Research Repository UCD

Title	Robot soccer anywhere: achieving persistent autonomous navigation, mapping and object vision tracking in dynamic environments
Authors(s)	Dragone, Mauro, O'Donaghue, Ruadhan, Leonard, John J., O'Hare, G. M. P. (Greg M. P.), Duffy, Brian R., Patrikalakis, Andrew, Leederkerken, Jacques
Publication date	2005-07-13
Publication information	Dragone, Mauro, Ruadhan O'Donaghue, John J. Leonard, G. M. P. (Greg M. P.) O'Hare, Brian R. Duffy, Andrew Patrikalakis, and Jacques Leederkerken. "Robot Soccer Anywhere: Achieving Persistent Autonomous Navigation, Mapping and Object Vision Tracking in Dynamic Environments." SPIE--The International Society for Optical Engineering, 2005.
Conference details	SPIE Opto-Ireland 2005, Dublin, Ireland, April 4, 2005
Publisher	SPIE--The International Society for Optical Engineering
Item record/more information	http://hdl.handle.net/10197/4485
Publisher's statement	One print or electronic copy may be made for personal use only. Systematic electronic or print reproduction and distribution, duplication of any material in this paper for a fee or for commercial purposes, or modification of the content of the paper are prohibited.
Publisher's version (DOI)	10.1117/12.608404

Downloaded 2023-04-07T10:52:29Z

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)

© Some rights reserved. For more information

Robot Soccer Anywhere: Achieving Persistent Autonomous Navigation, Mapping and Object Vision Tracking in Dynamic Environments

Mauro Dragone^a, Ruadhan O’Donoghue^a, John J. Leonard^{a,b}, Gregory O’Hare^a, Brian Duffy^a, Andrew Patrikalakis^b, and Jacques Leederkerken^b

^aUniversity College Dublin, Belfield, Dublin 4, Ireland;

^bMassachusetts Institute of Technology, Cambridge, MA 02139, USA

ABSTRACT

The paper describes an ongoing effort to enable autonomous mobile robots to play soccer in unstructured, everyday environments. Unlike conventional robot soccer competitions that are usually held on purpose-built robot soccer “fields”, in our work we seek to develop the capability for robots to demonstrate aspects of soccer-playing in more diverse environments, such as schools, hospitals, or shopping malls, with static obstacles (furniture) and dynamic natural obstacles (people). This problem of “Soccer Anywhere” presents numerous research challenges including: (1) Simultaneous Localization and Mapping (SLAM) in dynamic, unstructured environments, (2) software control architectures for decentralized, distributed control of mobile agents, (3) integration of vision-based object tracking with dynamic control, and (4) social interaction with human participants. In addition to the intrinsic research merit of these topics, we believe that this capability would prove useful for outreach activities, in demonstrating robotics technology to primary and secondary school students, to motivate them to pursue careers in science and engineering.

Keywords: Intelligent Agents, Mobile Robots, Navigation, Mapping, Robot Soccer

1. INTRODUCTION

Autonomous robot navigation in dynamic environments has been a central goal of research in the fields of artificial intelligence, robotics, and machine learning for several decades. The dream of researchers in these areas has been to enable mobile robots to operate in complex environments, for long periods of time, with minimal supervision, and with a high degree of reliability. Despite tremendous recent progress, however, the performance of autonomous mobile robots in typical everyday environments is disappointingly poor. In particular, current robots often fall far short in several key dimensions:

- **dynamic environments:** the ability to detect and track moving objects, and to detect and account for long-term changes in the environment;
- **persistence:** the ability to achieve improved performance from repeated operations over time in the same environment, learning from experience;
- **robustness:** the ability of systems to detect failures autonomously, and to recover from them in a reliable manner;
- **social interaction:** the ability to explicitly interact with other robots (and ideally with people as well) to perform complex tasks in unstructured environments.

To spur progress along these critical research dimensions, we have posed the challenge of enabling a team of robots to play soccer (football) in generic (unmodified) indoor environments. Unlike RoboCup soccer,¹ in which teams compete on color-coded pitches whose specifications are provided *a priori*, in our work we focus on operations in environments not specifically prepared for robots playing soccer. We pose the question: is it possible to simply show up with a team of robots at a school or hospital, turn the robots on, and “leave them to it” to demonstrate robotic soccer playing in a classroom or hallway.

