

An Experimental Study of Anticipation in Simple Robot Navigation

Birger Johansson¹ and Christian Balkenius¹

Lunds University Cognitive Science, Kungshuset Lundagård
222 22 Lund, Sweden

{birger.johansson, christian.balkenius}@lucs.lu.se

Abstract. This paper presents an experimental study using two robots. In the experiment, the robots navigated through an area with or without obstacles and had as their goal to shift places with each other. Four different approaches (random, reactive, planning, anticipation) were used during the experiment and the times to accomplish the task were compared. The results indicate that the ability to anticipate the behavior of the other robot could be to an advantage. However, the results also clearly show that anticipatory and planned behavior is not always better than a purely reactive strategy.

1 Introduction

The ability to anticipate the behaviors of others is something we take more or less for granted and we often do not appreciate the complexity of this ability. When attempting to build robots with anticipatory abilities, it becomes clear that this is far from trivial. Not only does the robot need to control its own movement, it also needs to predict what other robots or possibly humans will do. Moreover, it needs to use the anticipated behaviors of others in a sensible way to change its own behavior.

Consider the following real life situation of the near future. You have sent your personal shopping robot to the “Autonomous supermarket” to get your favorite chocolate cake. To get the cake, located at the other end of shop, your robot cannot chose the straight path toward the cake because of the shelves and other obstacles, including all the other personal shopping robots in the store. Instead some alternative strategy must be used.

One possibility would be to move around randomly in the store until it finds the chocolate cake, but this would probably result in a long period of aimless wandering before it gets to its goal. This type of random behavior is very inefficient and is seldom used in robot navigation, although is rather common in robot exploration.

It is obvious that better methods can be used. Instead of moving at random, the robot may try to move in the direction of the goal. This is a reactive place approach method where the robot reacts to the position of the chocolate cake and selects actions accordingly [2]. The problem with this approach is that the

robot cannot go straight to the goal because of the shelves and other robots in the store. It needs to apply an obstacle avoidance strategy when there is something in its way. For example, it may turn around and move in some other direction for a short while before turning toward the goal again. This type of reactive navigation has been widely used in robotics where the relation between the stimulus and response is often preprogrammed [8][13]. A number of rules are set up that must be fulfilled for an action to be executed. A problem with such a reactive approach to navigation is that the robot can easily get itself into situations where it becomes trapped.

Although the reactive strategy is more efficient than random movements it would be better to plan a path around the shelves based on knowledge of the layout of the store. This has traditionally been the most common way of dealing with robot navigation. This plan can use grids [21], potential fields [21][1][7], or some symbolic or geometric description of the environment. As long as the map of the shop is correct, the plan will also be correct and can be used to efficiently go to the cake.

Unfortunately, when the personal shopping robot reviews the map after a few seconds of moving according to the planned path, it realizes that the map is no longer accurate. The shelves are still where they are supposed to, but most of the other robots have moved and are not where the map indicates. As with the previous strategy, this makes it necessary to use some obstacle avoidance strategy to avoid colliding with the other robots which may limit the usefulness of the plan.

The solution to this problem is to include the movements of the other robots when the personal shopping robot makes its plan. This is however not trivial as it does not know where the other robots are heading. One reasonable assumption is that they will continue in the direction they have now, although this will only be true for a short while. Better predictions can be made if the robot knows the goals of the other robots. By anticipating the behaviors of others, it will be able to choose a better path and do not have to use the obstacle avoidance strategy as often. The better its ability to anticipate, the less it will need to use its alternative strategies.

Several different types of anticipatory behavior has been used in robotics and AI. First, it is possible to use an anticipatory mechanism to reduce the latency of a control system. For example, Behnke et al. [4] used neural network to reduce the control latency for the FU-fighter team in RoboCup. The control system had a delay of four frames (132 ms) and with a speed up to 2 m/s, this could result in an error between the actual position of the robot and the tracking of the robots of approximately 20 cm. By feeding a neural network with the position, orientation and motor commands from the last six frames to anticipate the current position, the influence of delay in the system was almost eliminated. A similar method has also been used to predict the location of a moving target for visual tracking [3].

