

Playing Soccer with Legged Robots

Manuela Veloso[†], William Uther[†], Masahiro Fujita[‡], Minoru Asada[§], and Hiroaki Kitano^{§§}

[†]Computer Science Department, Carnegie Mellon University, Pittsburgh PA 15213, USA

[‡]D21 Laboratory, Sony Corporation, 6-7-35, Kitashinagawa, Shinagawa-ku, Tokyo, 141 Japan

[§]Dept. of Adaptive Machine Systems, Osaka University, Suita, Osaka, 565 Japan

^{§§}Sony Computer Science Laboratory Inc., 3-14-13 Higashi-Gotanda, Shinagawa-ku, Tokyo, 141 Japan

Abstract

Sony has provided a remarkable platform for research and development in robotic agents, namely fully autonomous legged robots. In this paper, we describe our work using Sony's legged robots to participate at the RoboCup'98 legged robot demonstration and competition. Robotic soccer represents a very challenging environment for research into systems with multiple robots that need to achieve concrete objectives, particularly in the presence of an adversary. Furthermore RoboCup'98 offers an excellent opportunity for robot entertainment. We introduce the RoboCup context and briefly present Sony's legged robot. We developed a vision-based navigation and a Bayesian localization algorithm. Team strategy is achieved through pre-defined behaviors and learning by instruction.

1 Introduction

Problem solving in complex domains necessarily involves multiple agents, dynamic environments, and the need for learning from feedback and previous experience. Robotic soccer is an example of one such complex task where agents need to collaborate in an adversarial environment to achieve specific objectives.

Research in robotic soccer has been pursued along several different aspects of the problem: simulation, small-sized, medium-sized robots, and legged robots.¹

In this paper we present some of the initial work developed with the legged robots built by Sony at Carnegie Mellon University and Osaka University. The University of Paris VI is also participating in the RoboCup'98 games and demonstrations of the quadruped robots. Legged robots represent a remarkable advancement for robotics. In the particular con-

text of robotic soccer, legged robots provide a very interesting opportunity for robot entertainment.

2 The Concrete RoboCup-98 Setup

The legged robot as a robotic soccer player is a fully autonomous robotic system without global vision or wireless remote operation. In addition, in order to simplify the RoboCup competition, no modification of hardware by the three different competing teams was allowed. The Legged Robot Exhibition Match is therefore a software competition between robots with the same hardware platforms. Figure 1 shows the field for the Legged Robot Exhibition Match.




Figure 1: The RoboCup'98 Soccer Field for Legged Robot Competition.

The setup of the game includes the following characteristics:

Colored Environment: The important game items, e.g., the ball and the goals, are painted with different colors so that a stand-alone robot can process vision tasks in real time. We use a set of eight colors distributed in the UV color space. We selected the colors carefully so that a robot can easily distinguish the colored objects.

¹See [9] and <http://www.robocup.org/RoboCup/RoboCup.html> for additional information.

Team: Each team has three players. This is the minimum number for each team to do some “team play” or collaboration tasks, which is one of the research objectives of RoboCup.

Field Size: The field size ($2m \times 3m$), should provide enough space for six robots to navigate while dribbling and passing the ball.

Slanted Wall: The walls surrounding the field with 45° slanted. In addition, we also made a triangular slanted wall for every corner. These slanted walls effectively return the ball to the field when the ball is pushed against the border, which is a frequent situation at least in our practice with remote-controlled legged robots.

Landmarks: Six poles are used as landmarks for self-location. Each landmark is painted with two different colors so that six different poles can be painted with only three different colors. The robot should measure the angles of three different poles to get the self-location in the field.

In addition to the regular soccer game, we will evaluate software performance based on the RoboCup Physical Challenge [2, 8], which is defined by the RoboCup Challenge Committees.

3 A Quadruped Legged Robot

The robot soccer player used in RoboCup Legged Robot Exhibition Match is based on OPEN-R [5], which is proposed by Sony as a standard architecture for Robot Entertainment Systems [6]. The significant feature of OPEN-R is a decomposition technology for both of hardware and software modules. This technology enables us to build various kinds of robot styles such as a quadruped robot and wheel-based robot style, as well as various software configurations [7]. In addition, for software researchers, OPEN-R can provide a highly reliable robot platform so that they can concentrate on software development for a new image processing algorithm, posture control, agent architecture, and son on. Furthermore, they don't need to develop software from scratch, because OPEN-R provides some software modules such as color detection and walking control. This means that OPEN-R can accumulate the developed software as reusable components, and accelerate autonomous robot research.

