

Vision, Strategy, and Localization Using the Sony Legged Robots at RoboCup-98

Masahiro Fujita, Manuela Veloso, William Uther, Minoru Asada, Hiroaki Kitano, Vincent Hugel, Patrick Bonnin, Jean-Christophe Bouramoué, and Pierre Blazevic

■ Sony has provided a robot platform for research and development in physical agents, namely, fully autonomous legged robots. In this article, we describe our work using Sony's legged robots to participate at the RoboCup-98 legged robot demonstration and competition. Robotic soccer represents a challenging environment for research in systems with multiple robots that need to achieve concrete objectives, particularly in the presence of an adversary. Furthermore, RoboCup 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 predefined behaviors and learning by instruction.

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-size, medium-size, and legged robots.¹

In this article, we present some of the work developed with the legged robots built by Sony at Carnegie Mellon University (CMU), in the Laboratoire de Robotique de Paris (LRP), and at

Osaka University. These research groups participated at the RoboCup-98 games and demonstrations of the quadruped robots with teams named CMTRIO, LES TITIS PARISIENS, and BABY TIGERS, respectively. Legged robots represent a remarkable advancement for robotics. In the particular context of robotic soccer, legged robots provide an interesting opportunity for robot entertainment.

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, 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. The setup of the game includes the following characteristics:

Colored environment: The important game items, for example, 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 ultraviolet color space. We selected the colors carefully so that a robot could easily distinguish the colored objects.

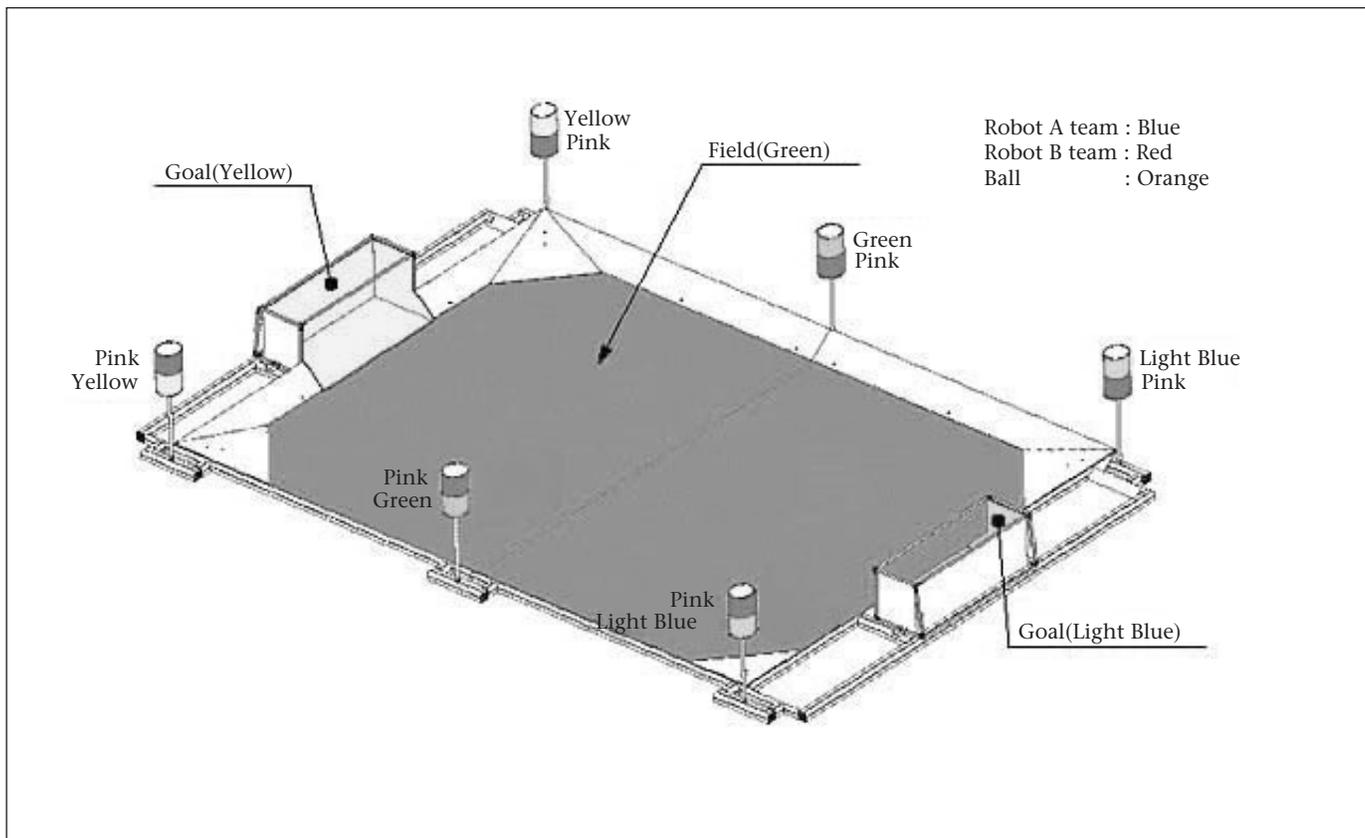


Figure 1. The RoboCup-98 Soccer Field for Legged Robot Competition.

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 (2 meters \times 3 meters) should provide enough space for six robots to navigate while they dribble and pass the ball.