A second type of anticipation concerns anticipation of the environment, for example the movement of other robots. Sharifi et al. [19] describe a system for

the simulation league of RoboCup where the future state is used to anticipate which robot will possess the ball next, while Veloso et al. [22] anticipate the state of the whole team. This means that a seemingly passive agent is not passive at all. Instead it actively anticipates opportunities for collaboration.

The anticipation of robot movement can also be based on observation. For example, Stulp et al. [20] model the goal keeper in RoboCup to be able to anticipate its behavior. Ledezma et al. [15] used a similar method to model the behavior of the other players based on their observed input and outputs. Usually, some type of communication between agents are used in anticipation, either a complete knowledge of world or broadcasting of individual plans but there is also work on cooperation without sharing information between agents [18].

Human-robot interaction can also merit from using anticipatory behavior. Sabanovic et al. [17] used a stationary robotic receptionist that provided information to visitors and enhances interaction through story-telling to study human-robot interaction. In this study, the robot receptionist turns toward people passing by and tries to interact with them. To be able to interact in an efficient way, the robot receptionist anticipates the position of people passing by.

The importance of anticipation has also been studied in the domain of computer games [14]. In human activities, Saad [16] pointed out the close connection between driving and anticipation, even stating that “driving is anticipating”.

Davidsson [9] used simulations to investigate the benefits of anticipation. Two different types of experiments were conducted. The first investigated competition between agents and in the second experiment, the agents were cooperative. In the experiments, the task of the agents was to pick up targets in a two dimension grid world in a particular order. By using a linearly quasi-anticipatory agent architecture, one agent could realize that it would not reach the target before the other agent and would instead start to move toward the following target. Only one of the robots used anticipatory behavior while the other one used reactive behavior. In the second experiment, the agents cooperate which lead to a decreased total time for fetching all target objects.

Although simulations can be very valuable in testing different strategies, a simulation must necessarily include a perfect model of the simulated environment. It will thus always be possible to make perfect predictions in a simulation if this is desired. It is well known that this can easily lead to solutions that are not useful when applied to robots that have to operate in the real world [6].

To evaluate the benefit of anticipation in mobile robots in the real world, we tested a number of strategies in three different environments with two robots. We compared a random and a reactive strategy with control methods based on planning with or without anticipation of the behavior of the other robot. The goal of the experiments was to test under what conditions the ability to anticipate would help the robots in a simple task. In addition, we tested three different methods to use the anticipated behavior of the other robot.




FIG. 1: *The three environments with different complexity. A and B: goal locations for the robots.*

2 Description of the Experimental System

2.1 Environment

The size of the experimental area was 2×2 m. Bricks marked with white color were placed in the area in two of the experiments as obstacles (Fig. 1). Experiment 1 used the empty environment, Experiment 2 used an environment with walls and in Experiment 3, obstacles were placed at random in the environment.

2.2 Robots

The robots used were two modified BoeBots (Parallax Inc., Rocklin, California). These robots are approximately 14 cm long and use a differential steering. No sensors on the robots were used in these experiments. Instead, each robot was marked with two colors that could be detected by a camera mounted 3.5 m above the robot area. This camera transmitted images to a computer that calculated the position and orientation for the two robots four times per second. This computer also performed all the computations for the two robots and controlled the robots via wireless bluetooth communication. In addition, it stored tracking data and collected all statistics for the experiments.

2.3 Control Systems

The control systems of the robots were built using the Ikaros framework¹. The interface components used included processing of the the video stream from the camera, color tracking to detect the position and orientation of the robots, and wireless communication. In addition, modules where built for reactive robot control, path planning, and anticipation.

¹ www.ikaros-project.org




FIG. 2: *The computer is using the overhead camera to track the two robots and transmits motor commands via bluetooth.*

Random Control A random control system was the first tested in the experiments. This system simply transmits random motor commands to the robot until it has reached the goal. The robot is instructed to turn toward a random orientation and then travel in this direction until an obstacle activates an obstacle avoidance system, in which case a new random direction is set. This is repeated until both robots have reached their goals.

Reactive Approach The next control system performs reactive approach where the robot always tries to go directly toward the goal. The desired path is calculated as the straight line between the current location of the robot and the goal location. This strategy will obviously have problems when there are obstacles in the way and to handle this situation a reactive avoidance system was added.

