The UNSW RoboCup 2001 Sony Legged Robot League Team

Spencer Chen, Martin Siu, Thomas Vogelgesang, Tak Fai Yik, Bernhard Hengst, Son Bao Pham, and Claude Sammut

School of Computer Science and Engineering
University of New South Wales
Sydney, Australia

Abstract. In 2001, the UNSW United team in the Sony legged robot league successfully defended its title. While the main effort in last year's competition was to develop sound low-level skills, this year's team focussed primarily on experimenting with new behaviours. An important part of the team's preparation was playing practice matches in which the behaviour of the robots could be studied under actual game-play conditions. In this paper, we describe the evolution of the software from previous years and the new skills displayed by the robots.

1 Introduction

In the RoboCup robot soccer tournament, the Sony legged league differs from the other leagues in that all the teams use identical hardware, the Sony ERS-210 robot. Thus, a team's strength is based on the software that implements low-level skills as well as game-play behaviour. UNSW's team in RoboCup 2000 made many significant advances in programming the low-level skills of the previous model, the ERS-111 robot. These included improved methods for vision, localisation and locomotion [3,2]. The goal of the 2001 team was to build on these skills to develop new behaviours for stronger game play.

The first step in developing the software was to port last year's code to the new robots. The ERS-210 has several advantages over the older model. The leg motors are substantially stronger, permitting the robots to move more quickly. The on board MIPS R4000 series processor is also much faster. However, the shape of the new robot is somewhat different from the previous robot. Since many of the skills to do with pushing the ball are dependent on the geometry of the body, substantial changes were required to the code, including abandoning some skills and inventing new ones.

Once the porting of the infrastructure modules for vision, localisation, locomotion and sound were completed, development of new behaviours could begin. A very important part of our development strategy was to play weekly matches between "last week's champion" and new code written during the week. The winning code retained the trophy for that week. We believe that observing and testing behaviours in realistic game play is vital. One of the major lessons we

A. Birk, S. Coradeschi, and S. Tadokoro (Eds.): RoboCup 2001, LNAI 2377, pp. 39–48, 2002.

learned from our 1999 entry [1], was that understanding the interactions of the robots on the field was paramount. Behaviours tested using a single robot are unlikely to work as expected when the robot is being interfered with by other players. Furthermore, close observation of a game can reveal where subtle adjustments to behaviours in particular situations can result in an advantage for the team.

In the remainder of this paper, we describe the modifications made to the infrastructure and we describe the new skills developed for this year's team. We refer the reader to the report of the 2000 team [3] for detailed descriptions of the previously developed software.

2 Vision

2.1 High Resolution

The vision system of the Sony robots can deliver images at low, medium or high resolution. In previous years, we limited ourselves to using medium resolution so that we could maintain a high camera frame rate. The faster CPU in the new model allowed us to process high resolution images with no loss of speed, almost keeping up with the maximum frame rate of the camera. The benefits of high resolution were that images contained larger and better defined blobs of colour which made object recognition and localisation more reliable and, in particular, detection of the other robots was improved. Furthermore, we were able to reliably detect the ball across the length of the field. Switching to high resolution involved minor changes to the image processing routines and changes to tools for classifying sample images and generating the colour lookup table.

2.2 Polygon Growing and Decision Trees

All of our object recognition relies on finding regions of specific colours. Thus, accurate colour classification is critical. Since colours may appear differently under different lighting conditions, we train the colour classifier on sample images taken on the competition playing field.

Starting in 1999 and until half way through the 2001 competition, we trained the system using a Polygon Growing Algorithm (PGA). The polygon based classifier was used during the pool matches, however, we were unhappy with its behaviour. The lights at the competition venue were very bright and caused some significant misclassifications. Earlier in the year, we had briefly experimented with using Quinlan's C4.5 decision tree learning algorithm [4] to construct the classifier. We decided to try this again during the rest day between the pool matches and the finals and found that the decision tree classifier was more reliable under the competition lighting conditions. We took the chance of using the decision trees in the final rounds and this gamble seems to have paid off.

The input to C4.5 is a set of examples, each of which is characterised by a list of attribute/value pairs and a class label. In this case, the attributes are the Y, U, V values of each pixel and the class label is the colour obtained from manually classifying sample images. The output of the learning algorithm is a decision tree that is easily translated into a large if-statement in C.

