

Development of an Augmented Environment and Autonomous Learning for Quadruped Robots

Hayato Kobayashi, Tsugutoyo Osaki, Tetsuro Okuyama,
Akira Ishino, and Ayumi Shinohara

Graduate School of Information Science, Tohoku University, Japan
{kobayashi,osaki,okuyama}@shino.ecei.tohoku.ac.jp
{ishino,ayumi}@ecei.tohoku.ac.jp

Abstract. This paper describes an interactive experimental environment for autonomous soccer robots, which is a soccer field augmented by utilizing camera input and projector output. This environment, in a sense, plays an intermediate role between simulated environments and real environments. We can simulate some parts of real environments, e.g., real objects such as robots or a ball, and reflect simulated data into the real environments, e.g., to visualize the positions on the field, so as to create a situation that allows easy debugging of robot programs. As an application in the augmented environment, we address the learning of goalie strategies on real quadruped robots in penalty kicks. Our robots learn and acquire sophisticated strategies in a fully simulated environment, and then they autonomously adapt to real environments in the augmented environment.

1 Introduction

The experiments on real robots, especially quadruped robots, take much more time and cost than those on PCs. It is also an enormous difference since we often need to consume a lot of energy for treating physical objects such as robots. The use of virtual robots is one of the efficient methods to avoid those difficulties, and so there are many studies on dynamic simulated environments [1,2,3,4,5]. However, since simulated environments cannot produce complete, real environments, we finally need to conduct experiments in the real environments where basic skills heavily depend on complex physical interactions.

In this paper, we propose an interactive *augmented environment*, which serves as a bridge between simulated environments and real environments. Augmented environments, also known as *Augmented Reality* (AR), involve the real environments combined with some data generated by computers. In our augmented environment, robots can obtain both the precise positions of physical objects and the imaginary positions of virtual objects despite being in a real environment. Moreover, human programmers can get debugging information overlaid on a soccer field instead of that in a console terminal, allowing them to concentrate on observing the behavior of the robots in the field. This environment, in a sense, plays an intermediate role between simulated environments and real

environments; considering the fact that we can simulate some parts of real environments.

There are several related studies of augmented environments utilizing real robots [6,7]. The study of Guerra et al. [7] is most related to our study. They proposed the CITIZEN Eco-Be! league (recently, also known as the physical visualization sub-league) where CITIZEN's mini robot Eco-Be! plays soccer in an augmented environment created by a display and a camera. This league gave new interesting and challenging research issues in RoboCup. Our augmented environment can be regarded as a generalization of the technique in the Eco-Be! league to the other real robot leagues. We can easily apply our techniques to different types of robots that are used in other leagues, although we focus only on Sony's quadruped robot AIBO in this paper.

As an application in the augmented environment, we address the learning of goalie strategies in penalty kicks. The goalie strategies involve the skills to save an incoming ball kicked by an opponent player, which is critical to not lose the game. The learning can be also regarded as the two dimensional extension of the study of Kobayashi et al. [8]. They studied the *autonomous learning* to trap a moving ball by utilizing reinforcement learning. The goal of the learning was to acquire a good timing in initiating its trapping motion autonomously, depending on the distance and the speed of the ball, whose movement was restricted to one dimension. The two dimensional extension would need the troublesome programs to find a rebounded ball and return it to the initial position, as well as an accurate localization system. In contrast, our robot in augmented environments can avoid these difficulties by adopting a virtual ball that can be positioned and controlled arbitrarily, where the precise positions of the robot and the ball are measured and calculated by the system. In this way, we can perform much more effective experiments, by using real robots.

The remainder of this paper is organized as follows. In Section 2, we describe our system to create an augmented environment for autonomous soccer robots. In Section 3, we propose an autonomous learning method of goalie strategies in penalty kicks and make our robot learn and acquire goalie strategies on their own. Finally, Section 4 presents our conclusions.

2 Augmented Soccer Field System

2.1 Overview of Our System

In this section, we describe the *Augmented Soccer Field System* for our augmented environment. This system consists of a pair of cameras and a pair of projectors in the ceiling with one computer to control them. The process in the computer is separated into two programs: a *recognition program* and a *virtual application*. The recognition program is a program that recognizes objects, such as robots and balls, in a soccer field based on the images from the ceiling cameras. The virtual application is a program that calculates real coordinates in the field based on the results of the recognition program and then sends the positions of the objects to the robots. It also detects collisions between the objects.