Planning System The planning system is responsible for path finding within the environment. To accomplish this an A* based navigation algorithm is used [12]. This is a grid-based navigation algorithm with full knowledge of the environment. It finds the shortest path to the goal by testing it in the grid-map. If it is unable to use the shortest path, the second shortest path is tested and so on, until a path has been found. Each robot uses the algorithm to find the best path through the robot area. The grid-map is divided into 32×32 elements with a status of either occupied or free. The planning system takes no account of where the other robots are located and only uses its own position, the desired position and the grid-map to find the path.

Anticipation System The anticipation system is similar to the planning system but also includes the movements of the other robots. If the other robots were stationary, the A* algorithm could register the other robots as obstacles. When the other robots are moving it becomes necessary to anticipate their position at each time-step in the future. To solve this, each robot has a model of the other robot. This model is built using each robot's own planning system, for example, robot A assumes that robot B would use the path that robot A would have used if it were located at the position of robot B and heading for the goal of robot B. Before robot A tries to find its own path, it updates its model of the other robot and uses this to find the path for robot B by stepping forward in the planning and checking if there is any collision. If there is a collision, the robot chooses an alternative path and tests if this is a valid. This is repeated until a valid path is found. It should be stressed that the individual paths are not shared between the robots. Only the start and goal positions are known by the other robot. With noise in the system this could lead to inaccurate models of the other robot and this could in turn lead to more activation of the reactive avoidance system. A similarly approach was presented by Guo [11].

An obvious problem arises with this approach. If both robots use the same method to find a valid path, it is possible for both robots to select the alternative path which will result in a collision. A way to avoid this problem is to assign a rank [10] to each robot where the robot with the highest rank always takes the shortest path. For example, let the robot with the longest distance to the goal have the higher rank and let the other robot replan its path around the more highly ranked robot. If the present robot has the lowest rank, we let A* see the other robots as an obstacle but only during that time step. This means that at just that time step there is an object at that position at some time steps later the obstacle has moved and the grid that was occupied in the first time step is free again. In the experiments, we tested three different ways to select the rank of each robot, (1) a fixed rank, (2) the robot closest to its goal would have the highest rank, and (3) the robot with the larger distance to its goal would receive the highest rank. Note that according to the last two strategies, the ranks of the robots may change when the robots move.

Reactive Avoidance A reactive avoidance system is placed on top of the other navigation systems and is activated if there is an obstacle too close to the robot. We divided the reactive area around the robot into 8 regions (Fig. 3). Three in front of the robot, one on each side of the robot and three behind the robot. The robot performs different types of avoidance behaviors depending on in which regions the obstacle was found. If an object is straight ahead, the robot turns on the spot until the obstacle has disappeared from the region and if an object is found to the left of the robot, it steers to the right to obtain a free path. Although the reactive avoidance system mainly helps the robot to reach its goal, it sometimes counteracts the control of the navigation system. For example, when the navigation system instructs the robot to turn right, the reactive avoidance system may detect an obstacle in that area and tell the robot to turn left instead.




FIG. 3: *The robot with the reactive field around it. The reactive field divide the surroundings into eight regions and different avoidance behaviors are activated depending on the location of the obstacle.*

2.4 Experimental Procedure

The task for the robots was to switch places with each other. One robot started at position A and the other started at position B (Fig. 1). When the first robot had arrived at its goal position, it waited for the other robot to reach its goal. The goal locations were subsequently switched and same procedure was repeated. With the switching procedure the robots was forced to interact as their path crossed each other. The presence of interaction is necessary to investigate anticipation in this form.

During this experiment, the time for each position switch was recorded together with the number of times the reactive avoidance system was used. Note that this was a cooperative task where it is the time for both robots to switch places that is recorded.

The worlds were chosen to have multiple possible path to the goal and the robots had to chose one depending on the strategy they currently used although the choose path could change during the experiment depending on the position of the robot. This is a form of coordination problem [5].

Six different strategies were tested: (RAND) random behavior, (APPR) reactive approach behavior, (PLAN) planning, (A-fixed) anticipation with fixed rank, (A-short) anticipation with higher rank for the robot closest to the goal, and (A-long) anticipation with lowest rank for the robot closet to the goal.

Each strategy was tested twice before the robots shifted to the next strategy. When all strategies had been tested two times, the procedure was repeated until in total 40 trials with each strategy had been run. In total, there were 240 trials in each experiment.

