AI Magazine Volume 21 Number 3 (2000) (© AAAI)

Using Reactive and Adaptive Behaviors to Play Soccer

Vincent Hugel, Patrick Bonnin, and Pierre Blazevic

■ This work deals with designing simple behaviors to allow quadruped robots to play soccer. The robots are fully autonomous; they cannot exchange messages between each other. They are equipped with a charge-coupled–device camera that allows them to detect objects in the scene. In addition to vision problems such as changing lighting conditions and color confusion, legged robots must cope with "bouncing images" because of successive legs hitting the ground. When defining task-driven strategies, the designer has to take into account the influences of the locomotion and vision systems on the behavior. Locomotion and vision skills should be made as reliable as possible. Because it is not always possible to simulate the problems encountered in real situations, the behavior strategy should anticipate them. In this article, we describe all the behaviors used to play soccer games on a soccer field surrounded with landmarks. Experiments were carried out at the 1999 RoboCup in Stockholm using the Sony quadruped robots (Fujita 2000).

Real quadruped robots cooperating or competing in collective entertainment games such as soccer (figure 1) can help a lot to deal with discrepancies between simulation and reality. In these cases, the actual behavioral strategy cannot only be based on behaviors emerging from simulation, where locomotion and vision skills are assumed to be always perfect. The strategy must also take into account the robots' potential in locomotion and vision (the visual sensor is the only exteroceptive sensor considered here). In fact, unlike wheeled robots, legged locomotion has a nonnegligible influence on the visual perception abilities of the machine. The recognition system must cope not only with changing lighting conditions and obstacle-strewn environments but also with "bouncing images." Because the robot never walks as well as in simulation and does not always perceive what it should in front of it, results can be quite different from what is expected.

As far as the behavioral strategy is concerned, the robot must be able to adopt some reactive behaviors and analyze the current situation so that it can adapt itself when it is changing.

Reactivity is absolutely necessary when the robot is facing unpredictable events, for example, if it loses the object it is tracking or if it falls down. Reactive behaviors are like reflexes: They make the machine react the same way when confronted with the same situation.

Adaptation is different from reactivity because it results from a kind of reasoning capability. It is also necessary for robots to be able to analyze a real situation so that they can make up their minds about what to do next. Here, the quality of the analysis determines the suitability of the behavior selected. For example, adaptation comes into play when reflexes lead to deadlocks. The machine should switch autonomously to a rescue behavior. Examples of adaptive behaviors can be slowing down to avoid obstacles or turning around a ball to bring it back to the desired location.

This work describes the entire strategy used to design behaviors of legged robots playing soccer. The article starts with the description of the influences of locomotion and vision on the behavior strategy. The strategy used to play soccer games on a colored field is then described in detail.