The most likely cause of the improved classification accuracy is that in the PGA, we arbitrarily divide the colour space into a fixed set of Y-planes and construct polygons for each plane. So, in a sense, the PGA is not fully 3D. Whereas, the decision tree classifier uses the full range of all the attributes in its classification. One drawback of using C4.5 are that it requires more training examples, particularly negative examples, to avoid over generalisation.

2.3 Blob Formation

Since all object recognition is based on finding blobs of particular colours, the blob formation algorithm is critical. However, it is also the most time consuming section of the code and causes a slow down in the camera frame rate. When the ball is close to the camera, it fills the image,creating a large patch of orange. When large blobs appeared in the image, our previous code would slow down from about 25fps to 19fps. A new blob forming algorithm was written to solve this problem. It requires two scans over an image. The first pass goes through each horizontal lines and creates run length information on continuous colour pixels. The second pass merges the same colour run length segments if they overlap each other. The the new blob forming algorithm can keep up with the maximum frame rate of the camera.

2.4 Extending Occluded Beacons

The playing field of the Sony legged robot league has six coloured landmarks or beacons placed at each corner of the field and at the centres of each side. These beacons are used by the robots to localise. When a beacon is partially obscured, the distance to it may be incorrectly calculated. The distance to the beacon is based on the Cartesian distance between the centroids of the two coloured blobs that make up a beacon. When one half of the beacon is partially occluded, the distance between the two centroids will appear smaller than its actual distance. In this case the distance to the object will be over estimated and result in poor localisation. To circumvent this problem, we extend the partially occluded blob and shift its centre, as in Figure 1.

2.5 Looking at the Ball at Close Range

One of the differences between the ERS-111 and the ERS-210 robots is that the older model's head has a longer "snout". Since the camera is in the tip of the snout, the ERS-111 loses sight of the ball as it comes under the chin and close to the robot's chest. The ERS-210's shorter snout permits the head to look down and track the ball much closer to the robot's body, which improves ball handling skills.

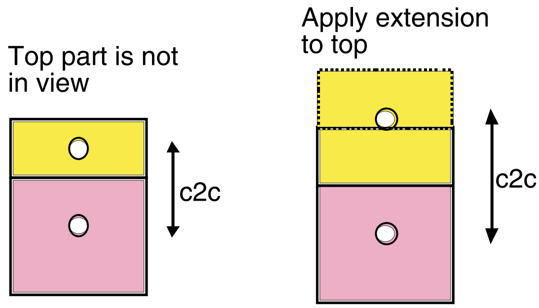

Fig. 1. Extension to Beacons

When the ball is close, the limited field of view of the camera means that the view of the ball is clipped. This affects the calculation of the distance to the ball since it is based on the size of the orange blob. When we recognise that clipping may be occurring, only the top of the ball's bounding box is used to estimate distance. Combining this with tracking the top of the ball rather than the centre of the ball, improves the accuracy of ball distance estimation when the ball is within 20cm of the robot.

This case is a good illustration of the fact that the overall level of performance of the robots results from an accumulation of skills for handling the many different situations that might occur during a game. So far, the only way we have of identifying all of these possible situations is by playing practice matches and carefully observing the behaviour of the robots. These observations also resulted in finding various bugs, including problems with the 3D transforms that take into account the orientation of the head relative to the robot's body.

3 Localisation and World Model

Each robot maintains a world model. Since we know the dimensions of the playing field, the world model is simply a map in which the goals and landmarks are fixed and the mobile objects are placed in the map at their estimated locations. A consequence of using only medium resolution in last year's vision system was the it was difficult to reliably detect the colours of the robots' uniforms. Therefore, the world model did not include the locations of other robots.

With the switch to high resolution, it became possible to add other players into the world model. Each robot now keeps track of the last seen team mate and opponent, allowing more complex team play, as described in section 6. Because the uniform of each robot consists of many patches and the robots are constantly in motion, it remains difficult to separate multiple robots and to accurately judge the distance to the robot. These are topics of on-going investigation.

Several other adjustments were made to the localisation system, including modifications to improve the accuracy of triangulation on two beacons and discarding distances to far away beacons, since these measurements are often unreliable.