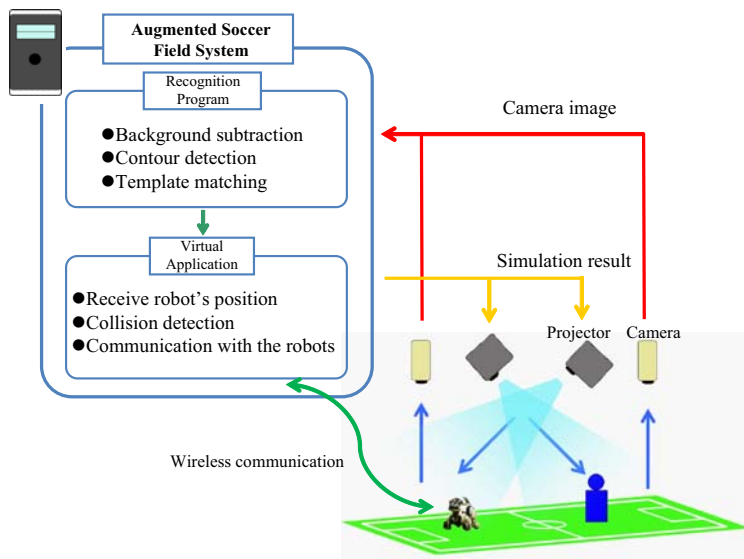


Fig. 1. Overview of our system, which is separated into two programs: A recognition program that recognizes objects with a camera and a virtual application that projects generated data with a projector

Fig. 1 shows the overview of our whole system. Each pair of cameras and projectors was set up on the ceiling located in the center of each half of the field so as to cover the whole field, since one pair can take images with a camera and projects generated data with a projector for a half of the field. We use Dragonfly2 of Point Grey Research Inc. as a camera and LP-XL45 of Sanyo Electric Co., Ltd. as a projector. Only the projector gives our room's lighting when the system is running, so that we can sufficiently recognize projected images on the field. Each projector projects the farther half of the field, since the distance from the ceiling to the field, which is about 2.5 meter height, is too small to cover the whole field.

2.2 Recognition Program

Capturing of Camera Images. Camera images are captured into the computer by utilizing the API for Dragonfly2, and then passed to an object recognition process with the Open Source Computer Vision Library (OpenCV). The processes for two cameras are multithreaded, since they are independent of each other.

Estimation of Camera Parameters. Captured camera images are quite distorted since our cameras have wide-angle lenses. The distortion of images can easily lead to some errors in the object recognition process for the images. Thus,

we need to estimate adjustment parameters for each camera so as to correct such distortion.

We also need to estimate different adjustment parameters (i.e., a projection matrix) for transforming the coordinates of recognized objects in camera images to those in the field. In our system, we obtain the projection matrix by referring 8 points that consist of 4 corners of the half field and 4 corners of the penalty area. Since our cameras and projectors are fixed in our room, it is sufficient to estimate those parameters only once unless their positions are changed.

Recognition Method. We utilize the *background subtraction* method [9] in order to recognize objects on the field. The reason comes from the fact that the object recognition based on colors is usually difficult because the colors can change depending on lighting conditions. The background subtraction method is designed to store a background image without any objects in advance, judge whether or not each pixel in a target image in process belongs in the background image, and separate the target image into a foreground image and the background image. In this paper, the background image means the image of our soccer field without any robots and balls.

We create a masked image that indicates the regions of the objects to be recognized by utilizing the background subtraction method and then identify each object by extracting the contour of the object. Since our robot has its direction (or angle), we identify the orientation of each robot by utilizing template matching.

Calculation of Background Subtraction. The most naive subtraction method is to calculate color subtraction between a current target image and the background image for every pixel and then judge whether the pixel belongs to the background image based on a certain threshold. In general, however, the criterion for each pixel should be changed depending on its location, since brightness varies with location. Thus, we suppose that the degree of color change for each pixel follows a normal distribution $N(\mu, \sigma^2)$. We calculate the mean μ and variance σ^2 of each pixel by capturing 100 background images and regard the pixel as background if the difference between μ and the color value of the pixel is greater than $c\sigma$ with some constant c . Each image has 3 color channels (RGB), and so we separately treat each channel, in order that the constant c of a channel is independent of those of the others. We define that the pixels regarded as background must satisfy the above condition in every channel.