In addition, hardware researchers can design their own hardware modules with an OPEN-R interface,

which can then attach to an existing OPEN-R system. In RoboCup-98 Paris, we prohibit modification of robot hardware, however, in principle, it is possible for us to build various styles of robots with various sensors and actuators which have OPEN-R interface.

For RoboCup-98 Paris, we will deploy legged robot with four legs and one head, each of which has three degree of freedom, and rich sensory channels, for example, a head has a color CCD camera, stereo microphone, touch sensors, and a loud speaker.

Most intelligent autonomous robots are implemented with a wheel-based mechanical configuration. A wheel-based robot has an advantage in simplicity of motion control, so that researchers can concentrate on vision, planning, and other high-level issues. However, since our goal is robot entertainment, we have a different emphasis. We believe that the capability of representation and communication using gesture and motion is very important in entertainment applications. Therefore, we chose a mechanical configuration of our robot as a quadruped-legged type, as shown in Figure 2.




Figure 2: Legged Robot

The merits of the quadruped-legged configuration are: (1) walking control of a quadruped is easier than that of a biped robot, and (2) when in a sitting posture, two hands are free to move, therefore, they can be used to present emotions or to communicate with a human.

Since each leg or hand has to be used for various purposes besides walking, we assign three degree of freedom (DoF) for each leg/hand. In addition, we add a tail and three DoF for neck/head so that the robot has enough representation and communication capabilities using motions. During the RoboCup games, legs are not necessary used for expressing emotions. However, they can be used for sophisticated control of balls, such as passing ball to the side or back, or engaged in deceptive motions.

One possible disadvantage of using a legged robot is that their moving speed is not as fast as wheel-based robots. In the future, the speed issue may be resolved when galloping is made possible. For now, legged robot will be played within dedicated league. Although serious hardware limitations exist, teams with efficient leg motion coordination may have major advantages in the game.

In general, it is difficult for a stand-alone (autonomous) robot to perform navigation and perception tasks in real time in a real world environment because of its limited computational power. Remotely controlled robots and systems with globally overlooking cameras can provide off-board the computational power. These solutions have disadvantages, such as the need for having a video transmitter for each robot in a remotely controlled system, and the lack of individual views with the global vision.

We believe that technologies to process images from each individual robot viewpoint without any global information will become very important in Robot Entertainment in future. Therefore, we decided to build the RoboCup System with stand-alone robots under local communication constraint. We successfully resolved two hardware issues to enable full on-board vision system for small size robots: (1) small camera size, and (2) large on-board processor power. We solved these problem by actually manufacturing a dedicated camera and a processor chip [8].

4 Vision-Based Navigation

Each legged robot for RoboCup'98 is equipped with a single perception sensor, namely a vision camera. The hardware-based vision processor provides a robust eight-color discrimination. Robots need to act solely in response to the visual input perceived. At Carnegie Mellon, we have therefore decomposed our work along the following aspects:

- Reliable detection of all of the relevant colors: orange (ball), light blue (goal and marker), yel-

low (goal and marker), pink (marker), light green (marker), dark blue (teammate/opponent), and dark red (opponent/teammate).

- Active ball chasing: the robot actively interleaves searching for the ball and localization on the field to evaluate both an appropriate path to the ball and final positioning next to the ball.
- Game-playing behaviors: robots play attacking and goal keeping positions.

In this section, we present our on-going research in addressing these issues.

4.1 Supervised Learning of UV Colors

The Sony legged robot has specialized hardware for the detection of colors. However, this hardware still requires pre-setting of appropriate thresholds in YUV color space for the desired colors. It is well known that color adjustments are highly sensitive to a variety of factors, such as lighting and shading.