As robots become more pervasive in our everyday environment, our interaction with these autonomous devices becomes inevitable. This necessitates a robot control paradigm which facilitates the development of some degree of social partnership between robots and between a robot and a human. This work removes many of the environmental constraints found in the RoboCup competitions, and introduces the human as an active participant in free soccer interaction in unstructured environments.

In addition to the numerous research challenges confronted in developing such a capability, a second motivation for work in this area is public outreach. Soccer Anywhere offers an ideal opportunity to excite and inspire young students to pursue careers in science technology. Mobile robots provide a compelling physical embodiment of many interesting and challenging aspects of information and communications technology. With a capability for Soccer Anywhere, we can connect directly with the everyday world of many young students. (Clearly, a sizable fraction of the Irish elementary and secondary population have a great familiarity with soccer at the everyday level, and this is perhaps equally true in many countries around the globe.)

Rather than the more structured approach being pursued by RoboCup (akin to professional soccer matches), in our work we are aiming at the unstructured simplicity of more “informal” soccer matches being played by as few as two people. The simple skills exhibited by primary school students in an “informal” game of soccer in a schoolyard are still far beyond the capabilities of many current robots.*

Clearly many technologies developed for RoboCup can be applied with success in less structured environments. Recent examples of relevant research pursued in the RoboCup domain include the work of Veloso and colleagues at CMU working with simulated robots, small-size platforms, legged robots, and Segway robots,^{2,3} and Röfer *et al.*, which describes the design of the World Champion German Team in the 2004 legged robot competition.⁴

While probabilistic localization given a prior map has been implemented for RoboCup, for example by Lenser and Veloso,⁵ to our knowledge, we are unaware of previous work that has integrated a real-time SLAM system with soccer-playing robots.[†] One can perhaps argue that full-blown SLAM is “overkill” for the soccer task, as purely reactive strategies will often suffice for implementation of basic single and multi-robot soccer behaviors. We believe, however, that the payoff from the integration of environment learning and adaptability into the agents will be significant. Imagine, for example, the following scenario. A team of robots is brought to a primary school, with no prior information about the environment supplied in advance. The robots are turned on and set out to explore the environment and to build a map of one the main corridors of the school and several of the adjacent classrooms. Subsequently, a soccer ball is introduced, and the robots proceed to play a soccer game, using the hallway as a playing field (pitch) and the doorways to two of the classrooms as “goals”.

Multi-robot coordination has been a very popular topic for research, both in RoboCup and in other domains. For motion coordination of multiple moving agents, examples include Stroupe *et al.*, who have developed Move Value Estimation for Robot Teams (MVERT), an action selection algorithm successfully demonstrated for cooperative target tracking,⁸ and Hidaka *et al.*, who have implemented optimal formation control for cooperative localization of mobile robots.⁹ Balch and Arkin have implemented reactive, behavior-based formation control for multi-robot teams.¹⁰ Spletzer and Taylor have developed dynamic sensor planning and control methods for optimal target tracking by multiple mobile robots¹¹

To work towards this capability of social robot cooperation in dynamic, unknown environments, we have developed a set of four low-cost mobile robots at UCD. Each system consists of motor control components from Evolution Robotics, Inc., a single-board linux computer with 802.11b wireless, and a low-cost webcam. Two of the four robots are equipped with SICK LMS laser scanners. Additionally, our effort makes use of four Nomad Scout robots, equipped with vision and sonar sensing. Currently, we are working to provide these robots with the ability to play soccer in campus buildings, and our subsequent goal is to take this robots to a number of schools in the area to attempt live demonstrations of this technology.

*Note: our work employs wheeled robots, with simple ball pushing devices. At present we completely neglect the mechanics of real physical soccer, namely the creation of legged robots that can kick, run, and jump, and focus instead on the sensing, navigation, and multi-robot task coordination aspects of the robot soccer task.

[†]A task formulated as “Almost SLAM” was provided by Röfer as a technical challenge for the 2004 legged robot competition,⁶ and was attempted by at least one team⁷

Figure 1. Left: UCD Soccer Anywhere robot. Right: SLAM map of first floor of the UCD Computer Science Building, produced using Atlas¹² with laser scan matching.