Extraction of Contour. We extract the contour of each object from a masked image calculated by the background subtraction method in order to identify the object. Since there can be unknown objects other than robots and balls in the masked image, we utilize the length of a contour line and the size of a contour area for eliminating unknown objects. Then we can obtain the minimum bounding box covering each identified contour as shown in Fig. 2 (a).

Identification of the Orientation of Robots. Although each robot is specified by a bounding box shaped like a non-regular rectangle, the orientation of

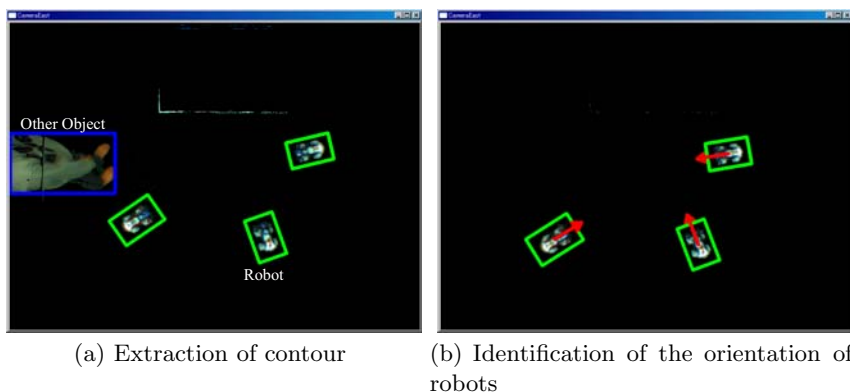


Fig. 2. (a) Results extracting robots as small rectangles and other objects as large rectangles. (b) Results of template matching. An arrow shows the orientation of each robot.

the robot can face either forward or backward. We capture the overhead image of the robot as a template and find a better matching orientation based on the template image. This makes it so that we can recognize multiple robots simultaneously without any markers as shown in Fig. 2 (b). Even if we must use different kinds of robots, we only have to exchange the template.

2.3 Virtual Application

The virtual application carries out physical simulations, such as collision detection based on the results of the recognition program and then projects the simulation results on the field. Thus, the real robots can play a soccer game even with a virtual ball. In addition, the virtual application can communicate with real robots, and so the robots can receive the ideal positions of real objects without an accurate localization system.

Simulator. Our simulator is based on Haribote [5], the dynamic simulator for AIBO, developed by utilizing the Open Dynamic Engine (ODE) and the Open Graphical Library (OpenGL). Although Haribote has a virtual robot represented as the exact model of AIBO, in order to represent the robots recognized by the ceiling cameras, we utilize the minimum rectangular solid for covering the robot. The reason is that the rectangular model can more easily synchronize real robots so that the real robot can manipulate a virtual ball in our simulator. Of course, we can utilize the exact model in order to verify some physical motions in the simulator as well.

Sending of the Positions of Objects. The virtual application always sends the precise positions of the physical objects recognized by the recognition program to real robots through UDP. The robots can receive the positions of the

ball and the other robots and replace noisy inputs with precise ones in their programs. This indicates that we can verify developed strategies in ideal situations in spite of using real robots.

Projection of the Simulator on the Field.

The virtual application projects the simulator on the field in order to visualize the positions of objects in the simulator. In addition, if we display some debugging information in the simulator, we can see the debugging information (e.g., the robot's position, the ball's position, and the robot's state) overlaid on the field instead of that in a console terminal. Although our projector projects data from an oblique perspective, this fact does not cause problems at all by utilizing the automatic trapezoidal correction in the projector.

Fig. 3 shows our simulator projected on the field by the projector. The upper robot is real, and the lower robot and the ball are virtual. In our system, both of the real robot and the virtual robot can manipulate the virtual ball.



Fig. 3. Our simulator that is projected on the field

2.4 Contributions

The main merit of our system is that we can easily find which program exerts a bad influence to the whole program. For example, if we simulate the robot's position and ball's position, we can verify the validity of strategy programs by utilizing the precise output of localization programs. In the same way, if we simulate the beacons' positions, we can verify the validity of localization programs by utilizing the precise output of vision programs.

Another merit of our system is the visualization of debugging information by the projector. For example, we can visualize the particles in localization programs using a particle filter algorithm as well as the objects recognized by vision programs. We can also visualize the current state of strategy programs and even the error messages of programs as strings.