Ordinary indoor lightning together with daylight from the large laboratory windows are used as light source for the experiment. The fact that no controlled light source was used could influence the tracking performance even if the tracking algorithm are dynamic for change in light intensity. To overcome this, one trial of each approach was run in sequence instead of all trials in one approach one at the time. So if there was an external influence on the tracking system it influenced all the approaches equally.

To ensure that the tracking system was accurate enough, each iteration that had a position or an angle that was impossible for the robot to reach was counted as a tracking error. This statistic was collected to make sure that the tracking system was reliable and not correlated with the performance of the robots.

In all the experiment the time for the both robots to reach its goal was stored but also the difference in time between the fastest and the slowest robot. Both times were used to compare the difference navigation approaches in each of the different environment. For the statistics, we used the logarithm of the recorded times to minimize the effect of outliers.

3 Results

3.1 Experiment 1

The behaviors of the robots in the different conditions are illustrated in Fig. 4. The environment in experiment 1 did not contain any obstacles. In this environment the tracking system had a mean position error of 1.4% and mean angle error of 4.3%. ANOVA ($F(1, 5) = 297.4, \alpha = 0.05$) showed a significant difference between the different strategies and a post hoc Tukey HSD ($df = 234, \alpha = 0.05$) was used to find the effects of the different strategies. As expected, the random behavior was significantly slower than all the other strategies ($p = 0.000020$). The reactive behavior was also significantly faster than the PLAN and the A-short behavior ($p = 0.0027$ and $p = 0.030$ respectively). There was no significant difference between the different anticipatory strategies.

When comparing the time difference between the fastest and slowest robot, ANOVA ($F(1, 5) = 73.20, \alpha = 0.05$) was used to find that there was a significant difference between the strategies ($p = 0.000020$) and a post hoc Tukey HSD ($df = 234, \alpha = 0.05$) was used to find the effects of the different strategies (Fig. 6). The difference between the times when the two robots reacher their goals was significantly larger for the reactive strategy ($p = 0.000020$).

3.2 Experiment 2

Experiment 2 used an environment with walls. In this environment the tracking error increased due to more obstacles. Shape and shadows from the obstacles make it harder to locate the robot features for exact position and angle tracking. The mean tracking error was 1.7% for the position and 6.8% for the angle of the robots.

ANOVA ($F(1, 5) = 303.9, \alpha = 0.05$) showed a significant difference between the strategies and a post hoc Tukey HSD ($df = 234, \alpha = 0.05$) was used to find the effects of the different strategies. Again, the random behavior was significantly slower than all the other strategies ($p = 0.000020$). The reactive approach strategy was significantly slower than the planning and anticipatory strategies ($p = 0.000020$ in all cases, Fig. 5 middle). There were no significant difference between the planning and the anticipation approaches.