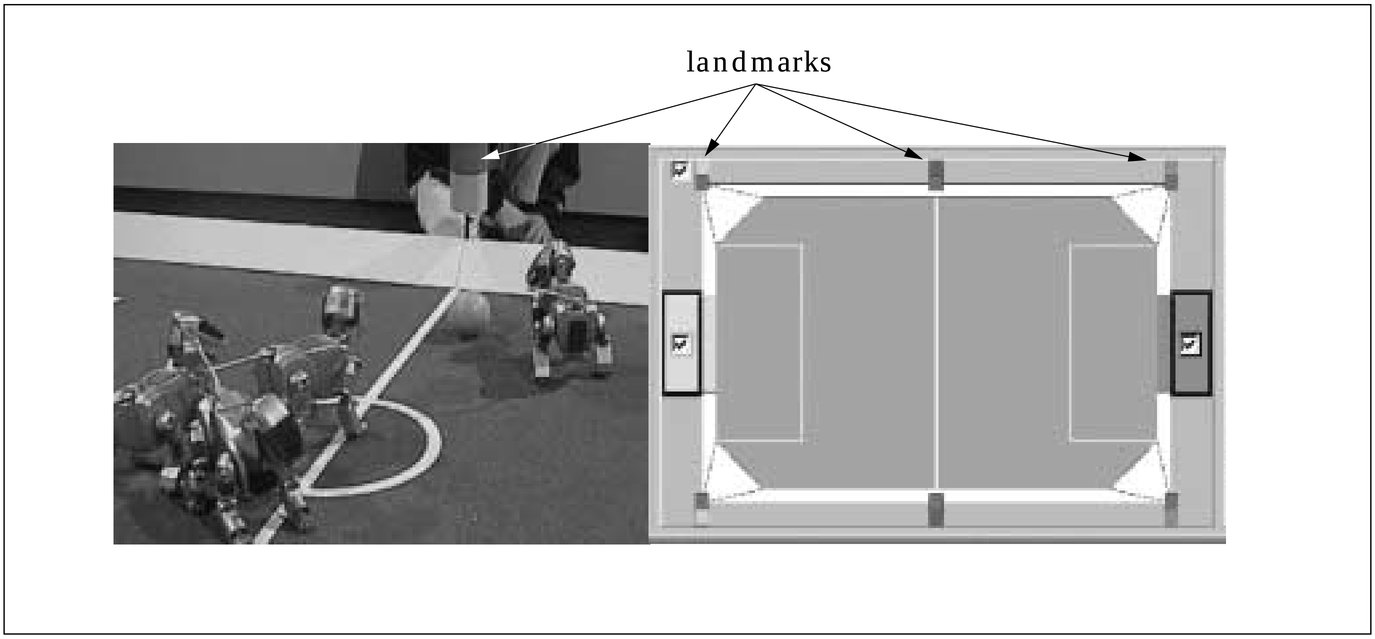

Figure 1. Sony Quadrupeds Playing Soccer and Top View of Soccer Field Surrounded by Six Landmarks (Poles).

Influence of Locomotion on Behavior

Locomotion is not omnidirectional. To change direction, the robot must compute the best-suited turning circle and switch to a turning pattern. If the turning circle is below the threshold, the machine should stop and switch to the turn-in-place mode in the desired direction (Hugel and Blazevic 1999). Therefore, if the supervision module needs to plan the trajectory of the body within a world reference frame, it must take these characteristics into account. Moreover, it must incorporate the possible delays for a new walking pattern to trigger. A supervision module was tested on simulation and gave satisfactory results.

The problem appears when trajectory planning is carried out on the real machine. In spite of the careful design of walking patterns, some leg slippage, drifts, and falls occur (Hornby et al. 2000; Hugel, Bonnin, and Blazevic 2000b; Hugel and Blazevic 1999). *Slippage* is the result of the nature of the interaction between the ground and the leg tip, which varies. Because the robot does not master these interactions, drifts cannot be avoided over a long distance (four or five times the length of the body). In addition, the robot might fall down when it bumps into walls or other objects in the scene. One reason for this behavior is that the locomotion is open-loop controlled and that the machine does not know how to deal with these situations. By adding some closed loops using exteroceptive data from force sensors, gyrometers, and accelerometers, the machine should be able to detect collision, loss of balance, and so on, and then react in real time to avoid falling down. In the worst case, where falls cannot be avoided, the machine should be aware of it and capable of getting back on its feet. Basic recovery behaviors have been implemented. Thanks to its three-axis accelerometer, the robot can detect on which side it has fallen and chose the right procedure to stand up.

When confronted with real situations, absolute trajectory planning does not provide good results because it relies on localizing procedures and, therefore, on the vision-recognition system. Experiments carried out on Sony quadrupeds show that absolute localization poses serious problems during motion (Lenser and Veloso 2000). Even when the robot is motionless, absolute localization sometimes yields poor results with bad accuracy, depending on the location on the soccer field. The strategy adopted by our team is to use relative positioning (Hugel, Bonnin, and Blazevic 2000a). Because precise positioning is not required, it does not matter if local motion is not accurate; the most important thing is to reach the goal. Therefore, the trajectory of the body is being corrected in real time by varying the turning circle according to the landmark captured while moving.