4 Locomotion

The walking style developed for the ERS-210 is based on last year's low stance walk originally developed for the ERS-111. However, the shape and dimensions of the robot are somewhat different. The shoulder joints are closer together and the diameters of the limbs are larger. Thus, the locomotion algorithms had to be tuned for the new robots. The limbs were extended further outward, which helped us to develop some new behaviours. For example, by having the front limbs extended further to the front, the ball is more likely to be trapped between the legs. This gives us the ability to turn with the ball (see section 5.1). One of the main features of the low stance walk is its stability. Our robots rarely fall over, while robots with a higher stance frequently fall when pushed. On the ERS-210, with the help of stronger motors, we are able to keep the same level of stability while raising the height of the walk slightly. This has improved sideway movement and consequently the general agility of the robots.

5 Skills

5.1 Paw Dribble and Turn Kick

Through playing many practise games we noticed that if one of the robots hits the ball with its front paw, often a straight powerful pass results. We decided to implement this as a paw dribble. This requires a very accurate approach to position the ball so that it connects with the paw for a straight kick. The walk is dynamically adjusted, based on the ball distance and heading, and in the final approach a downward tilt and rotation of the head is used for finer adjustment. This kick makes the game more dynamic and often allows the robot to push the ball past an opponent on the wing.

The turn kick was developed to allow the robots to kick the ball in directions 90° and greater from the robot's heading when it is able to trap the ball between its front legs. It is especially useful on the edge of the field when a perpendicular approach is often the only one possible. Once the ball is under the head, a quick turning motion is executed resulting in a powerful and reasonably accurate pass. There is a complex interplay between the exact turn parameters, and the position of the head. Thus, a substantial amount of experimentation is required to construct a reliable kick.

The very last goal of the tournament was scored when a UNSW robot approached the ball facing its own goal (a very dangerous thing to do) and then executed an almost perfect 180° turn to take the ball down the field and past the CMU goalie.

5.2 Chest Push and Double-Handed Kick

When the ball is close to the front of the robot's body, the body is thrust forward while the paws remain still on the ground, resulting in the chest pushing the ball forward. This kick is not a powerful one but it is easy to setup, reliable and quick for the robot to return to its original position. It is very effective near the target goal as it is able to push the ball past the goalie into the goal.

The "double-handed kick" was introduced in the 2000 code. This version brought both of the front arms up and pushed them down onto the ball to shoot. Because of changes in the robot's shape, this simple method no longer works. In the current version, the pressing force is changed to a squeezing force. The two arms are moved inward together, at floor level, to grab the ball and then the arms are moved up straight. When the ball is raised, it hits the head and is pushed downward under the arms. Now the ball is underneath the raised arms. The two arms now squeeze the ball out by moving both hands down quickly. This is very powerful and accurate and can score a goal across the full field length. However, to make the grab part reliable requires time to set up and during the game we usually do not have this luxury, so this kick was rarely effective. This is true of most of the kicks developed by all the teams.

6 Interaction with Other Robots

6.1 Team Mate Avoidance

During a game, robots tend to get into scrums with both team mates and opponents. Legs lock together and odometry is confused. Furthermore, other robots at close range block the field of view. To avoid these problems, our team has a "back off from team mate policy". When a team mate is detected within a range of about 20cm and the location of the ball is not known, the robot aborts the current action and takes a few steps backward. If ball position is known and a team mate is detected, the distance to the ball and the distance to the team mate are compared to determine which robot is closer to the ball. If the player itself is closer, than the team mate is ignored, and continues its attack, otherwise, the attack is aborted and the robot begins backing away from the team mate. As the robot backs away, it also turns to get behind the team mate and face towards the target goal. This often blocks an opponent's attack and also gives the player a better view of the situation. Since it tends to stay behind the scrum, the robot is often in a position to pick up a ball that comes loose.

At the same time that the player decides to back away from its team mate, it emits a high pitched sound signal telling the team mate that it is backing off. This signal is used by the other robot to decide in which direction to turn the ball.

6.2 Opponent Avoidance

The turn kick is useful in avoiding opponents and keeping the ball away from them. When a player has the ball in its possession and it detects an opponent at close range, it performs a turn kick. If it has heard a sound signal from its team mate, it knows that there is a supporting robot behind and to one side. so it turns in the direction of the team mate. If no signal was heard, it turns away from the opponent. This behaviour is effective in moving the ball away from an opponent and keeping possession of the ball, especially in a head-on situation.