3 Autonomous Learning of Goalie Strategies

In this section, we address the learning of goalie strategies in a penalty kick as an application of our system. We first describe the rule of penalty kicks and the learning method of goalie strategies. Next we experiment in our simulator so as to ensure the validity of the method, and finally experiment in our augmented environment so as to verify the effectiveness of our augmented soccer field system.

3.1 Penalty Kick

In the rules for the RoboCup standard platform league, the penalty kick is carried out with one attacker and one goalie. If the ball goes into the goal within the time

limit (1 minute), the penalty kick is deemed successful. If the time limit expires or if the ball leaves out the field, the penalty kick is deemed unsuccessful. In addition, if the attacker enters the penalty area then the penalty kick is deemed unsuccessful; if the goalie leaves the penalty area then a goal will be awarded to the attacker. The aim of our goalie is to acquire the best strategy to save the goal.

3.2 Learning Method

In this paper, we apply reinforcement learning algorithms [10]. Since reinforcement learning requires no background knowledge, all we need to do is to give the robots some appropriate rewards so that they can successfully learn goalie strategies.

The reinforcement learning process is described as a sequence of states, actions, and rewards,

$$s_0, a_0, r_1, s_1, a_1, r_2, \dots, s_i, a_i, r_{i+1}, s_{i+1}, a_{i+1}, r_{i+2}, \dots,$$

which is a reflection of the interaction between a learner and an environment. Here, $s_t \in S$ is a state given from the environment to the learner at time t ($t \geq 0$), and $a_t \in \mathcal{A}(s_t)$ is an action taken by the learner for the state s_t , where $\mathcal{A}(s_t)$ is the set of actions available in the state s_t . One time step later, the learner receives a numerical reward $r_{t+1} \in \mathcal{R}$, in part as a consequence of its action, and finds itself in a new state s_{t+1} .

- **State.** The state of our learner is represented by a list of 6 scalar variables (r, ϕ, V_v, V_h, X, Y). The variables r and ϕ represent the position of the ball on a polar coordinate and refer to the distance and angle from the robot to the ball estimated by our vision system. The variables V_v and V_h represent the velocity of the ball and refer to the vertical and horizontal velocity from the robot. The variables X and Y represent the relative position of the robot from its own side's goal and refer to the values of x -axis and y -axis on a right handed Cartesian coordinate system with its origin at the center of the goal, i.e., an absolute position (2700 mm, 0 mm), and a positive x -axis toward the center of the field, i.e., an absolute position (0 mm, 0 mm). In order to allocate the required minimum state space for our problem, we empirically limited the range of those state variables such that $r, \phi, V_v, V_h, X,$ and Y are in [0 mm, 1000 mm], $[-\pi/2$ rad, $\pi/2$ rad], [-50 mm, 0 mm], [-50 mm, 50 mm], [-500 mm, 1000 mm], and [-700 mm, 700 mm], respectively.
- **Action.** The action of our learner takes one of the following 10 exclusive actions. One is *save* to stop an incoming ball, which performs the motion to spread out robot's front legs. The action *save* cannot be interrupted for 350 ms until the motion finishes. Another is *stay* to prepare for an opponent's shooting, which moves its head to watch the ball without walking. The others are 8 directional walking actions (i.e., vertically, horizontally, or diagonally) to intercept an opponent's shooting.

- **Reward.** Our learner receives one of the following 5 kinds of rewards. A negative reward is also called a *punishment*.
 - *save_punishment* The punishment is -0.02, if the robot performs *save* action, since the time during *save* motion has the risk losing the ball.
 - *passive_punishment* The punishment is -0.0000001, if the robot performs an action other than *save* action, since the robot should take various actions in the initial phase of learning.
 - *lost_punishment* The punishment is -10, if the ball goes into the goal.
 - *save_reward* The reward is 0.5, if the robot stops the ball by *save* action, since it may be safest to save the goal.
 - *dist_reward* The reward is $1 - |ydist|/112.5$, if the robot saves the goal during a penalty kick, since it may be the safest saving. The variable *ydist* is the *y*-axis distance of the ball, which is limited in [-225 mm, 225 mm] based on the covering range 450 mm of *save* action. That is, *dist_reward* takes a value in [-1, 1] based on *ydist*.