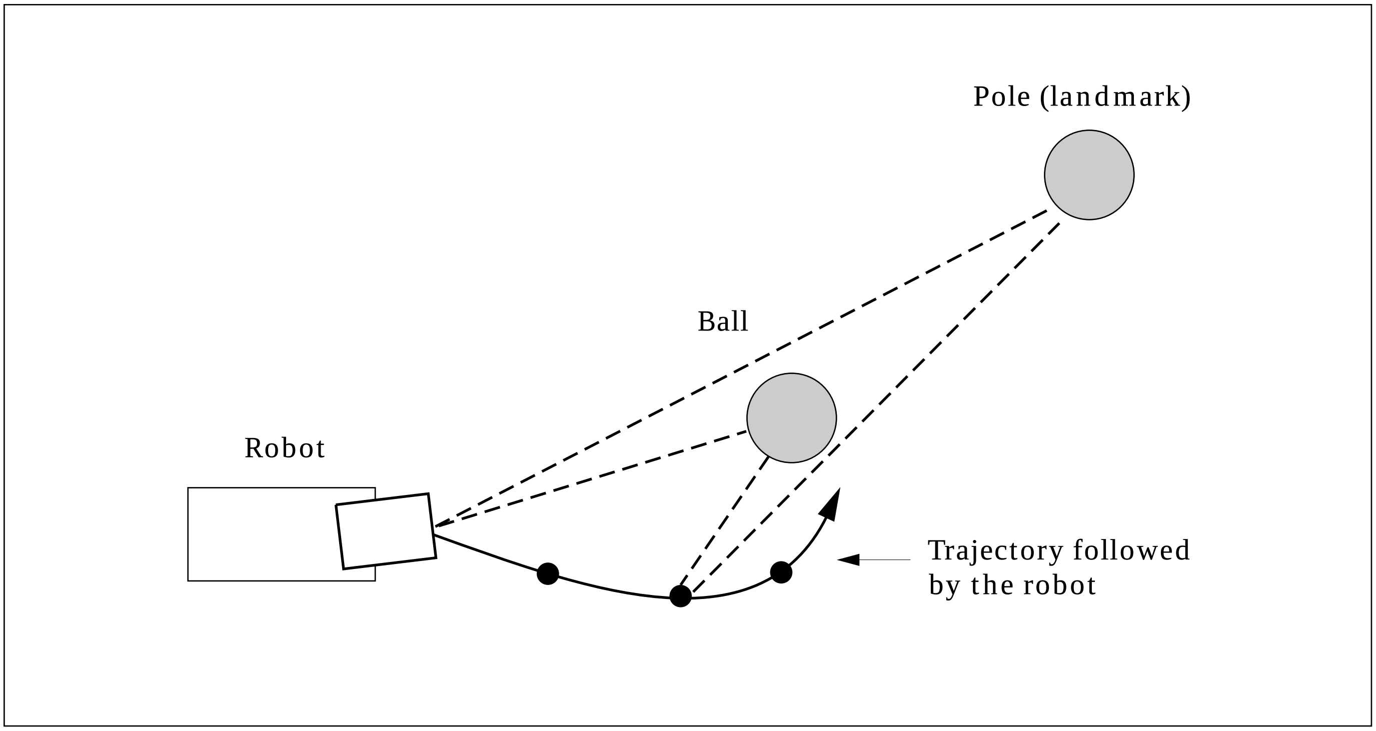

Figure 2. Online Trajectory Correction Using Landmarks Captured.

Influence of Vision on Behavior

The vision-recognition system is confronted with changing lighting conditions and bouncing images. The first difficulty can be dealt with using color-adaptive algorithms. Color thresholds can be changed online after analyzing the results on the field. However, the vision system cannot cope with the second difficulty. It often occurs that an object is detected in the first image acquisition and lost the time after.

Because it is not always possible to track an object in the successive images captured, it is essential to memorize the object in a certain amount of time. If the object is not refreshed when the timeout expires, the robot should switch to some different behavior. The first behavior consists of searching for the object lost. If it fails, and if the robot is completely lost, the next step is to begin an absolute localization procedure. Because it takes a lot of time, this behavior should only be used in critical situations. In other cases, relative positioning is preferable. For example, every time the robot can spot the ball, one goal or the other, or one of the landmarks surrounding the field, it can draw some information and adopt a more thoughtful behavior.

Two Kinds of Behavior

This section illustrates the strategies utilized to make quadruped robots play soccer successfully. Behaviors are separated into two groups. The first group includes basic behaviors that do not require an analysis of the situation. The second group contains more high-level behaviors that are based on interactions with the environment.

Basic Primitives

Two kinds of basic primitive are described here. The first kind regroups behaviors that can be seen as *reflexes*. Reflexes are always triggered when the robot is confronted with the same situation. The duration of such a behavior is limited. It is a kind of quick reaction to a similar event. Examples of reflexes are all the recovery procedures in case of a fall. Another reflex is where the robot turns its head toward the last position of the object lost. This reflex can help save searching time.

The second type of basic behavior is not limited in time. In fact, these behaviors need time to reach the objective set by the supervision module. Ball searching and ball tracking are examples of this kind of behavior. In ball searching, the goal consists of finding the ball again. In ball tracking, the objective is to get