Given that the legged robots inevitably act under many different conditions, at Carnegie Mellon we developed a method to automatically acquire the necessary YUV color thresholds. First we developed a tool to manually experiment with different boundaries in the UV space.² Secondly, and building upon the experience using the YUV-space tool, we developed a classification algorithm to automatically learn the thresholds that maximize the accuracy of the desired color detections. Our algorithm relies on supervised classification using a set of training and testing images. By moving the robot to different positions on the field, we accumulate a series of images. For each image, we manually classify the regions of the different colors. This manual labeling is easily done through an interface that we developed that overlays the original image and the supervised classification. Areas of the image can have their correct classification specified using tools similar to PC type paint programs.

Once the data has been classified, the YUV color thresholds are learned separately for each color using a conjugate gradient descent based algorithm. Each threshold is softened by replacing it by a sigmoid function:

$$C = \frac{1}{1 + e^{t(a-x)}}, \quad (1)$$

where C is the classification for this sigmoid, a is the value of this threshold and t is a variable equivalent to the current temperature in a simulated annealing

²This tool was developed by Kwun Han, for which we thank him.

algorithm. The classifications of the different thresholds are multiplied and gradient descent is performed on the sum-squared error. Initially the temperature is quite high. It is reduced gradually over time as the sigmoids are learned.

In our experiments, we generally use about twenty images for training. The learning algorithm converges in less than one hour achieving a high classification accuracy.

4.2 Bayesian Probabilistic Localization

The CMU legged Robot team uses Markovian Localization to determine the robot position on the field.

Relying on dead-reckoning in the legged robot for localization is completely unrealistic. The effects of movement actions are highly noisy and modeling them accurately did not seem feasible to us. However, to compensate to the highly unreliable dead reckoning, the field environment for RoboCup'98 includes several fixed colored landmarks. Therefore, at Carnegie Mellon, we developed a Bayesian localization procedure (e.g., [4, 3]). As the robot cannot keep enough markers in view at all times to calculate its location directly, our algorithm uses a probabilistic method of localization using triangulation based on two landmarks.

The field is discretized in grid locations. The continuous robot head angles are also discretized. We create a state space with these discretized grid cells and robot headings. Observations of the landmarks are combined with the state space for the position calculation. There are two passes to our localization algorithm. The first incorporates observations into our probability distribution. The second takes into account the movement actions selected. Incorporation of observations is based upon Bayes' Rule:

$$P(S_i|O) = \frac{P(S_i)P(O|S_i)}{\sum_j P(S_j)P(O|S_j)} \quad (2)$$

where $P(S_i)$ is the apriori probability that the robot is in state S_i , $P(S_i|O)$ is the posterior probability that the robot is in state S_i given that it has just seen observation O and $P(O|S_i)$ is the probability of observing O in state S_i .

To represent the probability distribution $P(S_i)$ we use a table of values. The table is a three dimensional mapping $X \times Y \times \theta \rightarrow S$. A table of values was chosen because of some of the distributions we wish to represent do not have a nice parametric form. For instance, given a uniform prior distribution, the observation of the angle between two markers gives a high probability

circle through the state space that is not representable by a Gaussian distribution.

Incorporation of movement is based upon a transition probability matrix. Given a previous movement M , for each state, the algorithm computes the probability that the robot be in that state:

$$P(S_i|M) = \sum_j P(S_j)P(S_j \Rightarrow S_i|M) \quad (3)$$

where $P(S_j)$ is the apriori probability of state S_j and $P(S_j \Rightarrow S_i|M)$ is the probability of moving from state S_j to state S_i given the movement M . It is assumed that the transition probabilities, $P(S_j \Rightarrow S_i|M)$, take into account any noise in M .

For example, imagine the robot sees an angle of 90° between two markers, turns to the left and then sees an angle of 90° between two more markers. Initially it does not know where it is - our prior distribution is flat. After its first observation, the projection of the state probability matrix onto the X, Y plane would be as shown in Figure 3(a). During the turn, the projection spreads over the X, Y plane - representing the increased uncertainty introduced by the dead reckoning as the robot turns (see Figure 3(b)). The second observation finally localizes the robot (see Figure 3(c)).




Figure 3: The positioning probability: (a) after the 1st observation; (b) after a 90° turn; (c) after the 2nd observation.

Our localization algorithm is invoked actively, i.e., the robot searches for two landmarks when the maximum state probability is below a pre-defined threshold. In our experiments, the robots localize themselves with high accuracy.