Algorithm 1 shows the autonomous learning algorithm of goalie strategies used in our study. It is a combination of the episodic SMDP Sarsa(λ) with the linear tile-coding function approximation (also known as CMAC). This is one of the most popular reinforcement learning algorithms, as seen by its use in Kobayashi et al. [8]. In our experiments, we define the period from the starting to the ending of a penalty kick as one *episode*.

Here, F_a is a *feature set* specified by tile coding with each action a . In this paper, we use 6-dimensional tiling, where we set the number of tilings to 32 and the number of tiles to about 400000. We also set the tile widths of r , ϕ , V_v , V_h , X , and Y to a quarter of the ranges, i.e., 250 mm, $\pi/4$ rad, 12.5 mm, 25 mm, 375 mm, and 350 mm, respectively. The vector $\vec{\theta}$ is a *primary memory vector*, also known as a *learning weight vector*, and Q_a is a *Q-value*, which is represented by the sum of $\vec{\theta}$ for each value of F_a . The policy ϵ -greedy selects a random action with probability ϵ , and otherwise, it selects the action with the maximum *Q-value*. We set $\epsilon = 0.001$. Moreover, \vec{e} is an *eligibility trace*, which stores the credit that past action choices should receive for current rewards. λ is a *trace-decay parameter* for the eligibility trace, and we set $\lambda = 0.9$. We set the *learning rate parameter* $\alpha = 0.18$ and the *discount rate parameter* $\gamma = 1.0$.

3.3 Experiments in Simulated Environments

We experiment in the simulated environment (or simulator) to ensure the validity of the algorithm in the previous section. In our simulator we do not treat the dynamics of robots, while the scales of the field, ball, and robots are set to the same values in real environments. Our virtual robot is represented as a rectangular solid, and if *save* action is performed then the width of the rectangular solid stretches out. State inputs have their ideal values without any noise, and action outputs are set to almost the same values as real robots. In experiments from now on, we assume that the attacker dribbles the ball to any point and then shoots it to the goal so as to confuse the goalie. Thus, we set the ball in the simulator to perform such behavior in the beginning of penalty kicks.

Algorithm 1. Algorithm of the autonomous learning of goalie strategies

```

while still not acquiring goalie strategies do
  initialize penalty shootout settings;
   $s \leftarrow$  a state observed in the environment;
  forall  $a \in \mathcal{A}(s)$  do
     $F_a \leftarrow$  set of tiles for  $a, s$ ;
     $Q_a \leftarrow \sum_{i \in F_a} \theta(i)$ ;
  end
   $lastAction \leftarrow$  an optimal action selected by  $\epsilon$ -greedy;
   $\vec{e} \leftarrow 0$ ;
  forall  $i \in F_{lastAction}$  do  $e(i) \leftarrow 1$ ;
   $reward \leftarrow 0$ ;
  while during a penalty kick do
    do  $lastAction$ ;
    if  $lastAction = guard$  then
       $reward \leftarrow save\_punishment$ ;
    else
       $reward \leftarrow passive\_punishment$ ;
    end
     $\delta \leftarrow reward - Q_{lastAction}$ ;
     $s \leftarrow$  a state observed in the environment;
    forall  $a \in \mathcal{A}(s)$  do
       $F_a \leftarrow$  set of tiles for  $a, s$ ;
       $Q_a \leftarrow \sum_{i \in F_a} \theta(i)$ ;
    end
     $lastAction \leftarrow$  an optimal action selected by  $\epsilon$ -greedy;
     $\delta \leftarrow \delta + Q_{lastAction}$ ;
     $\vec{\theta} \leftarrow \vec{\theta} + \alpha \delta \vec{e}$ ;
     $Q_{lastAction} \leftarrow \sum_{i \in F_{lastAction}} \theta(i)$ ;
     $\vec{e} \leftarrow \gamma \lambda \vec{e}$ ;
    forall  $a \in \mathcal{A}(s)$  s.t.  $a \neq lastAction$  do
      forall  $i \in F_a$  do  $e(i) \leftarrow 0$ ;
    end
    forall  $i \in F_{lastAction}$  do  $e(i) \leftarrow 1$ ;
  end
  if lost the goal then
     $reward \leftarrow lost\_punishment$ ;
  else
     $reward \leftarrow dist\_reward$ ;
    if  $lastAction = guard$  then
       $reward \leftarrow reward + save\_reward$ ;
    end
  end
   $\delta \leftarrow reward - Q_{lastAction}$ ;
   $\vec{\theta} \leftarrow \vec{\theta} + \alpha \delta \vec{e}$ ;
end