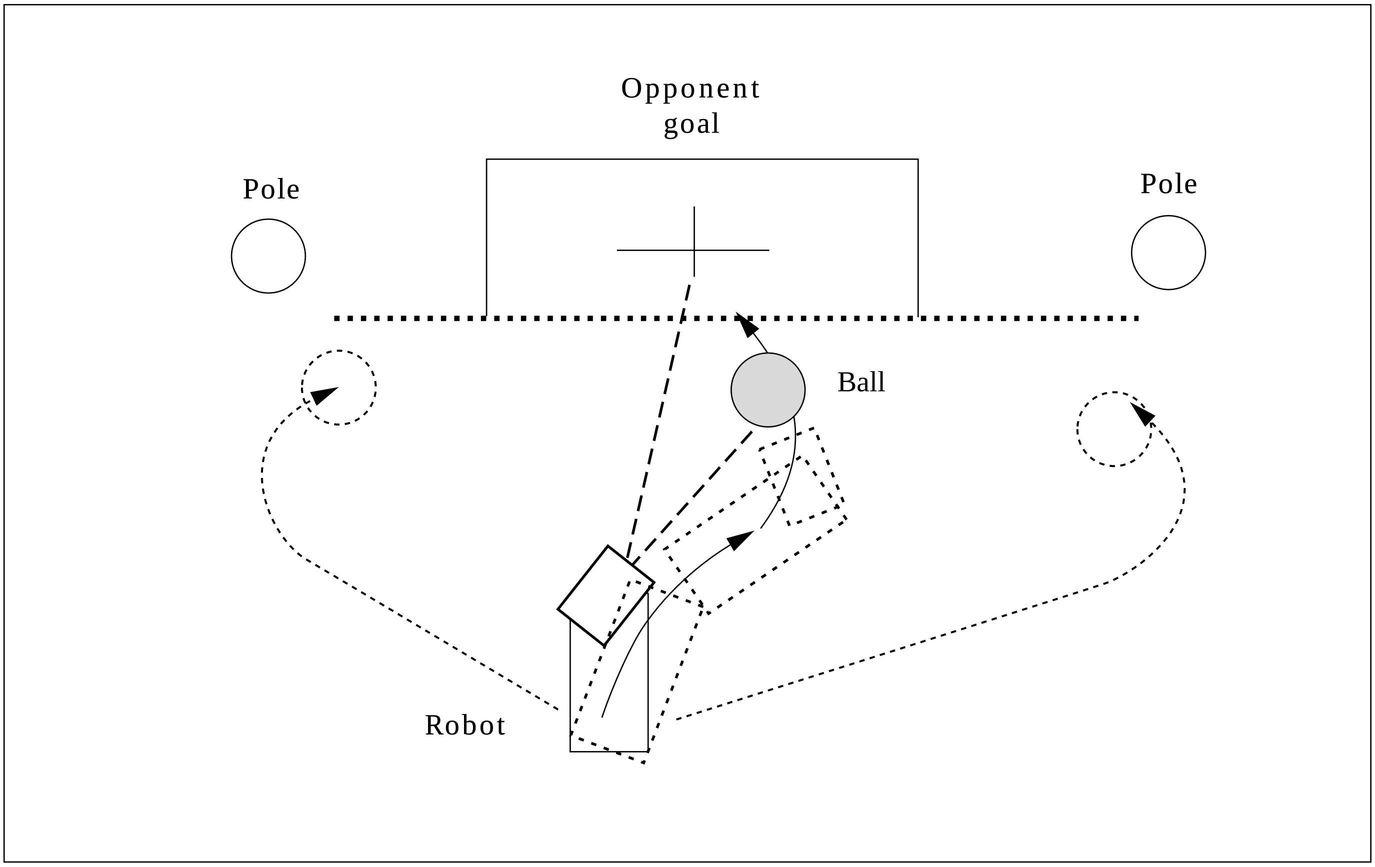

Figure 3. Attacker Behavior in Front of the Opponent Goal.
Direct attack of the opponent goal and attack from the corner.

closer to the ball. These behaviors are basic ones because no further analysis of the situation is required.

Higher-Level Behaviors Involving Situation Analysis

Four high-level behaviors have been designed for playing soccer.

The first high-level behavior consists of curving the trajectory according to the pole (landmark) captured (figure 2). The objective in the case of the soccer game consists of bringing the ball back to the opponent field. Every time a landmark is spotted, it is memorized, and the turning circle is corrected in such a way that the robot goes to the ball from the right or the left side, depending on which side of the field the pole is situated. The robot must have the ball in its field of view or have it memorized and have captured a landmark at the same time. The turning circle of the current trajectory applied depends on which pole is detected. The machine must therefore analyze the situation before triggering the behavior. In a soccer game, this behavior is used by a robot to defend its own goal or to attack the opponent's goal. In defense mode, if the robot can spot a corner pole of its own field and the ball at the same time, it can head for the ball and bring it back by moving behind it (*corner defense* behavior). In attacking mode, the robot can bring the ball back from the corners of the opponent side toward the opponent goal (*attack from the corner* behavior) (figure 3). Moreover, if the pole detected is a central one, the behavior selected is called *ball pushing into opponent field.*

The second behavior is used by players to attack the opponent goal. In this case, the robot tries to align the ball with the center of the adversary goal along its longitudinal axis. To trigger this behavior, the robot must spot the ball (the ball must be close enough) and the opponent goal. For this purpose, the robot uses the angles of the ball and the goal computed by the vision module. However, memorization should be tuned carefully. Refreshment rate must be sufficient; otherwise, the robot could miss the ball. Figure 3 shows the role of an attacker. This behavior is called *direct attack of the opponent goal.*

The third behavior consists of defending the robot's own goal when the robot spots the ball