We now proceed to provide a brief overview, with pointers to the relevant literature, for each the two key technologies being integrated in our project: SLAM and the Social Robot Architecture.

2. SIMULTANEOUS LOCALIZATION AND MAPPING (SLAM)

The starting point in our approach is to provide the capability for a single mobile robot to build a map of an unknown environment while concurrently localizing itself. This problem, known as SLAM,^{13,14} is now fairly well understood for single mobile robots operate in static environments of moderate size and complexity. A range of methods are available in the literature, and a recent extensive treatment survey is provided by Thrun *et al.*¹⁵ Two lines of recent research in SLAM are most relevant to the work presented here: (1) incorporation of tracking of moving objects into SLAM, and (2) the development of SLAM for multiple vehicles. Wang *et al.*¹⁶ and Hähnel *et al.*¹⁷ have developed methods that integrate tracking of moving objects into large-scale SLAM algorithms applied in 3-D environments. Detection of longer-term changes in the environment for persistent long-term operations in an open problem for research. There have been quite a few efforts that have addressed SLAM multiple vehicles; examples include Howard,¹⁸ Konolige,¹⁹ Nettleton *et al.*,²⁰ Thrun and Y. Liu,²¹ and Walter and Leonard,²²

The method for SLAM that we employ in our research is called Atlas, which was developed by Bosse, Newman, Teller and Leonard at MIT.¹² In this approach, efficient large-scale mapping is performed using a network of multiple local submaps. Figure 1(b) shows the output of Atlas using laser data for one of our robots mapping the first floor of the UCD Computer Science Building, which currently serves as our robot soccer test environment.

For Soccer Anywhere, SLAM provides an ability for multiple agents to “ground” their observations and actions with respect to the geometry of the local environment. We envision that by sharing maps, and making direct observations of the ball and of one another’s positions, that teams of robots can coordinate their actions, for example choosing real-world locations such as doorways as “goals”, and subsequently coordinating their actions to pass the ball back and forth to deliver the ball to the goal. This process must unfold in a manner that ensures the safety of other robots and people nearby, as well as without causing any physical damage to the environment.

Figure 2. The Social Robot Architecture, which consists of four elements: (1) the physical-sensorial level, (2) the reactive level, (3) the deliberative level and (3) the social level.²³

To enable cooperative and persistent operations, we are extending Atlas to allow repeated experiments over time in the same environment, by one or more robots. In this approach, submap networks from previous missions by either the same robot or a different robot can be pre-loaded into Atlas, and “relocation” is performed to combine the Development and testing of this capability is currently in progress. The task of coping with large databases of maps, accumulated over time, presents some interesting problems for future investigation, such as classifying long-term static environmental features, such as walls, such as objects that move, such as furniture.

3. HYBRID CONTROL ARCHITECTURE FOR DISTRIBUTED NETWORKS OF MOBILE AGENTS

To provide coordination and communication between different robots, our work makes use of the Social Robot Architecture (SRA), developed at UCD by Duffy and colleagues to facilitate explicit social interaction between a collective of heterogeneous robots.^{23,24} One of the demonstrators of the architecture was Robot Soccer, but the architecture was designed as a more general framework control of social interactions among multiple agents.

The SRA is a layered structure with four fundamental elements: the physical-sensorial level, the reactive level, the deliberative level and the social level. (See Figure 2.) The control architecture offers behavioural robustness through basic reflexes which guide robot responses to unexpected or dangerous events (e.g. halting in case of emergency or recovery from a collision). Such reflexes operate in parallel with more sophisticated *deliberative* behaviours such as obstacle avoidance, wall following, and for the robotic soccer scenario, dribbling, tracking and passing the ball. The Social Level provides for explicit social interaction through the agent communication language Teanga as described in Duffy et al. This is supported by social modeling mechanisms implemented at the Deliberative Level.