```

We made our robot on the simulator learn goalie strategies in 2000 episodes, and the experiment took about 20 minutes. Fig. 4 (a) shows the learning process




Fig. 4. Experimental results in our simulated environment and augmented environment, whose success rates indicate how many times the robot saved the goal in the past 200 episodes and 20 episodes, respectively

in the simulated environment in terms of the success rate of saving. The success rate of saving means how many times the robot saved the goal in the past 200 episodes. The figure indicates that the success rate reached more than 90% at 2000 episode. Actually, we repeated the same experiment 5 times and got the average of the final success rates, 93.32%. This result indicates that our virtual robot can successfully acquire a better goalie strategy in penalty kicks.

Fig. 5 shows an acquired goalie strategy (i.e., a state-action mapping) after 2000 episodes. This intuitively visualizes all of the better actions that the learner will select at each state in the environment. The figure indicates that *save* actions shown by squares are mostly on the way to the destination of a kicked ball so as to intercept the incoming ball. Although walking actions shown by arrows have seemingly no regularity, looking at the arrows only around the squares, we find that the directions of the walking actions mostly face to the states with *save* actions, so that the robot can move to intercept the ball if possible. The states with the other walking actions (far from the squares) seem not to be learned yet. The reason is that our virtual robot never loses the ball in the simulator, and so it cannot go through the experience of those states. In the case of the penalty kick, this may not be a big problem, since the robot knows the fixed initial positions of the attacker and the ball in advance.

3.4 Experiments in Augmented Environments

In this section we experiment in our augmented environment, where we set the ball as simulated and project to the field. Since a virtual ball is easily manipulated by our system, our robot does not need to perform troublesome tasks such as restoring the ball, and this advantage easily enables the autonomous learning of goalie strategies. Moreover, our robot can receive the ideal positions of the ball and the robot itself from our system. This also means that we can experiment using real robots in real environments, even if some basic programs such as vision and localization systems are premature or unfinished.

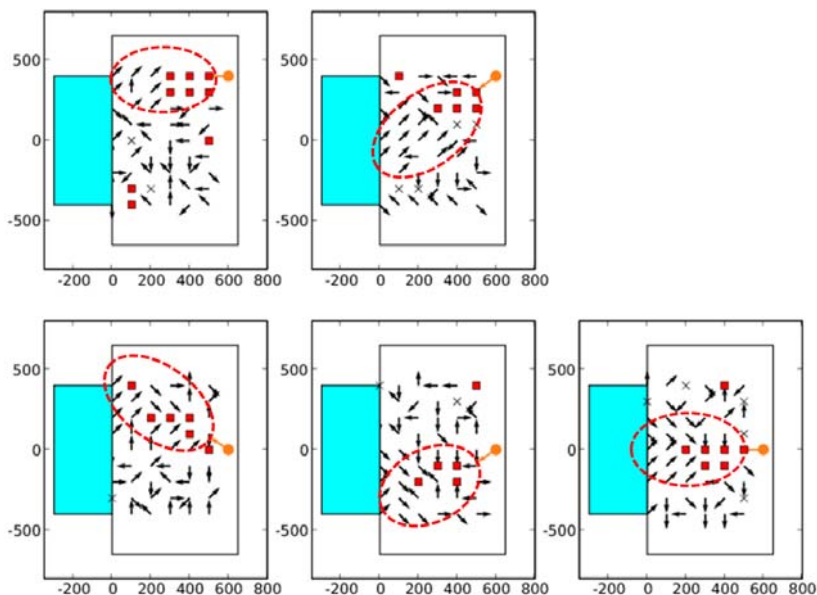


Fig. 5. State-action mapping that was acquired by the learning in the simulator. Each figure shows the left half field, and its coordinate shows the relative position of the robot from the goal, (i.e., X and Y in the state of our learner). Each pair of a circle and arrow shows the velocity of a ball, and each figure indicates which action our robot selects in the situation. A square, cross, and arrow show *save*, *stay*, and walking actions, respectively. In the area enclosed with a dashed circle in each figure, better actions are selected so as to intercept the incoming ball.