6.3 Sound

The ERS-210 has on board stereo microphones and a speaker. Using sine waves, we are able to drive the speaker at any frequency within the sampling rate limitations. To detect a tone frequency, we utilize the orthogonality of the complex Fourier series to extract the power coefficient of that particular frequency and are able to detect signals on several channels at the same time without effecting the strategy performance. Since there are stereo microphones, it is possible to roughly determine if the signal originated from the left or right side of the head. However, due to the noise of motors and collision with other robots, we are only able to separate signal from noise within 70cm radius of the head.

6.4 Goalie Avoidance

A new "keeper charging" foul was introduced this year. A player is removed from the field for 30 seconds if it charges the goal keeper while it is holding the ball. We were unsure what effect this rule wold have on the game since losing a player is obviously a serious disadvantage but it was not clear how often such a situation would arise. Furthermore, it is not possible for the robot's vision system to recognise if the goalie has possession of the ball. To make matters worse, judging the distance to a robot is very inaccurate.

Because of our uncertainty about handling this foul, we decided to run two different forwards, one which attempted to avoid the goalie when the ball appeared to be close to it and another that simply carried on as normal. Figure 2 explains how the goalie avoidance is implemented.

6.5 Kick Off

There are several alternatives for the kick off, depending on the position of the team mate and opponents. If the player kicking off sees a team mate to one side, it will always kick the ball in the team mate's direction at an angle of about 40°. This helps the team retain possession of the ball after kick off and is often make it possible to get behind the opponent's forwards.

before kicking off, the robot looks up to check if there are any robots ahead. If it sees an obvious gap on either side of the goal, it attempts to shoot directly using a head butt. If the player is placed on the kick off position facing one side, there are two options. One is to head butt the ball straight ahead to get the ball to the target end. The second option is to kick off at an angle if an opponent is ahead.

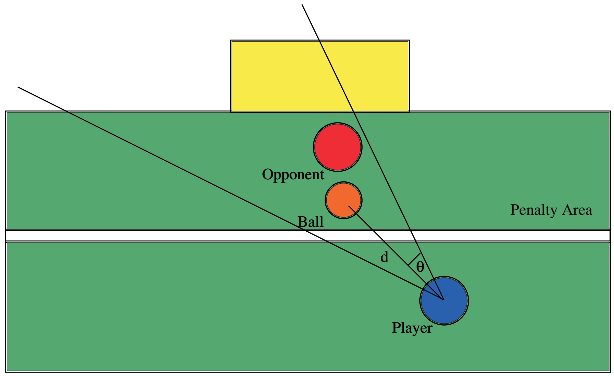

Fig. 2. If the ball is inside opponent's penalty area and the opponent robot is detected within an angle of θ from the direction of the ball, the player assumes the goalie is holding the ball and backs away from the keeper. The value of θ is inversely proportional to d, the distance to the ball.

7 The Challenge Competition

In addition to the normal soccer competition, the Sony legged league includes a set of "challenges". These are short tests of skill that may demonstrate abilities that are not necessarily used in the soccer matches.

7.1 Challenge 1

The first challenge in the 2001 competition was the "goalie" challenge. The goalie is placed on the side, at the centre line and has a maximum of 20 seconds to reach its defensive position within the penalty area. Once there, a ball is rolled down a ramp, aimed at the goal. The end of the ramp is placed at the centre of the field. The goalie must block the ball and clear it out of the penalty area.

We found that if the robot moves forward, near the penalty line, when it fully spreads its front legs to block a shot, it can cover almost the entire shooting angle. So the challenge 1 program simply moves the robot to a fixed position slightly behind the penalty line, facing directly forward. When the ball is released and approaches the robot, the spread motion is triggered. The robot does not attempt to move left or right to block the shot, nor is there any need to calculate the trajectory of the ball.

Robots that successfully completed the task were ranked according to the length of time taken for the robot to enter the penalty area. UNSW came second to the University of Pennsylvania.

7.2 Challenge 2

The second challenge tested the robot's ability to localise. Five marks were placed on the field. One robot was given the coordinates of these points and then instructed to move to each mark, in turn. At each mark, it was required to stop for three seconds and indicate that it knew it was over the mark by wagging its tail or beeping. Teams that successfully completed the task were ranked according to the total time taken.