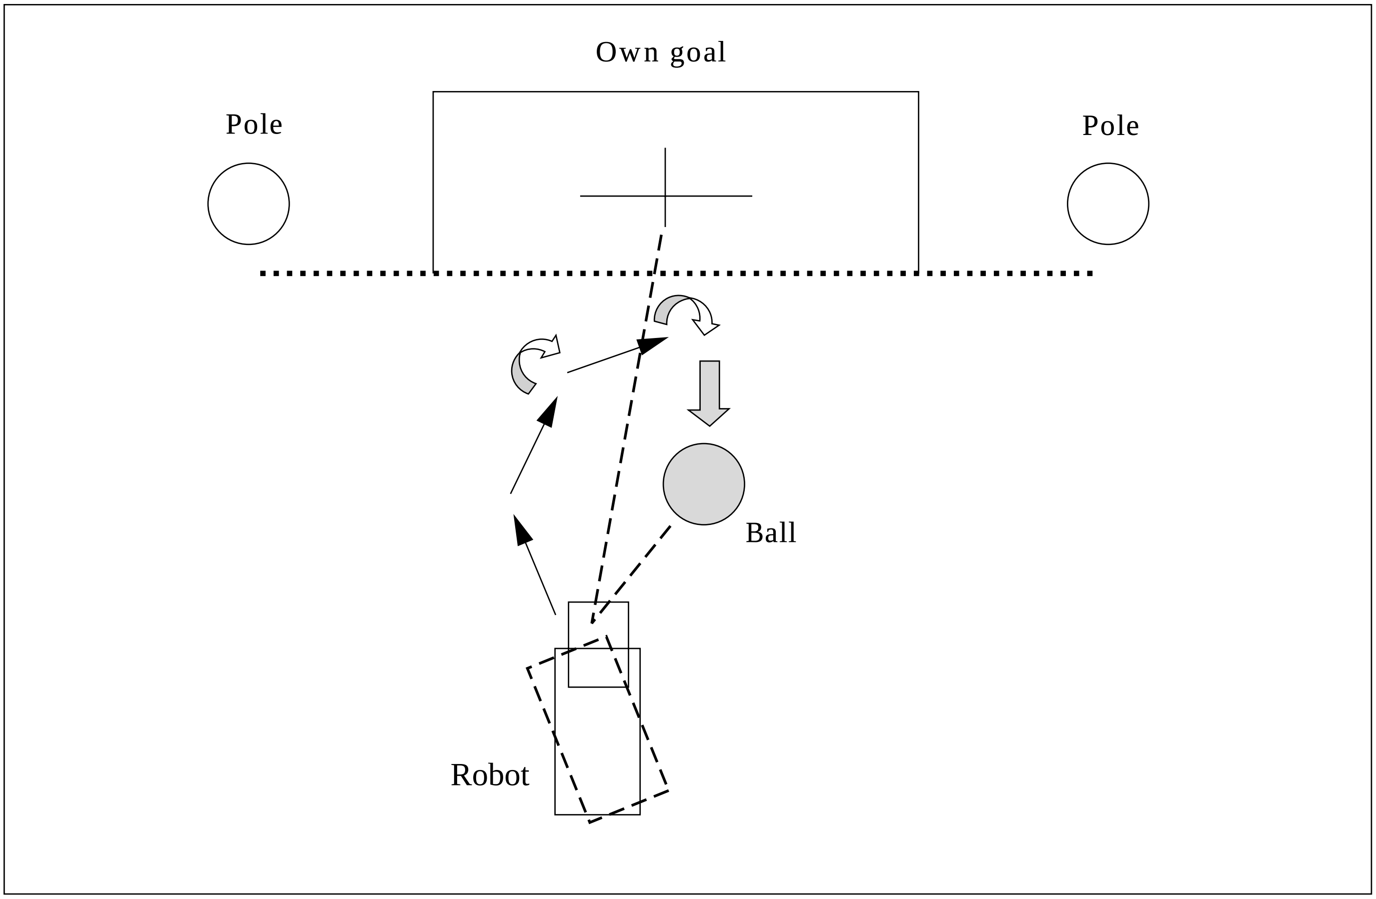

Figure 4. Turn-around-the-Ball Defense.

Successive turn-in-place and forward-motion sequences around the ball.

and its own goal at the same time, both of them being relatively close; this situation is dangerous. When all conditions are met, the robot starts to turn the ball around. The movement is a sequence of straight-line and turn-in-place motions, as shown in figure 4. During the procedure, the robot does not turn its head. If it loses the ball, the machine switches to turn-in-place motion until it sees the ball again, then it moves straight forward until losing the ball, and so on. Typically, the robot makes a turn of approximately 180° and can therefore push the ball to the opponent field. This behavior is called *turn-around-the-ball defense.*

The last behavior is specially designed for the goalkeeper. First, the robot directs its head toward the ball. If the ball gets close, the goalie aligns its body with the ball. Once the ball gets closer, the robot moves toward the ball to strike it and push it away. It then searches for its goal and goes to it. Its infrared sensor allows it to detect the walls of its goal; it then begins to search for the opponent goal by turning in place. Once it has found the opponent goal, it tries to spot the ball again. The procedure can restart. In this behavior, the robot has to estimate the distances to the ball and the goal. Experiments must be carried out to tune them because they depend on the lighting conditions.