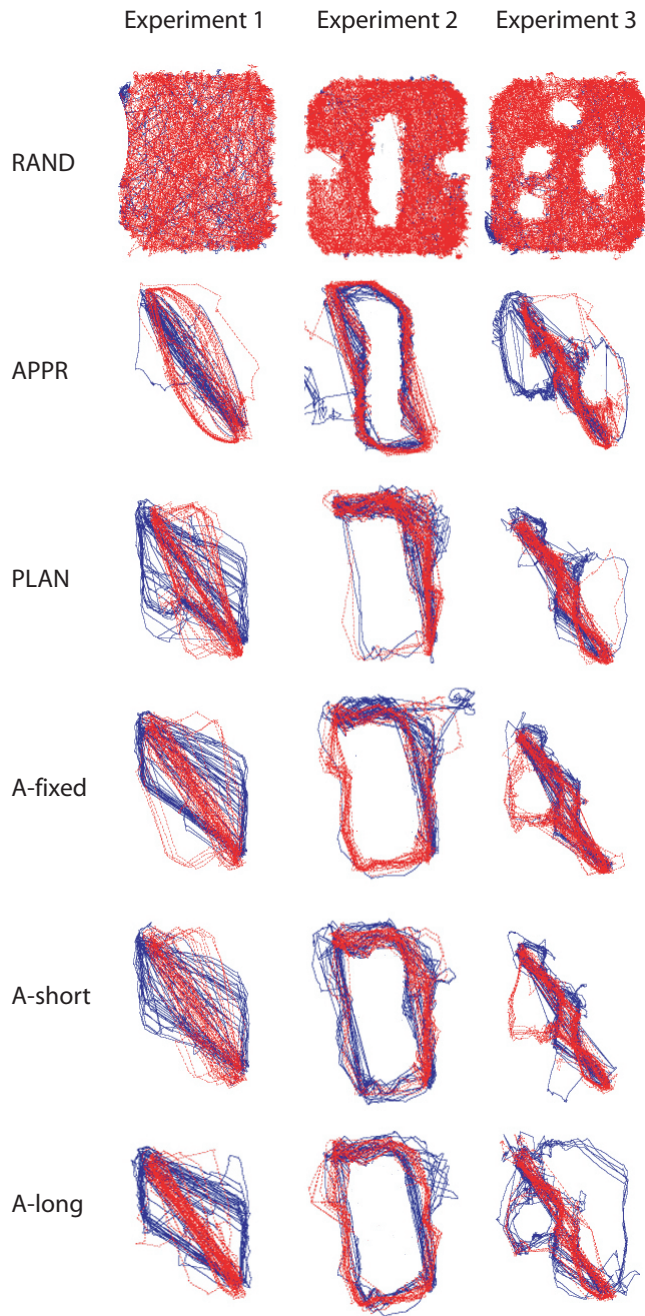


FIG. 4: Illustration of all the movement of the robots. In the experiments with the random behavior, all the available area is covered. It is easy to discern the obstacle location in experiment 2 and 3. Using the reactive approach behavior, less of the area is covered. With this behavior, the required movement has been reduced in comparison to the random approach behavior. Using the planning behavior, the robots will often take the same path which will result in a possible collision and extensive use of the reactive avoidance system. This is most clearly seen in experiment 2 where the robots often both select the top-right path. In the anticipation behaviors, the paths of the robots have more variation because the anticipation causes the robots to use different paths. Note that the robots balance the use of the two paths between the two goal locations. In experiment 1, one robot uses the diagonal path while the other moves to the left or right. The same pattern can be seen in experiment 2 and 3, most clearly in A-long in experiment 3.




FIG. 5: *Left. Time to reach the goals for the different strategies in the empty environment in experiment 1. In the empty environment, the reactive approach behavior (APPR) performed best. The error bars show the standard deviation. Middle. Switch time for the different strategies in the environment with walls in experiment 2. The anticipatory strategy where the robot with longest distance to its goal had highest rank (A-long) was most efficient. Right. Switch time for the different strategies in the environment with random obstacles used in experiment 3. Strategy A-short was fastest in this environment.*




FIG. 6: *Mean time difference between the first robot has reached its goal and the second robot reached its goal.*

To analyze the difference in the times when the two robots reached their goals, we used ANOVA ($F(1, 5) = 70.50, \alpha = 0.05$) and a post hoc Tukey HSD ($df = 234, \alpha = 0.05$). The time difference for the random approach was significantly larger than for the other approaches ($p = 0.000020$). The reactive approach had a larger time difference than the planning strategy ($p = 0.000020$), the A-fixed ($p = 0.031$), the A-short ($p = 0.036$), and the A-long strategy ($p = 0.0063$). The planning strategy was significantly faster than the A-fixed ($p = 0.0021$), the A-short ($p = 0.0017$), and the A-fixed strategy ($p = 0.0012$).

3.3 Experiment 3

The environment in the last experiment contained randomly placed obstacles. This environment had a tracking error of 3.5% for the position and 8.4% for the angle. The increase of error, compared to the previous experiments, are due to more shapes and shadows from the randomly placed obstacles.

ANOVA ($F(1, 5) = 208.8, \alpha = 0.05$) indicated a significant difference between the strategies and a post hoc Tukey HSD test ($df = 234, \alpha = 0.05$) was used to find the effects of the different strategies. Again, the random strategy was significantly slower than all the other strategies ($p = 0.000020$). The reactive approach strategy was significantly slower than the planing strategy ($p = 0.000021$), the A-fixed ($p = 0.000020$), the A-short ($p = 0.000020$), and the A-long strategy ($p = 0, 000045$). We found no significant difference between the different anticipation approaches.