Since we were not sure how accurately the robots would be required to locate the marks, we initially concentrated on making our localisation as accurate as possible. Later, maintaining this level of accuracy, we attempted to make the robot as fast as possible. Our strategy was to always face the robot towards the centre and pan the head slowly, to maximize the number of beacons in a frame. Once we were close to a point the speed was adjusted dynamically to allow for exact positioning. A shortest path algorithm was used to select the order of marks visited.

Once again, the UNSW robot was ranked second, this time to CMU's winning robot. Since the marks placed on the field were 8 cm in diameter, we were too conservative in our localisation and could have sacrificed some accuracy for greater speed.

7.3 Challenge 3

Challenge three was indeed challenging. None of the teams successfully completed this task. However, with some small adjustments to the starting conditions, the UNSW robot did demonstrate that the skill could be achieved. The task was to let two robots walk down the field, one on each side, passing the ball back and forth until one of the robots scores. Two unmoving "opponent" robots were placed on the middle of each half of the field as obstacles.

Our philosophy for Challenge 3 was to do it slowly but surely. This was necessary because it was a complex task, involving many stages, all requiring a high degree of reliability. The strategy involved grabbing the ball and then carrying it with the front paws in a customised walk. Once in position, we used the double-handed kick to ensure the ball made it to the other side and was positioned so that the other robot could pick it up. There was much experimentation to ensure that the robot would behave correctly. For example, the robot would back away from the edge to give it enough room to turn and align its kick. Unfortunately, in the actual challenge, a small error in localisation, as a result of the robot's vision being covered on boot up caused the first challenger to switch into its final state. The challenger did not pass to its team mate, instead going directly to the goal and shooting.

Although we did not earn any points, we demonstrated Challenge 3 again, with the robots' cameras uncovered at the start. On this occasion, both robots performed flawlessly.

On aggregate points the UNSW team also won the challenge competition.

8 Conclusion

Early in the competition, it was evident that two teams, UNSW and CMU, were clearly stronger than the others. Both teams had powerful, agile and stable locomotion. Both teams had good vision and localisation and both teams were recording high scores against their opponents. The UNSW vs CMU final was arguably the best contest seen in this league. Most games are won because of a mismatch in low-level skills. For example, one team may be faster, see the ball better, localise more accurately, etc. However, in the 2001 final, both teams had good low-level skills. The difference seemed to be more in game-play tactics.

One of the reasons for the strength of the UNSW team is due to that fact that we always test new behaviours under game-play conditions. We also run frequent practice matches that enable us to develop tactics that are more likely to work during a game. For example, several teams developed very good kicks. While these are useful to have, they are not game winners against a fast opponent that can block the ball. The UNSW team concentrated on small things like when the robot loses sight of the ball, turning in a direction that is most likely to knock the ball in the right direction. Other small behaviours dealt with minimising head movement when localising so that we do no lose sight of the ball and do not waste time looking around. While the paw-dribble and turn kicks were not "killer" shots, they tended to move the ball around and keep it away from opponents. Finally, the 2001 team improved the robot avoidance behaviours to ensure that team mates were well spread on the field, not interfering with each other and in positions where they could take advantage of loose balls.

It is clear that a great deal of hand-crafting was required to achieve championship winning performance. The real challenge is to devise means by which situation-specific behaviours can be learned. However, we believe there is still much for us humans to learn before we understand the problem well enough to be able to achieve this goal.

Acknowledgements

We thank Mike Lawther for his assistance in setting up the CVS repository, used to control the 2001 source code. Mike and Ken Nguyen also provided valuable support during the Seattle competition.

References

1. J. Dalgliesh and M. Lawther. Playing soccer with quadruped robots. Computer engineering thesis, School of Computer Science and Engineering, University of New South Wales, 1999.
2. B. Hengst, D. Ibbotson, S. B. Pham, J. Dalgliesh, M. Lawther, P. Preston, and C. Sammut. The UNSW Robocup 2000 Sony Legged League Team. In P. Stone, T. Balch, and G. Kraetzschmar, editors, *Robocup 2000: Robot Soccer World Cup IV*, pages 64–75. Springer-Verlag, 2001.
3. B. Hengst, D. Ibbotson, S. B. Pham, and C. Sammut. UNSW RoboCup 2000 Team Report. Technical report, School of Computer Science and Engineering, University of New South Wales, 2000. `http://www.cse.unsw.edu.au/~robocup`.
4. J. R. Quinlan. *C4.5: Programs for Machine Learning.* Morgan Kaufmann, San Mateo, CA, 1993.