Behavior Interaction

In the case of soccer, the interaction between the different behaviors is easily managed. In fact, some behaviors are naturally preemptive. The direct attack of the goal has a higher priority than other behaviors. If it were the behavior that curves the trajectory, the robot could miss the ball and the goal. In the case of defense, it is the turn-around-the-ball behavior that preempts all the others because this situation is the most critical. It can be seen as a rescue or last-chance behavior to prevent the other team from scoring. Figure 5 summarizes the behavior priorities. Figure 6 features the events that trigger the behaviors described in Two Kinds of Behavior.

The experience brought by soccer experi-

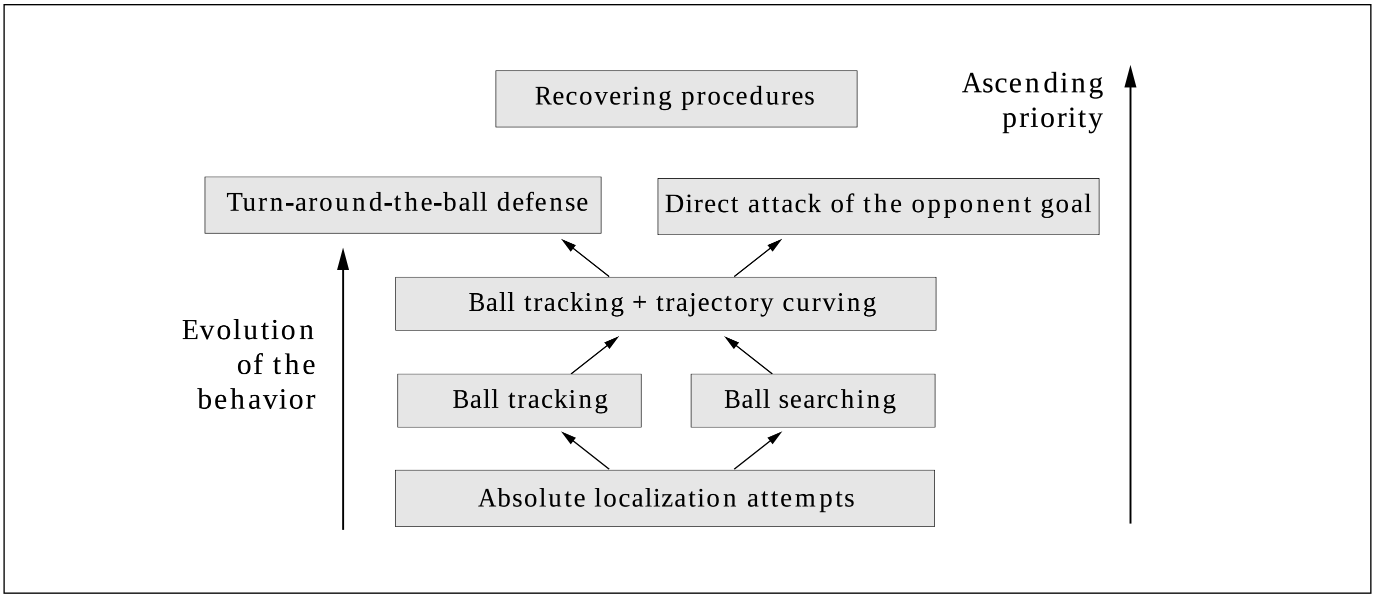
Recovering procedures
Ascending priority
Turn-around-the-ball defense
Direct attack of the opponent goal
Ball tracking + trajectory curving
Evolution of the behavior
Ball tracking
Ball searching
Absolute localization attempts

Figure 5. Scale of Behaviors' Priorities.

Example of soccer game.