While a number of researchers have recognized the desirable attributes of hybrid architectures that combine deliberative and reactive control approaches, a distinctive aspect of the Social Robot Architecture is the *explicit* incorporation of social interaction. In contrast, many other architectures employ only three layers. For example, Connell proposed a three level architecture known as the SSS architecture, that consisted of servo, subsumption, and symbolic layers.²⁵ Numerous other examples are described in the literature; Kortenkamp *et al.*, for example, provides a good overview of many related research efforts.²⁶

Our robot soccer implementation offers basic reflexes guiding robot responses to unexpected or dangerous events (e.g. halting in case of emergency or recovery from a collision) together with more sophisticated behaviours

Figure 3. Reactive layer of the Social Robot Architecture, which provides fast local response for tasks such as obstacle avoidance and ball-tracking.

such as obstacle avoidance, wall following, and for the robotic soccer scenario, dribbling, tracking and passing the ball.

In implementing the SRA, our work makes use of Agent Factory^{27, 28} (<http://www.agentfactory.com>), an integrated and toolled environment for the rapid prototyping of social intentional agents based upon belief-desire-intention agent theory.²⁹ Agent Factory provides a distributed environment in which multiple agents can employ practical reasoning techniques to deliberate upon their perceived situation, update their mental state and select the future line of action.

To enable collaboration among social agents, Agent Factory agents make use of speech act theory,³⁰ a formalism for accurate and expressive communication mechanism in multi-agent systems. This is done by performing a speech act (such a requesting, ordering, informing or promising), that is sending a message to one or more of their acquaintance in order to affect their mental states.

The present implementation of our behaviours for autonomous navigation is based on the curvature-velocity method.³¹ Obstacles in the environment are detected merging information coming from the onboard camera and from the particular combination of range sensors adopted from the robotic platform (sonar, infrared sensors and/or laser scanners). Goal-oriented behaviours are delivered through the intentional agent structures of Agent Factory.

Figure 3 depicts the reactive level which has the responsibility to monitor the onboard sensors (through modules called perceptrors). The sensory information is processed at this level from activities running in a multi-thread within the reactive controller. This results in clear agent events being generated and subsequently communicated to the agent’s deliberative level. For example, `sense(ball)` notifications from the reactive layer (perceptions) are converted into beliefs and augmented into the belief set providing the agent with an up-to-date model of its current perceived situation. This situation is then analyzed and the agent’s commitments are updated. Pre-existing commitments are analyzed and those pertaining to the current time frame are honored, 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.

For instance, a soccer playing robot may contain a simple rule such as:

```
Belief(sense(ball)) and Belief(close(ball)) => Commit(reach_object(ball)).
```

In this example, the agent, upon making visual contact with a nearby ball, will commit itself to reach for the ball activating the proper plan (in this case the single reactive behaviour `reach_object`).

Our implementation of the SRA benefits greatly from the virtual robotic workbench (VWB), a lightweight Collaborative Virtual Environment framework comprising an XML based, peer to peer networking protocol and

Figure 4. Graphical user interface of the virtual robotic workbench, used both for simulations and monitoring and control of hardware experiments. This example illustrates collaborative operation of two Nomad Scout mobile robots. The interface provides visualization of robot sensor data, communications between robots, and image feedback from cameras mounted on each robot, a full 3D visualization of the scene, and teleoperation control of any/all robots if desired.

a multimedia visualization suite. Operation of the virtual robotic workbench is illustrated in Figure 4, for an experiment involving two Nomad Scout mobile robots employed in a soccer task. The goal of the virtual robotic workbench is to offer researchers a common platform that can ease integration and replication issues. In contrast to traditional client-server solutions the virtual robotic workbench is modelled on a Collaborative Environment paradigm (as in IEEE-Std-1278-1993, *Standard for Distributed Interactive Simulation, DIS*) characterized by peer to peer communication between nodes replicating a common design. In contrast to other systems, we push the standardization one step further, in a collaborative environment paradigm embracing every aspect of our work from data communication to visualization tools. Each experiment is seen as a collaborative experience in which information about activities and interaction between all participants (robots, experimenters and environment) are collected and made available, to various degrees, within the framework.

4. EXPERIMENTAL RESULTS

To demonstrate the integration of SLAM with distributed agent control architecture summarized above, in this section we describe a representative experimental result achieved thus far, in which robotic ball tracking has been performed concurrently with SLAM running in real-time on a single robot.