Slanted wall: The walls surrounding the field are at a ± 45 -degree slant. 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 occurrence 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 evaluate software performance based on the RoboCup Physical Challenge (Asada et al. 1998; Kitano et al. 1998), which is defined by the RoboCup Challenge Committees.

A Quadruped Legged Robot

The robot soccer player used in RoboCup Legged Robot Exhibition Match is based on OPEN-R (Fujita and Kageyama 1997), which is proposed by Sony as a standard architecture for robot entertainment systems (Fujita and Kitano 1998). The significant feature of OPEN-R is a decomposition technology for both hardware and software modules. This technology enables us to build various robot styles, such as a quadruped robot and a wheel-based robot as well as various software configurations (Fujita, Kitano, and Kageyama 1998). In addition, for software researchers, OPEN-R can provide a highly reliable robot platform so that they can concentrate on software development for a variety of tasks, including new image-processing algorithms, posture control, and agent architecture. Furthermore, there is no need to develop software from scratch because OPEN-R provides some software modules, such as color detection and walking control. Thus, 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

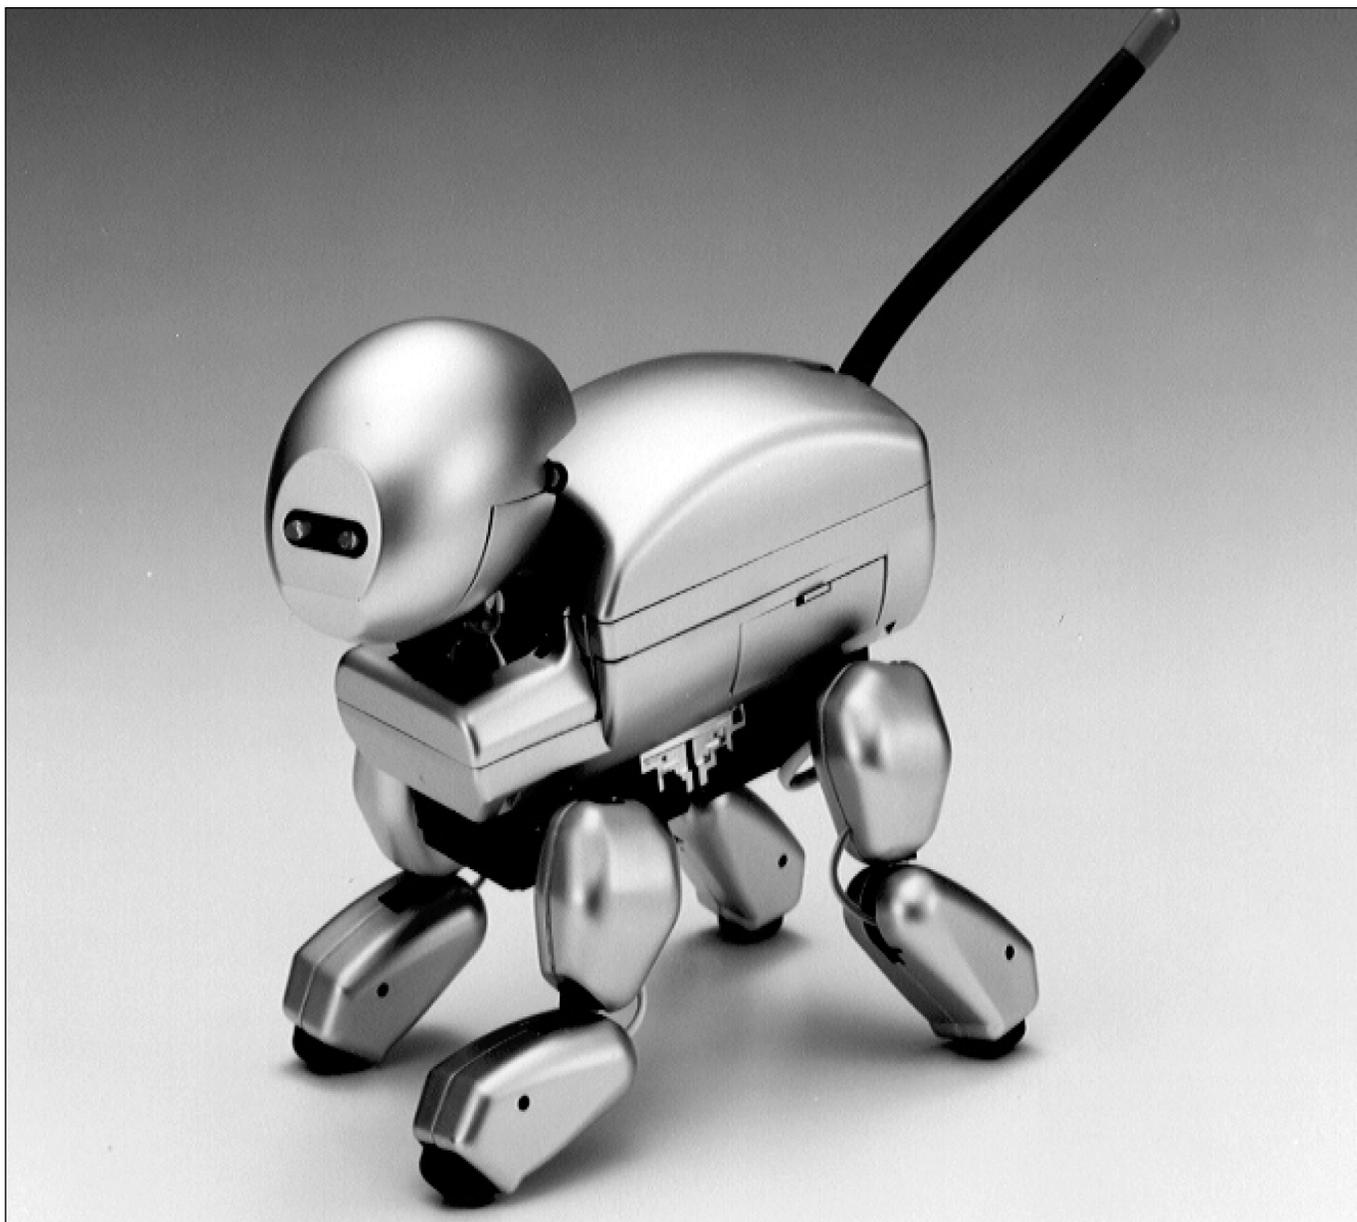


Figure 2. Legged Robot.

OPEN-R system. In RoboCup-98, we prohibited modification of robot hardware; however, in principle, it was possible for us to build various styles of robot with various sensors and actuators that have OPEN-R interface.

For RoboCup-98, we deployed a legged robot with four legs and one head, each of which has three degrees of freedom and rich sensory channels; for example, a head has a color CCD camera, a stereo microphone, touch sensors, and a loudspeaker.

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, because our goal is robot entertainment, we have a different emphasis. We believe that the capability of representation and communication using gesture and motion is important in entertainment applications. Therefore, we chose a mechanical configuration for our robot as a quadruped-legged type, as shown in figure 2.

The merits of the quadruped-legged configuration are that walking control of a quadruped is easier than that of a biped robot and that when in a sitting posture, two hands are free to move, therefore allowing them to present emotions or

communicate with a human. Because each leg or hand has to be used for various purposes besides walking, we assigned three degree of freedom for each leg-hand. In addition, we added a tail and three degrees of freedom for the neck-head so that the robot had enough representation and communication capabilities using motions. During the RoboCup games, legs are not necessarily used for expressing emotions. However, they can be used for sophisticated control of balls, such as passing a ball to the side or back or engaging in deceptive motions.

One possible disadvantage to using a legged robot is that their moving speed is not as fast as wheel-based robots. In the future, the speed issue might be resolved when galloping is made possible. For now, legged robots compete in a dedicated league. Although serious hardware limitations exist, teams with efficient leg-motion coordination can 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 the computational power off board. 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 for processing images from each individual robot viewpoint without any global information will become important in robot entertainment. 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 (Kitano, Fujita, and Kageyama 1998).

