The IITKanpur Tigers Soccer Team

Akhil Gupta, Rajat Raina, Amitabha Mukerjee and Harish Karnick
Department Of Computer Science and Engineering,
Indian Institute of Technology, Kanpur, India
{akhilg, rajatr, amit, hk}@cse.iitk.ac.in

Abstract

A champion-level robot soccer team must incorporate intelligent environment sensing and modeling, strong low-level skills such as ball interception and kicking, high level reasoning and team coordination through distributed planning. We have identified a wide range of research areas within robot soccer. Our major research thrust is in objective agent evaluation and heuristic optimization. We are also working on skill improvement, robust behaviour in uncertain environments, intelligent sensing, high level inter-player communication, opponent modeling and adaptation, and an exact and robust defensive technique.

1 Introduction

Robotic soccer presents a realistic multi-agent problem. The simulated agents must coordinate against a competitive group of adversaries, and use limited resources to achieve team goals.

We chose to work on robot soccer as it offers an open challenge, with immediate evaluation. Winning performance inevitably requires novel and practicable research. Research and methodologies that can work in the difficult domain of robot soccer carry solid proof of their applicability to other practical problems. Through our past work on the domain, we are also convinced that a large number of problems in robotics and artificial intelligence are related to those in robotic soccer.

2 Past Work

We have previously concentrated on machine learning approaches to robot soccer. [1] learns the shooting and interception skills using radial basis function networks. The interception skill is particularly interesting as it calculates the optimal intercept point based on the ball movement and also the fastest opponent location; such opponent sensitive interception is difficult to achieve exactly in analytical interception techniques. [2] learns the 2-on-2 dribbling skill using the trainer mechanism for reinforcement feedbacks. [3] looks at a human feedback driven evaluative and learning approach to simulated robot soccer. In preparing a complete team of intelligent soccer agents, we have tried to apply some of these ideas to actual game play.

3 Basic Client Architecture

Our overall approach to building a complete cooperative agent is roughly based on a layered approach [4]. We assume that the agent starts by learning low level skills such as shooting and intercepting the ball. It then builds on them to construct higher level skills, such as passing and dribbling.

Each agent maintains its own memory model, with coordinates, velocities and their confidence values for the ball and the other players. The memory is updated whenever a visual server message or an audio message from a teammate is obtained. We have adapted our low-level memory model from the publicly availabled CMUnited 99 team code, and the framework provided by the simple FCPAgent team.

The robot soccer problem must be divided among the teammates to allow good performance. At a high level, we assign roles to each agent on our team and assign home positions on the field to it. At each simulation cycle, the agent is faced with an action decision problem. The broad simplistic decision algorithm is given below.

$$
\begin{array}{l}
if\ (BallConfidenceLow()) \\
\quad Scan\ field\ for\ ball; \\
else\ if\ (isBallKickable()) \\
\quad Choose\ a\ with\ ball\ action; \\
else\ if\ (should\_I\_RunToBall()) \\
\quad Run\ to\ calculated\ interception\ point; \\
else\ if\ (BallWithOurTeam()) \\
\quad Perform\ offensive\ positioning; \\
else\ if\ (BallWithOpponentTeam()) \\
\quad Perform\ defensive\ positioning; \\
else\ Choose\ default\ action\ based\ on\ current\ role;
\end{array}
$$

We describe the individual research issues involved in greater detail below.

4 With-ball skills

We have currently used a heuristic approach to developing with-ball skills. An agent with the ball performs a heuristic search for the best action available in the next cycle. The types of actions considered are passing, dribbling, shooting, holding and looking. Thus, for example, several possible passes are evaluated according to analytic and tactical measures. In addition, each evaluated action is also assigned a confidence value based on the expected accuracy of the utility assignment.