We made our robot learn goalie strategies in 200 episodes, starting from the strategy acquired in the previous section, since it takes a lot of time and cost for real robots to perform a lot of actions. The experiment took about 45 minutes with 1 battery. Fig. 4 (b) shows the learning process in the augmented environment in terms of the success rate of saving. In the initial phase of the process, the success rate was only about 50% despite the initial strategy whose success rate was more than 90% in the simulator. This suggests that there is a certain gap between simulated environments and real environments, especially in terms of quadrupedal locomotion. In the latter phase of the process, the success rate increased to about 80%, which may be the maximum value since real robots cannot walk as smoothly as virtual robots. This also suggests that our system can fill in such a gap. We repeated the same experiment 3 times and got the average success rates in the first 50 episodes and the last 50 episodes, which were 56.7% and 70.7%, respectively. The strategies that were acquired in the augmented environment should work well even in real environments as long as our robot can accurately estimate the positions of the ball and the robot itself.

4 Conclusion

In this paper, we presented an interactive experimental environment to bridge the gap between simulated environments and real environments, and developed an augmented soccer field system to realize the environment. Moreover, we addressed the learning of goalie strategies in penalty kicks as an application of our system. By utilizing our system, our robot could autonomously learn and acquire better strategies without human intervention. All movies of the demonstration of our system and the experiments of the learning in our system are available on-line¹.

References

1. Zagal, J.C., del Solar, J.R.: UCHILSIM: A Dynamically and Visually Realistic Simulator for the RoboCup Four Legged League. In: Nardi, D., Riedmiller, M., Sammut, C., Santos-Victor, J. (eds.) RoboCup 2004. LNCS (LNAI), vol. 3276, pp. 34–45. Springer, Heidelberg (2005)
2. Laue, T., Spiess, K., Röfer, T.: SimRobot - A General Physical Robot Simulator and Its Application in RoboCup. In: Bredenfeld, A., Jacoff, A., Noda, I., Takahashi, Y. (eds.) RoboCup 2005. LNCS (LNAI), vol. 4020, pp. 173–183. Springer, Heidelberg (2006)
3. Zaratti, M., Fratarcangeli, M., Iocchi, L.: A 3D Simulator of Multiple Legged Robots based on USARSim. In: Lakemeyer, G., Sklar, E., Sorrenti, D.G., Takahashi, T. (eds.) RoboCup 2006: Robot Soccer World Cup X. LNCS (LNAI), vol. 4434, pp. 13–24. Springer, Heidelberg (2007)
4. Microsoft: Microsoft Robotics Studio (2007), <http://msdn.microsoft.com/robotics/>
5. Takeshita, K., Okuzumi, T., Kase, S., Hasegawa, Y., Mitsumoto, H., Ueda, R., Umeda, K., Osumi, H., Arai, T.: Technical Report of Team ARAIBO. Technical report, ARAIBO (2007), <http://araibo.is-a-geek.com/>
6. Sugimoto, M., Kagotani, G., Kojima, M., Nii, H., Nakamura, A., Inami, M.: Augmented coliseum: display-based computing for augmented reality inspiration computing robot. In: SIGGRAPH 2005: ACM SIGGRAPH 2005 Emerging technologies, p. 1. ACM, New York (2005)
7. da Silva Guerra, R., Boedecker, J., Mayer, N., Yanagimachi, S., Hirokawa, Y., Yoshikawa, K., Namekawa, M., Asada, M.: CITIZEN Eco-Bel! League: bringing new flexibility for research and education to RoboCup. In: Proceedings of the Meeting of Special Interest Group on AI Challenges, vol. 23, pp. 13–18 (2006)
8. Kobayashi, H., Osaki, T., Williams, E., Ishino, A., Shinohara, A.: Autonomous Learning of Ball Trapping in the Four-legged Robot League. In: Lakemeyer, G., Sklar, E., Sorrenti, D.G., Takahashi, T. (eds.) RoboCup 2006: Robot Soccer World Cup X. LNCS (LNAI), vol. 4434, pp. 86–97. Springer, Heidelberg (2007)
9. Piccardi, M.: Background subtraction techniques: a review. In: IEEE International Conference on Systems, Man and Cybernetics 2004, vol. 4, pp. 3099–3104 (2004)
10. Sutton, R.S., Barto, A.G.: Reinforcement Learning: An Introduction. MIT Press, Cambridge (1998)

¹ URL: <<http://www.shino.ecei.tohoku.ac.jp/~kobayashi/movies.html#goalie>>