CMTRIO: Vision-Based Navigation

Each legged robot for RoboCup-98 was 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 University (CMU), we have therefore decomposed our work according to the following aspects:

Reliable detection of all the relevant col-

ors: The colors are orange (ball), light blue (goal and marker), yellow (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 ongoing research in addressing these issues.

Supervised Learning of Ultraviolet Colors

The Sony legged robot has specialized hardware for the detection of colors. However, this hardware still requires presetting 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 Second, building on 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 have 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:

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

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 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 20 images for training. The learning algorithm converges in less than one hour achieving a high classification accuracy.

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 for the highly unreliable dead reckoning, the field environment for RoboCup-98 included several fixed colored landmarks. Therefore, at CMU, we developed a Bayesian localization procedure (for example, Burgard et al. [1998] and Elfes [1989]). Because the robot cannot keep enough markers in view at all times to calculate its location directly, our algorithm used 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 pass incorporates observations into our probability distribution. The second pass takes into account the movement actions selected. Incorporation of observations is based on Bayes's 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 a priori 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 some of the distributions we want to represent do not have a nice parametric form. For example, 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 on 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 a priori 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 degrees between 2 markers, turns to the left, and then sees an angle of 90 degrees between 2 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 3a. During the turn, the projection spreads over the X, Y plane, representing the increased uncertainty introduced by the dead reckoning as the robot turns (figure 3b). The second observation finally localizes the robot (figure 3c).

Our localization algorithm is invoked actively; that is, the robot searches for two landmarks when the maximum state probability is below a predefined threshold. In our experiments, the robots localize themselves with high accuracy.

Role-Based Behaviors

Following up on our experience with the small-size RoboCup wheeled robots (Stone and Veloso 1999; Veloso et al. 1998), we developed different behaviors based on positioning on the field. As of now, robots play two different roles, namely, (1) attacking and (2) goal keeping.

The procedure used by an attacking robot consists of the following steps: (1) find the ball; (2) localize attacking goal; (3) position behind the ball, aligned with the goal; and (4) shoot or pass.

The procedure used by the goal-keeping robot is a simplified version of the CMUNITED 1997 goal keeper and consisted of the following steps: (1) find the ball, (2) remain close to the goal, (3) move sideways aligned with the ball, and (4) clear the ball when it gets close to it.