In an inaccurately sensed environment, agent reasoning must consider not only the utility but also the current game state and the confidence in that action's success. Thus, it may be preferable to try only high confidence actions in defensive positions. We partially remedy this problem by providing an *expectation-based sensing* approach - a sensing direction may be chosen according to an apparent high utility but low confidence action. The agent can also reason about appropriate body and neck turns required to sense the environment in that direction. To improve the quality of soccer play, the higher utility actions are searched more finely. This allows greater number of forward passes to be discovered, and an inventive style of soccer play. Because of our evaluate-sense-evaluate cycle, we have consistently recorded high team possession against good teams. We also believe that such sensing is invaluable since the audio message bandwidth is reduced and low-level field information can no longer be exchanged through audio.

Besides the above, we are working on a *dodge* skill, where the with-ball agent moves the ball from one side to other in the kickable area in such a way as to confuse the defensive opponent. This would indeed be an interesting form of opponent skill anticipation - the anticipation of its expected weaknesses.

5 Sensing and dealing with uncertainty

Obtaining an accurate world model is crucial for good quality play. The sensing strategy in our agents is driven primarily by expectations of high reward or high information content. Defenders position themselves strategically so that they can watch the ball and the offensive opponent simultaneously.

The movements of the goalie are especially interesting in this regard. It is crucial for the goalie to almost always look at the ball, and its positioning must be very accurate to save the opponent from scoring. We propose to model the goalie's motion as a motion planning problem in which a robot (the goalie) must move from one point to another (under certain constraints) while keeping another moving target always in field of view. We expect to derive a combination of forward and backward dashes to optimize the task.

6 Without ball positioning

Most current soccer agents tackle interception as an opponent independent problem. However, it is clear (as in [1]) that optimal interception depends not just on the ball movement, but also on the fastest opponent. The agent must be able to reason about trying risky interceptions, or about getting into a defensive position instead of dashing towards the ball. The problem is complicated by the inaccurate sensing capabilities; however, we hope that such an intelligent interception will improve the quality of low-level soccer play drastically.

Robotic soccer scorelines are increasingly tending towards real soccer - most matches between good teams end up as low scoring games. For this reason, we adopt an exact and robust approach to defensive positioning. Our defensive system is based on the method of global optimization from local perspectives, similar to that of Tsinghuaeolus [5]. Here, each defender tries to optimize the global defensive utility for its team by pairing the defenders with defensive tasks. However, a direct application of this idea leads to poor behavior and defensive blunders, as the approach is very sensitive to the accuracy of the world model. Thus, we have modified the strategy to ensure that opponents with the highest perceived danger values are guarded first, and then the remaining defensive tasks are algorithmically paired with the defenders. The advantage of the approach is that it automatically fits to the opponent style of play; in a way, it optimizes the team's defensive behavior.

Offensive positioning is slightly less difficult, mainly because defensive positioning must be very exact, and because the teammate's algorithms are all known before a match. We aim to develop set plays through locker room agreements, especially for stop-ball situations. These set plays essentially consist of pre-coordinated actions and improve the overall cooperation among the agents.

7 High-level communication

With the reduction in audio bandwidth, it is no longer feasible to use audio for exchanging low level field information. Thus, we use audio only for critical and high-information messages. We use high level primitives such as *receive-pass* and *player-close-behind* for information exchange. We also use audio to maintain reliable game state parameters such as *team-in-possession* and *last-team-to-kick*, which are used heavily by the other modules of the agent (eg the team in possession parameter decides whether a without-ball agent should position defensively or offensively, last team to kick is used by the goalie to decide if it can catch the ball).

An agent that believes it has sufficient information to convey sends an audio message only if it perceives its audio message to be of comparable importance as its teammates. For this purpose, the information is evaluated based on whether it contains a critical primitive or a reliably observed change in a game parameter (such as team in possession), and on the relatively importance of the part of the field accurately sensed by that agent.

