped with reflex behaviours, enabling high reactivity while performing complex behaviours. The ease of development and application of experiments has demonstrated the viability and potential of this architecture and the underlying *social robot philosophy*.

The *virtual robotic workbench* facilitates the remote specification and subsequent observation of social robotic experiments.

ACKNOWLEDGEMENTS

We gratefully acknowledge the support of Enterprise Ireland through grant No.'s SC/96/621 COMmitment Evolution and DYnamics (COMEDY) and SP/97/08 Initiative for Mobile Programmable Agent-based Computing Technologies (IMPACT).

REFERENCES

[Arkin and Balch, 1998] Ronald C. Arkin, Tucker Balch, Co-operative Multi-agent Robotic Systems, *Artificial Intelligence and Mobile Robots, MIT/AAAI Press, 1998*

[Austin, 1969] John L. Austin. How To Do Things With Words. Oxford University Press, 1969.

[Billard and Dautenhahn, 1997] Aude Billard, Kerstin Dautenhahn. Grounding communication in situated, social robots, *Towards Intelligent Mobile Robots*, Manchester 1997

[Bond and Gasser, 1988] Alan Bond, Les Gasser eds. Readings in Distributed AI, San Mateo, CA., 1988.

[Brooks, 1986] Rodney A. Brooks. A Robust Layered Control System for a Mobile Robot, *IEEE Journal on Robotics and Automation*, 2(1) 1986

[Cai *et al*, 1995] A. Cai, T. Fukuda, F. Arai, T. Ueyama, A. Sakai. Hierarchical Control Architecture for Cellular Robotic System – Simulation and Experiments, *IEEE Int. Conf. On Robotics and Automation, 1995,*

[Cohen & Levesque, 1995]. Philip R. Cohen & Hector J. Levesque, *Communicative Actions for Artificial Agents,* 1995, Proceedings of the International Conference on Multi-Agent Systems, AAAI Press, San Francisco, June, 1995.

[Collier, 1996] Rem Collier. The realisation of Agent Factory: An environment for the rapid prototyping of intelligent agents. *M.Phil., Univ. of Manchester, 1996.*

[Dudek *et al*, 1996] Dudek, G., Jenkin, M., Milios, E., and Wilkes, D., A Taxonomy for multi-agent robotics, *Autonomous Robots, Vol 3, 375-397, 1996*

[Duffy *et al*, 1999] B.R. Duffy, C.F.B. Rooney, R.P.S. O'Donoghue, R.W. Collier, G.M.P. O'Hare, Towards Social Robots *Cognitive Science For The New Millenium, Dublin, 1999*

[Durfee *et al*, 1989] Edmund H. Durfee, Victor R. Lesser, D.D Corkhill, Trends in co-operative distributed problem solving, *IEEE: Knowl. Data Eng. 11(1), 63-8, 1989*

[Dreyfus, 1979] Hubert L. Dreyfus. What Computers Can't Do. *Harper, 1979.*

[FIPA 97] FIPA 1997 Spec. Part 2: Agent Communication Language, "http://drogo.cselt.stet.it/ufv/leonardo/fipa/spec/"

[Fukuda *et al*, 1989] Toshio Fukuda, S. Nakagawa, Y. Kawauchi and M. Buss. Structure Decision for Self Organising Robots Based on Cell Structure-CEBOT, *IEEE: Rob. & Autom., Scottsdale Arizona, 1989,*

[Jennings, 1993] Nick R. Jennings, Specification and implementation of a Belief-Desire joint intention architecture for collaborative problem solving. *Int. Jour. of Intel. and Co-op. Info. Sys. Vol. II no3, 1993.*

[O'Hare and Abbas, 1995] Gregory M. O'Hare and S. Abbas, Commitment Manipulation within Agent Factory, *Proc. of Decent. Intel. & MAS, Cracow, Poland, 1995*

[O'Hare *et al*, 1999] O'Hare, G.M.P., Duffy, B.R., Collier, R.W, Rooney,C.F.B., O'Donoghue, R.P.S., (1999) Agent Factory: Towards Social Robots, *First International Workshop of Central and Eastern Europe on Multi-agent Systems (CEEMAS'99), St. Petersburg, Russia*

[O'Hare and Jennings, 1996] Gregory M. O'Hare and Nick R. Jennings, (Editors.), Foundations of Distributed Artificial Intelligence, *Wiley Interscience Pub., New York, 1996, 296 pages. ISBN 0-471-00675*

[O'Hare *et al*, 1998] Gregory M. O'Hare, Rem Collier, J. Conlon and S. Abbas. Agent Factory: An Environment for Constructing and Visualising Agent Communities, *Proc. AICS98., 1998*

[Lakoff and Johnson, 1980] George Lakoff and Mark Johnson, Metaphors We Live By, *Uni. of Chicago Press, 1980.*

[Rao and Georgeff, 1991] Anand S. Rao and Michael P. Georgeff. Modelling Rational Agents within a BDI Architecture, *Prin. of Knowl. Rep. & Reas., San Mateo, CA., 1991*

[RoboCup] The RoboCup Official web site http://www.robocup.v.kinotrope.co.jp

[Rooney *et al*, 1999] Rooney, C.F.B., O'Donoghue, R.P.S., Duffy, B.R., O'Hare, G.M.P., Collier, R.W (1999) The Social Robot Architecture: Towards Sociality in a Real World Domain, *Towards Intelligent Mobile Robots 99, Bristol, UK.*

[Shen, 1996] Wei-Min Shen. YODA, Young Observant Discovery Agent, *Proc. AAAI-96*, Portland, Oregon.

[Steels, 1994] Luc Steels. Building Agents with Autonomous Behavior Systems, The `artificial life' route to `artificial intelligence'. *Building situated embodied agents. Lawrence Erlbaum Associates, New Haven. 1994*