LES TITIS PARISIENS: Walking and Strategy

To play soccer, the three Sony pet robots of each team must behave autonomously. With regard to the control system of these machines, the French RoboCup team focused on three points that appear of primary importance, namely, (1) the vision-recognition system, (2) the design of walking patterns, and (3) the strategy chosen to play the soccer games. Thanks to its visual-sensing capacity, every

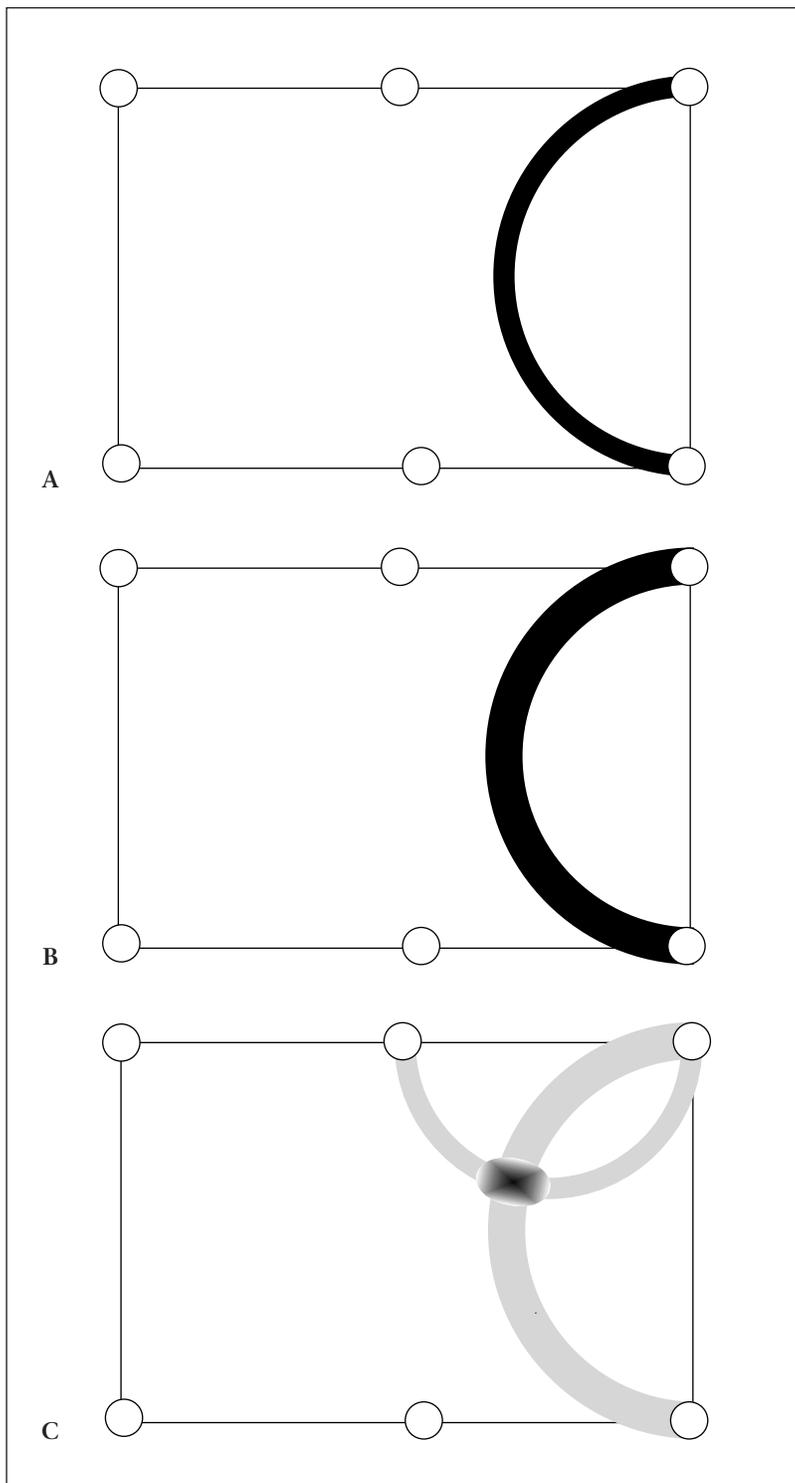


Figure 3: The Positioning Probability.

A. After the first observation. B. After a ± 90 -degree turn. C. After the second observation.

robot can make a virtual representation of the real world on the field. In addition to the sense of sight, the robot must master some efficient walking patterns that give it the ability to

move and change direction as quickly as possible. The two first points refer to basic skills the robot has to acquire; without them, no further development in the matter of behavioral strategies is possible. Assuming the robot owns these essential skills, the third point consists of designing a central decision-making system, taking into account objectives and information inferred from real-time scene recognition and object-detection history.

Vision-Recognition System

For recognition purposes, the French team used an interpretation and a three-dimensional localization procedure of connected components extracted from the result of the Sony color-detection hardware. Through optimization of image data access and processing (Pissaloux and Bonnin 1997), the on-board image treatment procedure gave information of position and orientation of all objects in the scene with respect to the robot at a frequency of 15 images a second. Experience during exhibition matches showed that lighting conditions posed serious difficulties because of reflections or saturation. Unfortunately from time to time, confusion between objects occurred. To solve this problem, temporal filtering was used to ensure temporal coherence of high-level data such as presence and three-dimensional positions of objects. To improve the system in the case of sudden changes of lighting conditions in the environment, the French team thought of the introduction of dynamically adaptive color threshold matching by using multispectral image data fusion (Bonnin, Hoeltzener, and Pissaloux 1995), which requires more computing power.