4.3 Role-Based Behaviors

Following up on our experience with the small-size RoboCup wheeled robots [11, 12], we developed different behaviors based on positioning on the field. As of now, robots play two different roles, namely attacking and goal keeping.

The procedure used by an attacking robot consists of the following steps: (i) find the ball; (ii) localize attacking goal; (iii) position behind the ball, aligned with the goal; (iv) shoot or pass.

The procedure used by the goal keeping robot is a simplified version of the CMUnited'97's goal keeper and consists of the following steps: (i) find the ball; (ii) remain close to the goal; (iii) move sideways aligned with ball; (iv) clear the ball when it gets close to it.

5 Behavior Acquisition by Teaching

The final goal of the Legged Robot Project in the Osaka University is to establish the methodology to acquire behaviors for team cooperation in the RoboCup context from the interactions between the legged robots through multi sensor-motor coordinations. The desired behaviors can be categorized into three levels: a basic level, a basic cooperation level, and a higher team cooperation level. In this section, we briefly explain our first step for the first level skill acquisition with preliminary results. That is behavior acquisition by direct teaching.

The most fundamental feature of the legged robot is that they move by their four legs (12 DOFs), which is quite different from conventional mobile robots (2 or 3 DOFs). From a viewpoint of sensor-motor learning and development, multi sensory information and multi DOFs control should be established simultaneously, that is, affecting each other, the sensory information is abstracted and the multi joint motions are well coordinated at the same time [1]. However, it seems very difficult for artificial systems to develop both together. Our goal is to design such a method.

From our experiences on robot learning, we realized that the number of trials by real robot is limited and a good trade-off between computer simulations and real robot experiences is essential for good performance. However, the computer simulation of the legged robots seems difficult to build, then we decided to adopt a direct teaching method in order to reduce the number of trials by real robot.

First, we collect test data, pairs of an action command given by human trainer and the sensory information during the action execution. Next, we apply

C4.5 [10] to the test data to extract rule sets. Then, the validity of the rule sets are checked against test data by applying the rule sets. Specifications of the data for shooting skill acquisition are as follows:

- **action command** forward, backward, left-shift, right-shift, left-rotation, and right rotation (20 degs/sec): these abstracted action commands will be decomposed into more primitive motor commands in future.
- **sensory information** head direction (rad.) and image features of both the ball and the goal in the observed 88×60 image: area (pixels), position (x, y coordinates), bounding rectangle (x, y coordinates of corners), height, and width. (See Figure 4.)
- **training position and sampling rate:** initial positions of direct teaching by serial line connection are evenly distributed in the field heading the goal, and the sampling rate is 300ms. One trial from the initial position to the goal takes about 10 seconds.

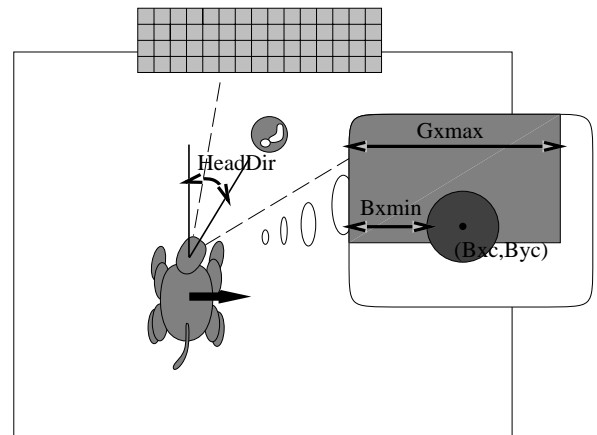


Figure 4: Typical situation of Right-Shift motion

We have collected about 740 pairs of an action command and sensory information for training data to make rule sets, and 500 pairs for test data to check the validity of the rule sets. Both sets of data are obtained in the same manner starting from the similar initial positions, but individual pairs are different from each other trial by trial.