The agent design script that was utilized in this experiment is shown Figure 5. The script is written in Agent Factory Agent Programming Language (AF-APL).²⁸ The first section of the script declares the set of six actuators and six perceptrs used by the agent. The next section of the script specifies the set of six commitment rules that drive the actions of the agent. The first commitment rule installs the drivers for the specific hardware of the robot and the behaviors that are active on the robot. The second rule activates the “stop” behavior on startup. The third commitment rule specifies a commitment rule that will be activate when the robot is close to

```

ACTUATOR PRISM.RobotAgent.InstallModulesActuator
ACTUATOR PRISM.RobotAgent.ActivateBehaviorActuator
ACTUATOR PRISM.RobotAgent.VRWRequestActuator
ACTUATOR PRISM.RobotAgent.VRWInformActuator
ACTUATOR actuators.AdoptActuator
ACTUATOR PRISM.RobotAgent.Talk
PERCEPTOR perceptors.SelfPerceptor
PERCEPTOR perceptors.TimePerceptor
PERCEPTOR PRISM.RobotAgent.VRWPerceptor
PERCEPTOR PRISM.RobotAgent.VRWMessagePerceptor
PERCEPTOR PRISM.RobotAgent.ActionPerceptor
PERCEPTOR PRISM.RobotAgent.EventsPerceptor

BELIEF(start) => COMMIT(Self,Now,setup(PRISM.RobotCtrl.drivers.MoosLaser,
    PRISM.RobotCtrl.drivers.Camera,PRISM.RobotCtrl.Behaviors.MoveFree,
    PRISM.RobotCtrl.Behaviors.MoveTo,PRISM.RobotCtrl.Behaviors.FollowRightWall,
    PRISM.RobotCtrl.Behaviors.FaceObject,PRISM.RobotCtrl.Behaviors.Step,
    PRISM.RobotCtrl.Behaviors.Stop,PRISM.RobotCtrl.Behaviors.TurnRight,
    PRISM.RobotCtrl.Behaviors.TurnLeft,PRISM.RobotCtrl.Behaviors.Scan,
    PRISM.RobotCtrl.Behaviors.Wander,PRISM.RobotCtrl.Behaviors.Die))

BELIEF(start) => COMMIT(Self,Now,ActivateBehavior(Stop))

BELIEF(sensing(ball)) & BELIEF(close(ball)) & !BELIEF(current(FaceObject))=>
    COMMIT(Self,Now,ActivateBehavior(FaceObject(Object,ball,MaxVelocity,0,w,450,
        aw,300,PID.P,1000,PID.D,0.2,PID.I,0)))

BELIEF(sensing(ball)) & BELIEF(close(ball)) & !BELIEF(current(FaceObject)) =>
    COMMIT(Self,Now,Talk("I am close to the ball",1,100,200))

BELIEF(sensing(ball)) & BELIEF(distant(ball)) & !BELIEF(current(Wander))=>
    COMMIT(Self,Now,ActivateBehavior(FaceObject(Object,ball,MaxVelocity,100,w,450,
        aw,300,PID.P,1000,PID.D,0.2,PID.I,0)))

BELIEF(sensing(ball)) & BELIEF(distant(ball)) & !BELIEF(current(Wander)) =>
    COMMIT(Self,Now,Talk("I am distant, I move toward the ball",1,100,200))

```

Figure 5. Agent design script, written in Agent Factory Agent Programming Language (AF-APL), for an experiment integrating real-time SLAM and soccer ball tracking.

the ball. It will cause the robot to turn toward the ball (if it is not already facing it), and specifies the control parameters for that behavior. In this case, the behavior is implemented with a PID controller, with the specified parameters. The next rule, for the same conditions, activates the text-to-speech specify to provide feedback on its next action. The final two rules for this example will be active when the robot is far from the ball (greater than one meter from the ball); in this case, the robot will also move forward towards the ball, with a maximum velocity of approximately 25 cm per second. An appropriate text-to-speech message is provided with the final rule.

The resulting map is shown in Figure 6. This map corresponds to an experiment of 10 minutes in duration, during which the robot made repeated traversals of the corridor, guided by a human moving the soccer ball through the environment. Movies of the experiment can be accessed on the web at: <http://oe.mit.edu/~jleonard/robotsoo>

Figure 6. Map of part of the UCD Computer Science Building produced by a robot executing soccer ball tracking behaviors while concurrently performing SLAM.