Walking Skills

The second basic skill is the walking capacity of the quadruped robot. First, emphasis was put on making the robot walk quickly in forward motion. After testing quasistatic walking patterns, the French team decided to use and adapt a special regular symmetric gait called *crawl gait*, about which McGhee developed a well-documented theory in the late 1960s (McGhee 1968; McGhee and Frank 1968). In this kind of gait, each leg reproduces the same trajectory with a phase lag with respect to the consecutive leg on the same side of the body. Opposite legs are half a total cycle period of the leg trajectory out of phase. Crawl is completely defined by the phase lag and the duty factor parameter, called β , which is the fraction of time of the cycle period that the leg spends in contact with the ground during the traction phase. Values of β greater than $3/4$ mean that

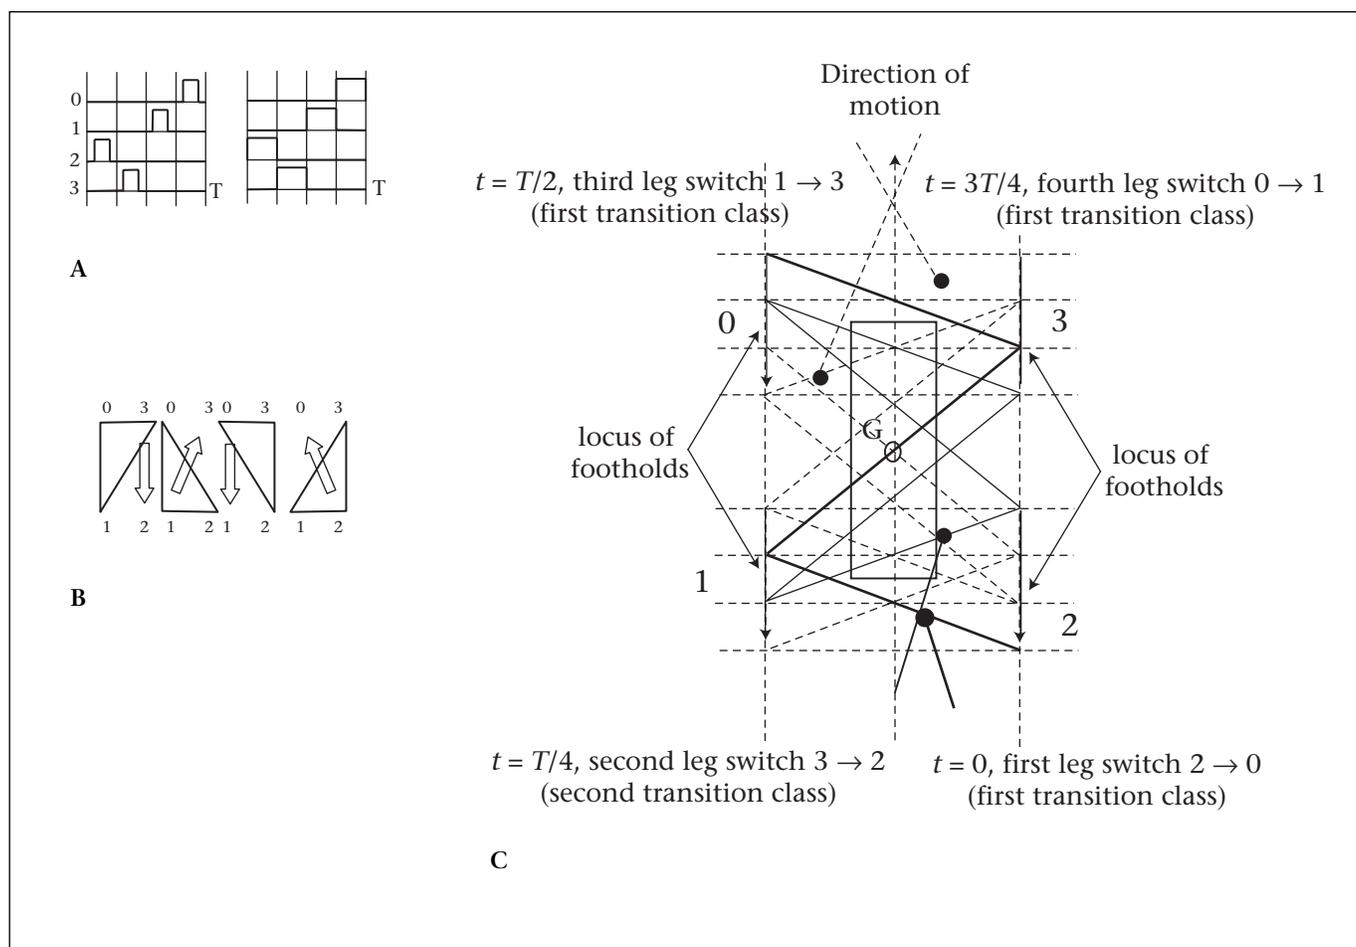


Figure 4. Achieving the Necessary Walking Skills: Forward Crawl Gait.