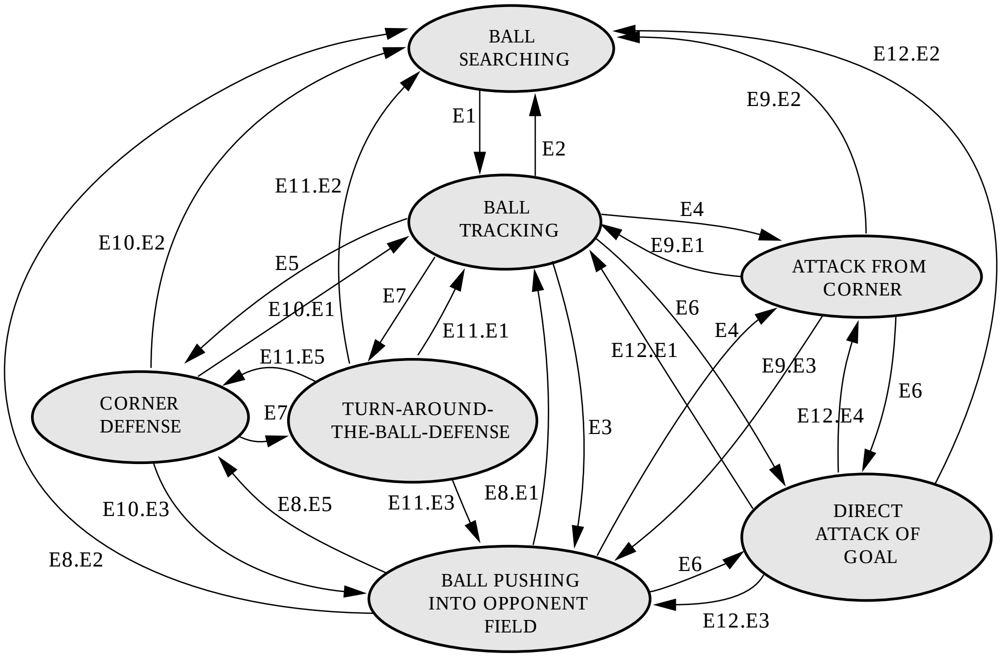

Events:

E1. Ball captured
E2. Ball lost
E3. Ball and central pole spotted
E4. Ball and corner pole of opponent field spotted
E5. Ball and corner pole of partner field spotted
E6. Ball and opponent goal close
E7. Ball and partner goal close
E8. Time-out after detecting central pole
E9. Time-out after spotting corner pole of opponent side
E10. Time-out after spotting corner pole of partner side
E11. Turn-around-the-ball defense achieved
E12. Time-out after detecting opponent goal close

Priorities('<' means 'is of lower priority than')
E1<E3<E5<E7
E1<E3<E4<E6

Figure 6. Behavior Interactions.

Goalkeeper and recovery behaviors not represented.

ments has taught us that behaviors should not be limited to classic basic ones, such as turn to the right, turn to the left, go to a specific location, stop to localize, and achieve a fixed number of steps. It is better to design some adaptive behaviors that interact with the environment. Decisions should be made while moving, and relative positioning should be used most of the time.

Conclusion and Perspectives

The experiments carried out on quadruped robots playing soccer allowed us to test behaviors in real situations. It helps us become aware of real implementation problems. Vision and locomotion skills should be taken into account in the design of behaviors. In the first series of experiments, communication between robots was not possible. However, the next version of quadruped prototypes should be provided with wireless communication links. Behaviors could then be improved because robots would know the relative or absolute positions of others. Information could then be crossed to increase the accuracy of positioning.

Another way to improve the overall behavior is to use *gyrodometry.* This technique consists of locally using gyrometer data to correct drifts from odometry. The robot can then estimate its displacement with enough precision over a few more steps.

In conclusion, a strategy that combines reactive and adaptive behaviors should give satisfactory results, provided that the interaction between both is well managed.

Acknowledgments

We are very grateful to Sony Corporation and in particular to the Digital Creature Laboratory in Tokyo, Japan, for lending us four prototypes of the Sony pet robot called AIBO and helping us all along the project of the RoboCup challenge.

References

Fujita, M. 2000. Digital Creatures for Future Entertainment Robotics. In Proceedings of the 2000 IEEE International Conference on Robotics and Automation, 801–806. Washington, D.C.: IEEE Computer Society.

Hornby, G. S.; Takamura, S.; Yokono, J.; Hanagata, O.; Yamamoto, T.; and Fujita, M. 2000. Evolving Robust Gaits with AIBO. In Proceedings of the 2000 IEEE International Conference on Robotics and Automation, 801–806. Washington, D.C.: IEEE Computer Society.

Hugel, V., and Blazevic, P. 1999. Toward Efficient Implementation of Quadruped Gaits with Duty Factor of 0.75. In Proceedings of the 1999 IEEE International Conference on Robotics and Automation, 2360–2365. Washington, D.C.: IEEE Computer Society.

Hugel, V.; Bonnin, P.; and Blazevic, P. 2000a. Reactive and Adaptive Control Architecture Designed for the Sony Legged Robot League in RoboCup 1999. In Proceedings of the 2000 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2000). Washington, D.C.: IEEE Computer Society. Forthcoming.

Hugel, V.; Bonnin, P.; and Blazevic, P. 2000b. Speed Control for Quadruped Using Duty Factor. In Proceedings of the 2000 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2000). Washington, D.C.: IEEE Computer Society. Forthcoming.

Lenser, S., and Veloso, M. 2000. Sensor Resetting Localization for Poorly Modeled Mobile Robots. In Proceedings of the 2000 IEEE International Conference on Robotics and Automation, 1225–1232. Washington, D.C.: IEEE Computer Society.

Vincent Hugel graduated from the Ecole des Mines de Nancy in 1994. He received his specialization degree in robotics at INSTN (CEA, FRANCE) in 1995. He has been involved in the design and implementation of walking algorithms on hexapods and quadrupeds. This work was part of his work toward his Ph.D., which he received in November 1999. His e-mail address is hugel@robot.uvsq.fr.