The number of rules obtained is about 30, and typical ones for forward(F), left-rotation (LR), and right-shift (RS) are the following:

Forward:	Left-Rotation:	Right-Shift:
BallArea>56	BallArea>49	BallArea<=40
HeadDir>-.14	HeadDir>.52	Ballycen>8
HeadDir<=.37	HeadDir<=1.10	BallXmin>3
GoalXmin<=11	BallWidth<=11	HeadDir<=-.24
		GoalArea<=665
		GoalXmax<=64

Figure 4 shows the typical situation of right-shift motion.

Due to the inaccurate teaching by the human trainer (actually, the human trainer is not experienced yet to do this sort of teaching), the learning algorithm is not completely reliable yet. The following table indicates a confusion matrix showing where some miss-classification of the training cases occurs:

(a)	(b)	(c)	(d)	(e)	(f)	<-classified as
111	13	3	--	4	--	(a): LEFT
14	--	--	2	9	--	(b): LEFT_TURN
--	--	8	16	9	--	(c): RIGHT_TURN
--	--	20	151	34	--	(d): RIGHT
1	--	3	13	91	--	(e): FORWARD
--	--	--	--	--	--	(f): BACK

Note in that several miss-classifications are shown (in particular LEFT_TURN does not have any correct classifications). We are planning to skill up teaching and also to use more training data to construct robust classification. Generalization is one of the big issues of direct teaching, and explanation-based learning seems one alternative to solve the problem. These generalization techniques are under the investigation.

6 Conclusion

In this paper, we reported on some of our work using the Sony quadruped legged robots to play robotic soccer. We briefly described the components of Sony's legged robots. We then presented our vision-based navigation, Bayesian localization, role-based behaviors, and behavior acquisition by teaching.

References

[1] M. Asada. An agent and an environment: A view on having bodies - A case study on behavior learning for vision-based mobile robot. In *Proceedings of 1996 IROS Workshop on Towards Real Autonomy*, pages 19-24, 1996.

[2] Minoru Asada, Yasuo Kuniyoshi, Alexis Drosgoul, Hajime Asama, Maja Mataric, Dominique Duhaut, Peter Stone, and Hiroaki Kitano. The

RoboCup physical agent challenge: Phase-i. *Applied Artificial Intelligence*, 12, 1998.

- [3] W. Burgard, A.B. Cremers, D. Fox, D. Haehnel, G. Lakemeyer, D. Schulz, W. Steiner, and S. Thrun. The interactive museum tour-guide robot. In *Proceedings of AAAI-98*, Madison, WI, July 1998.
- [4] Alberto Elfes. *Occupancy Grids: A Probabilistic Framework for Robot Perception and Navigation*. PhD thesis, Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA, 1989.
- [5] M. Fujita and K. Kageyama. An Open Architecture for Robot Entertainment. In *Proceedings of the First International Conference on Autonomous Agents*, 1997.
- [6] M. Fujita and H. Kitano. Development of an Autonomous Quadruped Robot for Robot Entertainment. *Autonomous Robots*, 5, 1998.
- [7] M. Fujita, H. Kitano, and K. Kageyama. Reconfigurable Physical Agent. In *Proceedings of The Second International Conference on Autonomous Agents*, 1998.
- [8] H. Kitano, M. Fujita, S. Zrehen, and K. Kageyama. Sony legged robot for robocup challenge. In *Proceedings of the 1998 IEEE International Conference on Robotics and Automation*, pages 2605-2612, 1998.
- [9] H. Kitano, M. Tambe, P. Stone, M. Veloso, S. Coradeschi, E. Osawa, H. Matsubara, I. Noda, and M. Asada. The robocup synthetic agent challenge 97. In *Proceedings of IJCAI-97*, Nagoya, Japan, August 1997.
- [10] J. R. Quinlan. *C4.5: Programs for Machine Learning*. The Morgan Kaufmann Series in Machine Learning, 1993.
- [11] Peter Stone and Manuela Veloso. Task decomposition and dynamic role assignment for real-time strategic teamwork. In *Proceedings of the ATAL workshop (Agent Theories, Architectures and Languages)*, Paris, July 1998.
- [12] Manuela Veloso, Peter Stone, Kwun Han, and Sorin Achim. CMUnited: A team of robotic soccer agents collaborating in an adversarial environment. In Hiroaki Kitano, editor, *RoboCup-97: The First Robot World Cup Soccer Games and Conferences*. Springer Verlag, Berlin, 1998.