A. Leg state diagram for the crawl gait with $\beta > 0.75$. B. Sequence of supporting polygons with circulation of events of takeoff and landing of legs. C. Successive support polygons for the forward crawl gait at fractions of cycle time of 0, $T/4$, $T/2$, and $3T/4$.

four legs always keep contact with the ground; in this case, quasistatic balance is guaranteed because the projection of the center of gravity remains strictly inside the stability polygon with a certain stability margin. To speed up motion, the idea is to adapt the crawl gait with duty factor of $3/4$. If β is less than $3/4$, there are phases in the gait where only two legs are on the ground, which appears difficult to master in a first approach using the pet robots. Therefore, we have conducted studies on crawl with β equal to $3/4$. We took special care in mastering the four transition “leg switches” (figure 4). Here, we introduced a specific sideways motion of the center of gravity to avoid losses of balance. As a matter of fact, transitions that imply the takeoff of a rear leg immediately followed by the landing of the corresponding diagonally opposite front leg are critical; at these moments, the center of gravity is exactly located on the diagonal edge of the support polygon. After adapting forward motion, the imple-

mentation of right and left turns allows the robot to change direction. Moreover, a special rotation motion around the center of mass has proved useful to make quick revolution turns without bumping into obstacles. Next developments will focus on achieving gaits with phases involving two legs on the ground; these gaits should increase the motion speed of the quadruped.

Behavioral Strategy

The third and last point to think of concerns the behavioral strategy of the quadruped pet playing soccer. Every team is composed of three robots and faces each other. A whole game can be considered as a multiagent system. Because they cannot communicate with each other, the only way for them to get information on the field is highly dependent on their embedded vision system, which makes the behavioral strategy more difficult to design. In this case, the idea was to allow the robot to

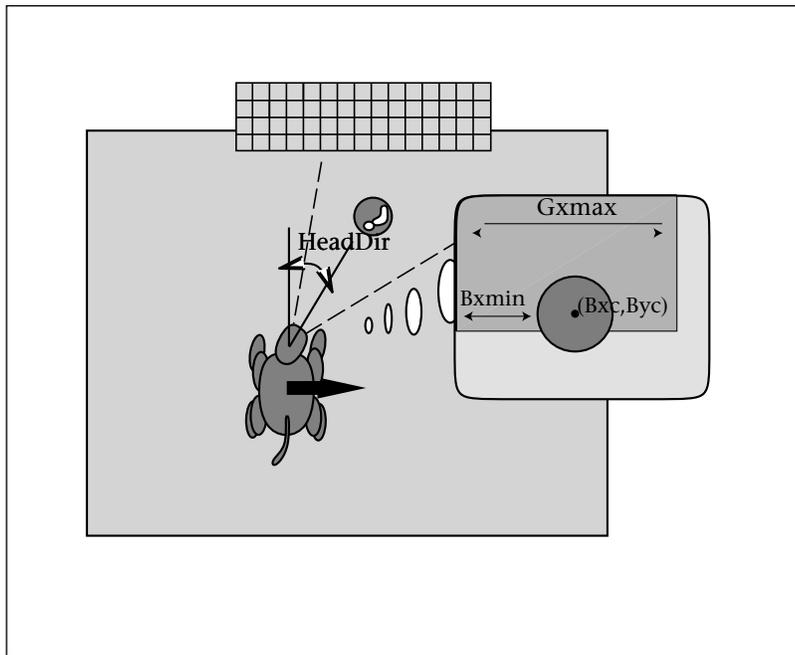


Figure 5. Typical Situation of Right-Shift Motion.

quickly switch the current planning behavior for a more reactive one and vice versa. For example, the pet should be able to generate a trajectory to kick the ball, protect its own goal, and give up all its plans at the next decision cycle if the situation has changed.

For behavioral strategy, we used a bucket brigade-type algorithm: Each agent managed a rule system, and a weight is associated with each rule. Applying a rule means that an agent chooses a role that is a specific behavior: as a defender or as a kicker. During the game, weights are dynamically modified. The weight of a rule is increased if applying this rule allows the pet to meet with the objective required, for example, kick the ball or score a goal; otherwise, the weight is decreased. Moreover, playing several games in simulation allows the pets to learn how to select the right role for a given game situation by adjusting the rules weights. In the 1998 RoboCup competition, every pet robot had to construct its own representation of the environment and tried to match it with the real one, knowing its absolute position with enough accuracy. In the next RoboCup challenge, Sony will provide new quadruped prototypes equipped with wireless communication interfaces, allowing pets to exchange information with each other. Therefore, the reliability of positioning and orientation data will increase, and the use of improved behavioral strategies assuming data transmission between agents will be possible.

BABY TIGERS: Behavior Acquisition by Teaching

The final goal of the Legged Robot Project at Osaka University was 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: (1) basic, (2) basic cooperation, and (3) higher team cooperation. 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 degrees of freedom), which is different from conventional mobile robots (2 or 3 degrees of freedom). From the viewpoint of sensor motor learning and development, multisensory information and multi-degree-of-freedom control should be established simultaneously; that is, affecting each other, the sensory information is abstracted, and the multijoint motions are well coordinated at the same time (Asada 1996).

However, it seems 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 seemed difficult to build. We then decided to adopt a direct teaching method to reduce the number of trials by real robot.