Patrick Bonnin graduated from ENS de CACHAN, France, and received his aggregation of applied physics in 1985. He received his master degree in 1986 in remote sensing from Paris VII University and his Ph.D. in image processing in 1991 from Paris VII University. He has been an assistant professor since 1992 at Paris Nord University - IUT de Villetaneuse. From 1985 to 1996, he worked at ETCA (French Defense Laboratory) in image processing, data fusion, and parallel processing. He joined the LRP in 1998. His current research interests are a real-time vision system for robotics, segmentation of color images, and stereoscopic vision. His e-mail address is bonnin@robot.uvsq.fr.

Pierre Blazevic graduated from ENS de CACHAN, France. He received his aggregation of applied physics in 1985, a master degree in robotics in 1986, and a Ph.D. from Paris VI University in 1991. He has been an assistant professor since 1992 at Versailles University. He joined the LRP at its creation in 1987. His current research interests are wheeled and legged locomotion and force-feedback devices for teleoperation. His e-mail address is blazevic@robot.uvsq.fr.

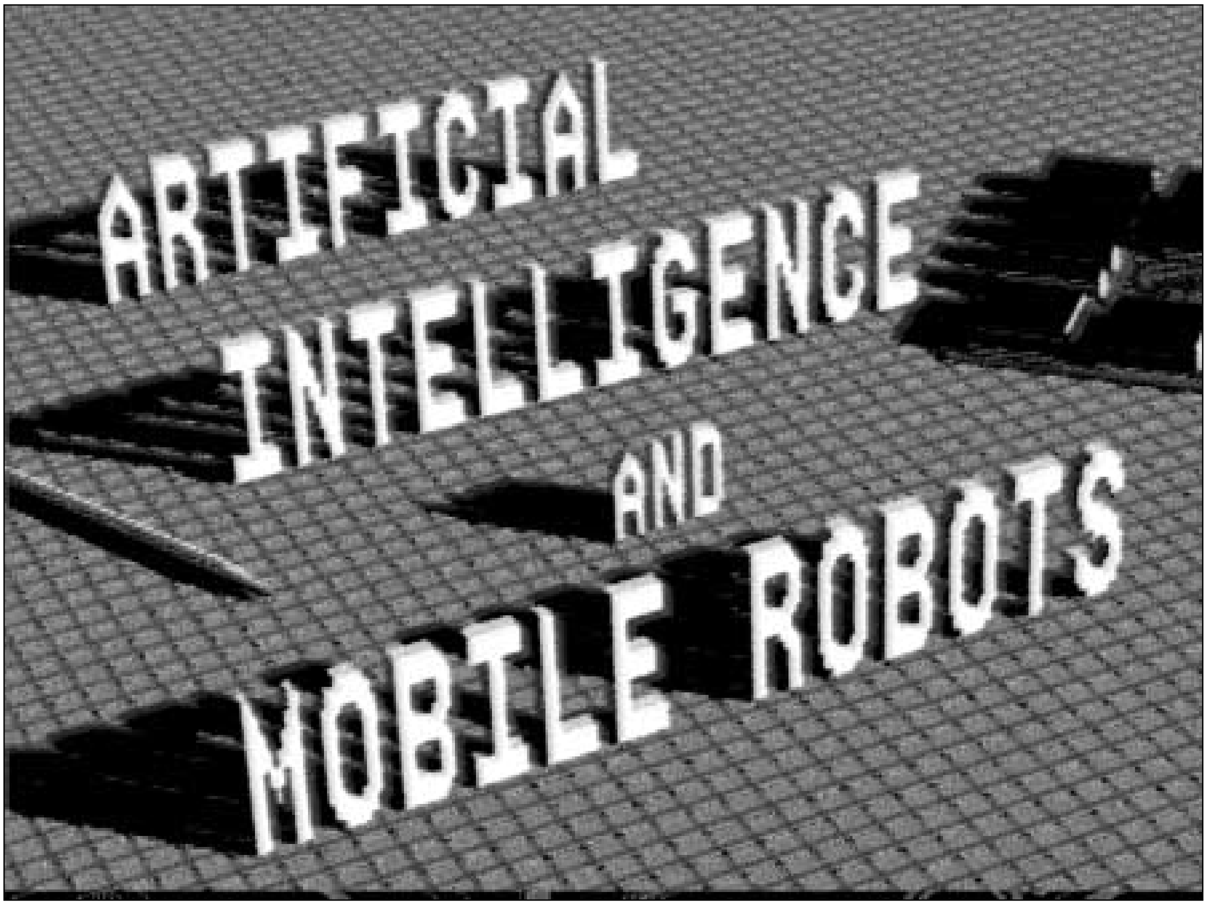

ARTIFICIAL
INTELLIGENCE
AND
MOBILE ROBOTS