8 Agent Evaluation and Specification

8.1 Agent evaluation

Objective evaluation of an agent (or of the team as a whole) is an extremely desirable tool. The most obvious evaluation mechanism of looking at game scorelines is too naive as it evaluates only a small subset of the agent's skill set. We have looked at extracting more interesting measures such as team possession directly through the perceived memory model of the players. The reliability of the measures is ensured by the fact that these game state parameters are reliably maintained by the visual and audio messages received in each cycle.

For example, to calibrate audio usage, we have looked at the logfiles to calculate the fraction of simulation cycles in which an audio message could have been placed without colliding with another audio message in another cycle. We have also tried to evaluate the fraction of critical (as perceived by the sender) audio messages that were actually transmitted without collisions. There is an obvious tradeoff between the above two characteristics - a simple flooding would use all the bandwidth but would block out most critical messages, and a very selective send mechanism would use only a fraction of the available bandwidth.

One of our important research goals is to derive similar evaluative measures for other skills - to look at successful passes and mispasses, shots at goal, defensive blunders etc.

8.2 Declarative specification of heuristics

Successful soccer agents often require a large body of soccer knowledge to be supplied explicitly as heuristics. Often, the heuristics are buried deep inside the agent logic, and it is difficult to change the behavior during testing, or during the game. The ability to accurately specify behavior would be desirable; it would allow independent testing of several competing heuristics and would be amenable to change even during a game by simply changing the set of heuristics to use.

For such a heuristic scheme, declarative specification of heuristics is a clear way out. Certain teams such as Robolog have already considered a declarative scheme, and we propose to take it a step further by optimizing the heuristics for best quality play.

8.3 Optimization of heuristics

We believe that once an evaluative measure is available for a skill S, the skill learning for S can be largely automated. Indeed, the declarative heuristics used in the agent logic can be automatically varied by the agent during its learning exploration phase. Each combination of heuristics can be evaluated according to the appropriate evaluation measure, and the best evaluation can be chosen. This technique offers the best of both worlds - the intuitiveness of heuristics and the exactness of a machine learning approach.

We further aim to link previous work on groundtruth models of evaluation [3]. Human experts are polled and asked to suggest best continuations for agents from given sample field positions. The feedback is considered to represent the optimal strategy for an agent from a given position. We believe that the evaluative benefits of such feedback can be immense; tactical and strategic benefits are often difficult to capture in automatic evaluations. Such feedback can thus be used to evaluate proposed heuristics in order to find the optimal solution.

9 Results and Future Work

Our team has demonstrated high quality play with the ball, often maintaining ball possession against good teams. The heuristic approach has its drawbacks - the agent plays in the way it has been taught to; there is a human-like pattern to its play and it is not necessarily suited for software soccer agents. Our novel ideas of intelligent interception, agent evaluation and heuristic optimization promise improved skills and better play. They bring in a machine learning element which allows an agent to optimize its own play. Coordination through high level communication, set plays and without ball positioning also lead towards a more cooperative style of play.

References

[1] Sreangsu Acharya, *Learning Radial Basis Function Based Soccer Strategies using Ideal Opponent Model*, M.Tech. thesis, Indian Institute of Technology Kanpur, 2000.

[2] Akhil Gupta, Rajat Raina and Amitabha Mukherjee, *Towards a "Perfect" Dribbling Strategy for Robot Soccer*, International Conference on Computational Intelligence, Robotics and Autonomous Systems, Singapore, 2001.

[3] Tamhant Jain, Sreangsu Acharyya and Amitabha Mukerjee, *Designing Reactive Behaviours for Soccer Agents using Groundtruth Models*, International Symposium on Autonomous Robots and Agents (ISARA 2000), Singapore, 2000

[4] Peter Stone, *Layered Learning in Multi-Agent Systems*, Carnegie Mellon University, 1998.

[5] Cai Yunpeng, Chen Jiang, Yao Jinyi and Li Shi, *Global Planning from Local Eyeshot: An Implementation of Observation-based Plan Coordination in Robocup Simulation Games*, Description paper for Tsinghuaeolus Robocup Champion 2001.