First, we collected test data, pairs of an action command given by a human trainer, and the sensory information during the action execution. Next, we apply C4.5 (Quinlan 1993) to the test data to extract rule sets. Then, the validity of the rule sets were 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 degrees a second) (these abstracted action commands will be decomposed into more primitive motor commands in the future)

Sensory information: Head direction (radius) 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 (figure 5)

Training position and sampling rate: Initial positions of direct teaching by serial line

connection evenly distributed in the field heading the goal; the sampling rate, 300 microseconds (one trial from the initial position to the goal takes about 10 seconds)

We have collected about 740 pairs of 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 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 as follows:

Forward: BallArea > 56, -0:14 \Leftarrow HeadDir < 0:37, GoalXmin \Leftarrow 11

Left-Rotation: BallArea > 49, 0:52 < HeadDir \Leftarrow 1:10, BallWidth \Leftarrow 11

Right-Shift: BallArea \Leftarrow 40, BallYcen > 8, BallXmin > 3, HeadDir \Leftarrow -0:24, GoalArea \Leftarrow 665, GoalXmax \Leftarrow 64

Figure 5 shows the typical situation of right-shift motion.

Because of 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. Table 1 indicates a confusion matrix, showing where some misclassification of the training cases occurs.

Note that several misclassifications are shown (in particular, LEFT_TURN does not have any correct classifications). We are planning to improve the teaching skills and also 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 for solving the problem. These generalization techniques are under investigation.

Conclusion

In this article, 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.

a	b	c	d	e	f	\Leftarrow 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

Table 1. Confusion Matrix.

In RoboCup-98, the research groups from CMU, LRP, and Osaka University played soccer, competing against each other by implementing what they have developed on the common legged robot platform. However, winning the games should not be our first priority. The real purpose of RoboCup is to advance and promote research and development of autonomous robots all together by setting a common target and goal as a standard problem, not to work on discrete interests. In the future, by taking advantage of the merit in the Sony Legged Robot League of the common platform, we hope to integrate or compare various algorithms and outcomes from all research groups.

Notes

1. See Kitano et al. (1997) and www.robocup.org/RoboCup/RoboCup.html for additional information.
2. This tool was developed by Kwun Han, for which we thank him.

References

- Asada, M. 1996. An Agent and an Environment: A View on Having Bodies—A Case Study on Behavior Learning for Vision-Based Mobile Robot. In Proceedings of the 1996 IROS Workshop on toward Real Autonomy, 19–24. Washington, D.C.: Institute of Electrical and Electronics Engineers.
- Asada, M.; Kuniyoshi, Y.; Drogoul, A.; Asama, H.; Mataric, M.; Duhaut, D.; Stone, P.; and Kitano, H. 1998. The RoboCup Physical Agent Challenge: Phase-I. *Applied Artificial Intelligence* 12(2–3): 251–263.
- Bonnin, P.; Hoeltzner, B.; and Pissaloux, E. 1995. A New Way of Image Data Fusion: The Multi-Spectral Cooperative Image Segmentation. In Proceedings of the IEEE International Conference on Image Processing (ICIP 95), 572–575. Washington, D.C.: IEEE Computer Society.
- Burgard, W.; Cremers, A. B.; Fox, D.; Haehnel, D.; Lakemeyer, G.; Schulz, D.; Steiner, W.; and Thrun, S. 1998. The Interactive Museum Tour-Guide Robot. In Proceedings of the Fifteenth National Conference on Artificial Intelligence, 11–18. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Elfes, A. 1989. Occupancy Grids: A Probabilistic Framework for Robot Perception and Navigation. Ph.D. dissertation, Department of Electrical and Computer Engineering, Carnegie Mellon University.
- Fujita, M., and Kageyama, K. 1997. An Open Architecture for Robot Entertainment. Paper presented at the First International Conference on Autonomous Agents, 5–8 February, Los Angeles, California.
- Fujita, M., and Kitano, H. 1998. Development of an Autonomous Quadruped Robot for Robot Entertainment. *Autonomous Robots* 5(1): 7–18.
- Fujita, M.; Kitano, H.; and Kageyama, K. 1998. Reconfigurable Physical Agent. Paper presented at the Second International Conference on Autonomous Agents, 10–13 May, Minneapolis, Minnesota.
- Kitano, H.; Fujita, M.; Zrehen, S.; and Kageyama, K. 1998. Sony Legged Robot for RoboCup Challenge. In Proceedings of the 1998 IEEE International Conference on Robotics and Automation, 2605–2612. Washington, D.C.: IEEE Computer Society.
- Kitano, H.; Tambe, M.; Stone, P.; Veloso, M.; Coradeschi, S.; Osawa, E.; Matsubara, E.; Noda, I.; and Asada, M. 1997. The RoboCup Synthetic Agent Challenge 97. In Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence, 24–29. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.
- McGhee, R. B. 1968. Some Finite-State Aspects of Legged Locomotion. *Mathematical Biosciences* 2:67–84.

McGhee, R. B., and Frank, A. A. 1968. On the Stability Properties of Quadruped Creeping Gaits. *Mathematical Biosciences* 3:331–351.

Pissaloux, E., and Bonnin, P. 1997. On the Evolution of Parallel Computers Dedicated to Image Processing through Examples of French Computers. In *Digital Signal Processing, Volume 7*, ed. R. Y. Gadda, 13–27. San Diego, Calif.: Academic.

Quinlan, J. R. 1993. *c4.5: Programs for Machine Learning*. Morgan Kaufmann Series in Machine Learning. San Francisco, Calif.: Morgan Kaufmann.