5. CONCLUSION

This paper has described an ongoing effort to develop a team of robots capable of playing soccer in unstructured environments such as schools, hospitals, and office buildings. This problem has many shared aspects with conventional RoboCup soccer, but also poses additional challenges of acquiring and sharing maps of unknown environments and operating in environments with mobile and static obstacles (such as people and furniture). Map acquisition and robot pose tracking is performed using the Atlas framework¹² for large-scale SLAM. The coordination of multiple agents is achieved via the Social Robot Architecture, a distributed agent based environment for heterogeneous networks of agents with mobility.²⁴ These efforts are greatly facilitated by two valuable software tools, the virtual robot workbench developed by Dragone at UCD, and the Mission Oriented Operating Suite developed by Newman at MIT and Oxford.³²

Our accomplishments to date include: (1) hardware development of a team of low-cost mobile robots; (2) software integration of MOOS with the virtual robotic workbench; and (3) integration of SLAM for a single agent with basic soccer-playing behaviors, such as ball tracking. This has enabled us, for example, to “drive” a robot around the computer science building at UCD, by moving a soccer ball through the environment, with the mobile robot tracking the ball while concurrently using SLAM to map the environment.

Our work in progress seeks to achieve real-time cooperative ball tracking and ball passing, sharing of maps between agents, and long-term persistent autonomous operations. Achieving this will require progress on dealing with communications latencies and bandwidth restrictions. Additionally, we are investigating the trade-offs between agent-centered vs. world-referenced representations, and reactive vs. symbolic decision-making for mobile distributed sensor networks.

ACKNOWLEDGMENTS

The authors gratefully acknowledge support from Science Foundation Ireland. This research is based upon works supported by the Science Foundation Ireland under Grant No. 03/IN.3/1361 and under an E.T.S. Walton Visitor Award held by J. Leonard.