In the time difference between the robots we used ANOVA ($F(1, 5) = 30.43, \alpha = 0.05$) to find a significant difference between a strategy and post hoc Tukey HSD ($df = 234, \alpha = 0.05$) There was a significant difference between the random approach and then the other approaches ($p = 0.000020$). The reactive approach had a larger time difference than PLAN ($p = 0.003870$), A-fixed ($p = 0.009715$), and A-short ($p = 0.000020$). We found no significant difference between the different anticipation approaches.

4 Discussion

We have tested a number of behavioral strategies in robots in three simple environment with varying complexity to investigate the usefulness of anticipatory abilities in a real robot situation. As expected, a reactive approach behavior that always tries to move in the direction of the goal performs well in an empty environment but is much worse when there are obstacles in the way. Also, all strategies were much better than a random behavior.

In some cases, some of the anticipatory strategies were more efficient than the planning strategy in the sense that the robots would avoid taking routes where they may collide or interact, but the merit of anticipation clearly depended on how anticipation was used and in what environment. There were no significant differences between the time used for the planning and anticipatory strategies.

With anticipation, the robot will often take a longer path than with only planning and if something goes wrong during the avoidance of the anticipated obstacles, the robot will loose much time. Had the robots been more accurate when they attempt to follow their planned paths, we expect that anticipation would have been better in most cases as is the case in simulation [9]. It can clearly be seen that under optimal conditions, anticipatory behavior is very efficient, but it is very sensitive to different disturbances.

In experiment 2 and 3, the difference between the time when the first robot reached its goal and the time when the second did was correlated with the complexity of the navigation strategy. The reactive approach had a larger time

difference than the planning and anticipatory approaches, except in experiment 3 where there were no significant difference between REACT and A-long. A-long is a strategy that helps the robot that is furthest away from the goal. This could seem as a very noble approach, where the weakest robot is always allowed to take the shortest path. However, after a while, the weaker robot has advanced closer to the goal and thereby loses the benefit from the other robot. This switching can lead to strange behavior where the robot suddenly stops and turns in another direction because it has lost the rank and the other robot now possesses the area in front of the first robot.

The second experiment indicated that the planning behavior had a larger time difference than the anticipatory behaviors. In experiment 1 we could not find this effect. It is possible that without the obstacles, the effect was not as easy to distinguish even if it was still present.

In the experiments, the robots had total knowledge of the environment as well as the position and goal of the other robot. It only had to anticipate the movement of the other robot. In such a situation, it may be more advantageous to make a collective plan for all the robots. What we are aiming at in the future, however, is the situation where the robots do not have full access to the environment. In this case, the robots must explore the environment to learn about different paths and the positions of the other robot. As they will not know the goal of the other robots, they must be inferred from its observed movements. In this case, we will be able to explore different learning methods and different strategies for observing the behavior of the other robot.

In the future, we want to change the experimental set-up in three ways. First, we want to increase the exactness of the control system to allow more precise movement control of the robots. This will probably lead to a greater advantage for the anticipatory strategies. Second, we want to make the environments more complex. In this experiment, we used extremely simple environments to allow all the different behaviors to complete the task. This simpleness reduces the benefit of anticipatory behaviors. If we had instead used a more maze-like environment with lots of long and narrow passages, the anticipatory behaviors would have gained enormously with respect to the other strategies since a collision would be more costly. In the current environments, the cost of interaction between the robots was not high enough to punish behavior where the robots interfered with each other. Third, we will also use a larger number of robots which will also increase the benefits of anticipation.

In conclusion, we have presented experimental results with two robots in different environments that show that the ability to anticipate the behavior of the other robot has the potential to make the behavior of the robots more efficient. However, this is highly dependent on the complexity of the environment and the accuracy of the control of the robots.

5 Acknowledgement

We would like to thank Anna Balkenius for helpful comments on the manuscript and Henrik Johansson for the 3D drawing support. This work was supported by the EU project MindRaces, FP6-511931.