Stone, P., and Veloso, M. 1999. Task Decomposition and Dynamic Role Assignment for Real-Time Strategic Teamwork. In *Intelligent Agents V—Proceedings of the Fifth International Workshop on Agent Theories, Architectures, and Languages (ATAL-98)*, eds. J. P. Müller, M. P. Singh, and A. S. Rao, 293–308. Lecture Notes in Artificial Intelligence 1555. Heidelberg: Springer-Verlag.

Veloso, M.; Stone, P.; Han, K.; and Achim, S. 1998. CMUNITED: A Team of Robotic Soccer Agents Collaborating in an Adversarial Environment. In *RoboCup-97: The First Robot World Cup Soccer Games and Conferences*, ed. H. Kitano, 242–256. Berlin: Springer Verlag.



Masahiro Fujita is a senior research scientist at Digital Creatures Laboratory, Sony Corporation. He received a B.A. in electronics and communications from the Waseda University, Tokyo, in 1981, and an M.S.

in electrical engineering from the University of California at Irvine in 1989. His research interests include modular robotics, agent architecture, visual and audio perception, neural networks, and evolutionary systems. His e-mail address is mfujita@pdp.crl.sony.co.jp.



Manuela M. Veloso is associate professor of computer science at Carnegie Mellon University (CMU). She received her Ph.D. in computer science from CMU in 1992. She received a B.S. in electrical engineering

in 1980 and an M.Sc. in electrical and computer engineering in 1984 from the Instituto Superior Tecnico in Lisbon. Veloso's long-term research goal is the effective construction of teams of intelligent agents where cognition, perception, and action are combined to address planning, execution, and learning tasks, in particular, in

uncertain, dynamic, and adversarial environments. Veloso has developed teams of robotic soccer agents in three different leagues that are RoboCup world champions: simulation (1998), CMU-built small-wheeled robots (1997 and 1998), and Sony four-legged dog robots (1998). Veloso was awarded a National Science Foundation Career Award in 1995 and the Allen Newell Medal for Excellence in Research in 1997. Her e-mail address is veloso@cs.cmu.edu.

William Uther is a doctoral student in computer science at Carnegie Mellon University. He received his B.Sc. with honors in computer sciences from the University of Sydney in 1994. Uther, who implemented the CMTRIO RoboCup-98 team, is currently investigating how to use reinforcement learning for agent control and how to automatically decompose problem spaces based on experience. His e-mail address is will+@cs.cmu.edu.



Minoru Asada received a B.E., an M.Sc., and a Ph.D. in control engineering from Osaka University, Osaka, Japan, in 1977, 1979, and 1982, respectively. From 1982 to 1988, he was a research associate of control engineering, Osaka University, Toyonaka, Osaka. In April 1989, he became an associate professor of mechanical engineering for computer-controlled machinery, Osaka University, Suita, Osaka. In April 1995, he became a professor in the same department. Since April 1997, he has been a professor in the Department of Adaptive Machine Systems at the same university. From August 1986 to October 1987, he was a visiting researcher at the Center for Automation Research, University of Maryland. He received the 1992 Best Paper Award at the Institute for Electrical and Electronics Engineers (IEEE)/Robotics Society of Japan (RSJ) International Conference on Intelligent Robots and Systems (IROS92) and the 1996 Best Paper Award from the RSJ. He was a general chair of the IEEE/RSJ 1996 International Conference on Intelligent Robots and Systems (IROS96). Since early 1990, he has been involved in RoboCup activities. His team was the first champion team, with the University of Southern California, in the middle-size league at the first RoboCup. His e-mail address is asada@ams.eng.osaka-u.ac.jp.

Hiroaki Kitano is a senior researcher at Sony Computer Science Laboratories, Inc., and a project director at Kitano Symbiotic Systems Project, Japan Science and Technology Corporation. He received a B.A. in physics from the International Christian



University, Tokyo, and a Ph.D. in computer science from Kyoto University. Since 1988, he has been a visiting researcher at the Center for Machine Translation at Carnegie Mellon University. Kitano received the

Computers and Thought Award from the International Joint Conferences on Artificial Intelligence in 1993. His research interests include RoboCup, computational molecular biology, engineering use of the morphogenesis process, and evolutionary systems.

Vincent Hugel graduated in 1994 from the Ecole des Mines de Nancy, one of the leading French engineering schools. He received his specialization degree in robotics at the National Institute of Nuclear Science and Technology, 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 Ph.D. thesis, which he received in 1999. His e-mail address is hugel@robot.uvsq.fr.

Patrick Bonnin received his Aggregation of Applied Physics in 1985, a Master in 1986 in remote sensing from Paris VII University, and a 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 the French Defense Laboratory in the field of image processing, data fusion, and parallel processing. He joined the Laboratoire de Robotique de Paris 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.

Jean-Christophe Bouramoué is currently working as an AI engineer for a French software company. He received his Master in AI in 1998 (D.E.A I.A.R.F.A - Paris VI). His domains of interest are object-oriented design and modeling, human-machine interface design, and multiagent systems.

Pierre Blazevic received his Aggregation of Applied Physics in 1985, a Master in robotics in 1986, and a Ph.D. from Paris VI University in 1991. He has been an assistant professor at Versailles University since 1992. He has been with the Laboratoire de Robotique de Paris since 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.