REFERENCES

1. H. Kitano, M. Asada, Y. Kuniyoshi, I. Noda, and E. Osawa, "RoboCup: The robot world cup initiative," in *Proceedings of the First International Conference on Autonomous Agents (Agents'97)*, W. L. Johnson and B. Hayes-Roth, eds., pp. 340–347, ACM Press, (New York), 5–8, 1997.
2. M. Veloso, "Autonomous robot soccer teams," **33**(1), pp. 8–12, 2003.
3. M. Veloso, "Robot soccer." <http://www-2.cs.cmu.edu/~robosoccer/main/>.
4. T. Röfer, I. Dahm, U. Düffert, J. Hoffmann, M. Jüngel, M. Kallnik, M. Löttsch, M. Risler, M. Stelzer, and J. J. Ziegler, "GermanTeam 2004," in *Eight International Symposium on RoboCup (Robot World Cup Soccer Games and Conferences)*, Springer, 2005. To appear.
5. S. Lenser and M. Veloso, "Sensor resetting localization for poorly modelled mobile robots," 2000.
6. T. Röfer, "Technical challenges for the RoboCup 2004 legged league competition," 2004. <http://www.tzi.de/4legged/bin/view/Website/History>.
7. T. Röfer, 2005. Personal Communication.
8. A. Stroupe, R. Ravichandran, and T. Balch, "Value-based action selection for exploration and dynamic target observation with robot teams," in *IEEE International Conference on Robotics and Automation*, (New Orleans, Louisiana), Apr 2004.
9. Y. S. Hidaka, A. I. Mourikis, and S. I. Roumeliotis, "Optimal formations for cooperative localization of mobile robots," in *Proc. IEEE Int. Conf. Robotics and Automation*, (Barcelona, Spain), April 2005.
10. T. Balch and R. Arkin, "Behavior-based formation control for multi-robot teams," 1997.
11. J. Spletzer and C. J. Taylor, "Dynamic sensor planning and control for optimally tracking targets," *Int. J. of Robot. Res.* **22**, pp. 7–20, 2003.
12. M. Bosse, P. Newman, J. Leonard, and S. Teller, "An atlas framework for scalable mapping," *Int. J. of Robot. Res.* **23**, pp. 1113–1139, December 2004.
13. R. Smith, M. Self, and P. Cheeseman, "A stochastic map for uncertain spatial relationships," in *4th International Symposium on Robotics Research*, MIT Press, 1987.
14. J. J. Leonard and H. F. Durrant-Whyte, "Simultaneous map building and localization for an autonomous mobile robot," in *Proc. IEEE Int. Workshop on Intelligent Robots and Systems*, pp. 1442–1447, (Osaka, Japan), 1991.
15. S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*, MIT Press, 2005.
16. C.-C. Wang, C. Thorpe, and S. Thrun, "Online simultaneous localization and mapping with detection and tracking of moving objects: Theory and results from a ground vehicle in crowded urban areas," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2003.
17. W. B. D. Hähnel, D. Schulz, "Map building with mobile robots in populated environments," in *Proc. IEEE Int. Workshop on Intelligent Robots and Systems*, 2002.
18. A. Howard, "Multi-robot mapping using manifold representations," in *IEEE International Conference on Robotics and Automation*, pp. 4198–4203, (New Orleans, Louisiana), Apr 2004.
19. K. Konolige, D. Fox, B. Limketkai, J. Ko, and B. Stewart, "Map merging for distributed robot navigation," in *Proc. IEEE Int. Workshop on Intelligent Robots and Systems*, 2003.
20. E. Nettleton, S. Thrun, H. Durrant-Whyte, and S. Sukkarieh, "Decentralised slam with low-bandwidth communication for teams of airborne vehicles," in *Proceedings of the International Conference on Field and Service Robotics*, (Lake Yamanaka, Japan), 2003.
21. S. Thrun and Y. Liu, "Multi-robot SLAM with sparse extended information filters," in *Proceedings of the 11th International Symposium of Robotics Research (ISRR'03)*, Springer, (Sienna, Italy), 2003.
22. M. R. Walter and J. J. Leonard, "An experimental investigation of cooperative SLAM," in *Proceedings of the Fifth IFAC/EURON Symposium on Intelligent Autonomous Vehicles*, (Lisbon, Portugal), July 2004.
23. B. R. Duffy, *The Social Robot*. PhD thesis, National University of Ireland, University College Dublin, 2000.
24. B. R. Duffy, G. M. P. O'Hare, R. O'Donoghue, C. Rooney, and R. Collier, "Reality and virtual reality in mobile robotics," in *1st International Workshop on Managing Interactions in Smart Environments MANSE'99*, 1999.

25. J. H. Connell, "SSS: A hybrid architecture applied to robot navigation," in *Proc. IEEE Int. Conf. Robotics and Automation*, pp. 2719–2724, IEEE, 1992.
26. D. Kortenkamp, R. P. Bonasso, and R. Murphy, eds., *Artificial Intelligence and Mobile Robots*, AAAI Press, 1998.
27. G. M. P. O'Hare, "The agent factory: An environment for the fabrication of distributed artificial systems," in *Foundations of Distributed Artificial Intelligence*, G. M. P. O'Hare and N. R. Jennings, eds., Wiley Interscience Publishers, 1996.
28. R. W. Collier, *Agent Factory: A Framework for the Engineering of Agent-Oriented Applications*. PhD thesis, National University of Ireland, University College Dublin, 2001.
29. A. S. Rao and M. P. Georgeff, "Modelling rational agents within a BDI architecture," in *Principles of Knowledge Representation and Reasoning*, San Mateo, CA, 1991.
30. J. Searle, *Speech Acts: an Essay in the Philosophy of Language*, Cambridge University Press, 1969.
31. R. Simmons, "The curvature velocity method for local obstacle avoidance," 1996.
32. P. Newman, "MOOS — a mission oriented operating suite," tech. rep., Department of Ocean Engineering, Massachusetts Institute of Technology, 2002.
33. J. Bruce, T. Balch, and M. Veloso, "Fast and inexpensive color image segmentation for interactive robots," 2000.
34. J. L. de la Rosa, A. Oller, J. Vehi, and J. Puyol, "Soccer team based on agent-oriented programming," *Robotics and Autonomous Systems* (21), p. 176, 1997.