References

1. R. Arkin, "Motor schema-based mobile robot navigation." *International Journal of Robotics Research*, vol. 8, no. 4, pp. 92–112, August 1989.
2. C. Balkenius, *Natural Intelligence in Artificial Creatures*. Lund University Cognitive Studies, 1995, no. 37.
3. C. Balkenius and B. Johansson, "Event prediction and object motion estimation in the development of visual attention," in *Fifth International Conference on Epigenetic Robotics*, B. L., K. F., K. H., Y. H., K. J., M. G., N. J., S. G., S. G., and B. C., Eds., 2005, pp. 17–22.
4. S. Behnke, A. Egorova, A. Glove, R. Rojas, and M. Simon, "Predicting away robot control latency." in *RoboCup*, 2003, pp. 712–719.
5. C. Boutilier, "Planning, learning and coordination in multiagent decision processes," in *Sixth Conference on Theoretical Aspects of Rationality and Knowledge*, 1996, pp. 195–210.
6. R. A. Brooks, "Cambrian intelligence: the early history of the new ai," *MIT Press*, 1999.
7. R. Brooks, "A robust layered control system for a mobile robot," *IEEE Journal of Robotics and Automation*, vol. RA-2, no. 1, pp. 14–23, 1986.
8. R. A. Brooks, "Intelligence Without Reason," in *Proceedings of the 12th International Joint Conference on Artificial Intelligence (IJCAI-91)*, J. Myopoulos and R. Reiter, Eds. Sydney, Australia: Morgan Kaufmann publishers Inc.: San Mateo, CA, USA, 1991, pp. 569–595.
9. P. Davidsson, "Learning by linear anticipation in multi-agent systems," *Distributed Artificial Intelligence Meets Machine Learning*, vol. 1221, pp. 62–72, 1996.
10. M. Erdmann and T. Lozano-Perez, "On multiple moving objects," *Algorithmica*, vol. 2, pp. 477–521, April 1987.
11. Y. Guo and L. Parker, "A distributed and optimal motion planning approach for multiple mobile robots," in *Proceedings IEEE International Conference on Robotics and Automation*, 2002.
12. N. J. R. B. Hart, P. E. Nilsson, "A formal basis for the heuristic determination of minimum cost paths," *IEEE Transactions on Systems Science and Cybernetics SSC4*, vol. 2, pp. 100–107, 1968.
13. J. U. Korein and J. Ish-Shalom, "Robotics," *IBM Systems Journal*, vol. 1, pp. 96–106, 1987.
14. J. E. Laird, "It knows what you're going to do: Adding anticipation to a quakebot," in *Proceedings of the Fifth International Conference on Autonomous Agents*. Canada: ACM Press, May 2001, pp. 385–392.
15. S. A. B. D. Ledezma A., Aler R., "Predicting opponent actions by observation," *RoboCup 2004*, pp. 286–296, 2004.
16. F. Saad, "In-depth analysis of interactions between drivers and the road environment: contribution of on-board observation and subsequent verbal report," in *Proceedings of the 4th Workshop of international cooperation on theories and concepts in traffic safety*, University of Lund, 1992.

17. M. M. S. R. Sabanovic, S., "Robots in the wild: Observing human-robot social interaction outside the lab." in *International Workshop on Advanced Motion Control*, Istanbul Turkey, March 2006.
18. S. Sen, M. Sekaran, and J. Hale, "Learning to coordinate without sharing information," in *Proceedings of the Twelfth National Conference on Artificial Intelligence*, 1997, pp. 509–514, (Reprinted from *Proceedings of the National Conference on Artificial Intelligence*, 1994).
19. H. A. A. Sharifi, M. Mousavian, "Predicting the future state of the robocup simulation environment: heuristic and neural networks approaches," *Systems, Man and Cybernetics*, vol. 1, pp. 32–27, 2003.
20. M. B. M. Stulp, F. Isik, "Implicit coordination in robotic teams using learned prediction models," in *IEEE International Conference on Robotics and Automation*, 2006.
21. S. Thrun, J.-S. Gutmann, D. Fox, W. Burgard, and B. Kuipers, "Integrating topological and metric maps for mobile robot navigation: A statistical approach," in *AAAI/IAAI*, 1998, pp. 989–995.
22. M. Veloso, P. Stone, and M. Bowling, "Anticipation as a key for collaboration in a team of agents: A case study in robotic soccer," in *Proceedings of SPIE Sensor Fusion and Decentralized Control in Robotic Systems II*, P. S. Schenker and G. T. McKee, Eds., vol. 3839, Bellingham, September 1999, pp. 134–143.