Team Yarra: holistic soccer

Andrew Jennings, Chirag Shah, Harit Jani, Maulik Upadhya and Murtuza Kapadia

RMIT University, Melbourne, Australia

ajennings@rmit.edu.au

Abstract. We describe the objectives and principles of Team Yarra, a new simulator league team. Our approach explores the use of a more strategic play selection mechanism. Using a game state representation based on an existing method used for human soccer games, we also explore the possibility of tuning for improvement of play selection.

1 Introduction

Team Yarra is a new team formed in 2002 with participants all relative newcomers to the simulation league. As we are still in the process of creating our team, many of the descriptions here will change as we progress in development. As do all beginners, we attempt to study the best of the teams in past competitions and to learn from them. We are building our team upon the excellent foundation provided by the UVA TriLearn team [7,9].

In this paper we present our approach under three main headings. Firstly we look at the representation of the game state, and how we can manipulate that representation. We present a new approach that attempts to give more weighting to the total game state, and to avoid exclusive focus on the player with ball possession. Next we consider approaches to cooperation, and consider alternatives to the prevailing model of global scripting of team behaviour. Finally we consider the prospects for learning, and look forward to a higher level of adaptability in strategic behaviour.

2 Representation and Reasoning

There are several natural representations of the soccer game state. We can firstly take a literal representation of the location and velocity of each player in the game as our starting point. This detailed view can form the basis of decision making, but without some further abstraction it may not form a solid basis. As Stone [1] indicates, there is a necessity for considering the right level of abstraction for both learning and decision making. It is not appropriate to consider both detailed dribbling and strategic movement of players at the same time.

If we wish to consider strategic behaviour then we must firstly adopt the view that perhaps the most important consideration is "..what to do when you don't have possession of the ball" [2]. How can we guide the players who do not have possession? Here we adopt an approach derived from observation of human soccer games [3] that works on the basis of "possession" of regions on the field.

If we consider the behaviour of a soccer player, based on velocity and heading there is a space of the field which that player can reach within a set time. Figure 1 illustrates a typical elliptic "capture region".

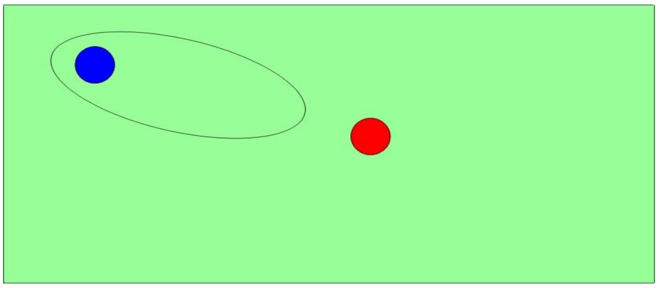

Figure 1. Regions of capture for a player.

When a human player is in full forward flight she cannot quickly move sideways, so it is much more feasible to capture parts of the field in the heading direction. It is more difficult to turn suddenly and capture territory. This introduces the concept of "ownership of space" on the field. Of course we could resort to extreme ownership of space surrounding the defensive goal, for example, but this would leave us with little opportunity to score and still vulnerable to a sudden attack.

How might we characterise the whole field with player and ball positions? This issue is important for deriving strategies: a compact representation that captures the essential aspects of the game state is of great advantage. Figure 2 illustrates the division of the field into regions controlled by each team.

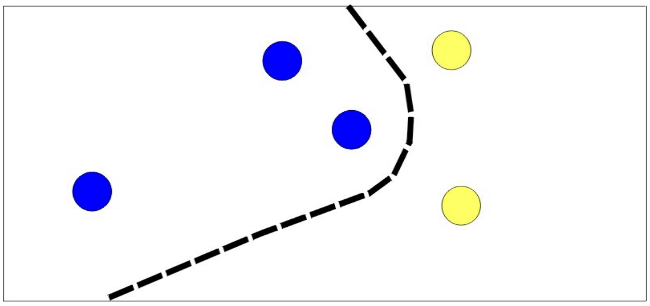

Figure 2. Voronoi diagram differentiation of regions of team ownership.

Here the dark team has control of the territory at the left. The team controls that territory to which it can deliver a player in advance of the opposition. The diagram of control corresponds to a Voronoi diagram [4] generated on a team basis. For every point on the field we determine which player is closest, and label that point with the team corresponding to the player.

The process of generating the Voronoi diagram for each game state is not difficult computationally, but can be further enhanced by considering issues of resolution. How accurately do we need to know the region state to guide strategy? Figure 3 illustrates a coarse discretisation of the field that may still be useful as a guide to strategy. We are experimenting with the choice of the correct resolution for optimal game play.

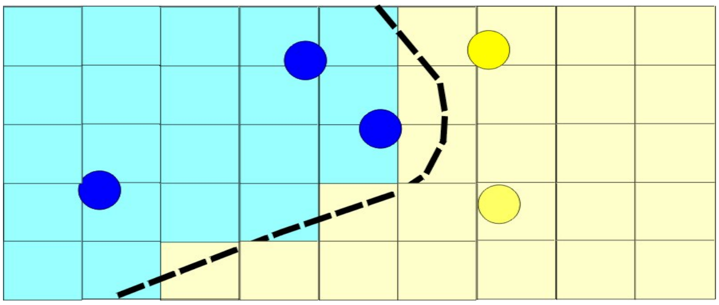

Figure 3. A discretisation of the regions of ownership.

We could characterise this approach as holistic in nature: an attempt to look at the whole field position when generating player moves. So instead of being concerned only with the direct movement of the ball, we attempt to deal with the state of the field in a more global manner. We attempt to move beyond the point where "In most circumstances, backward passes are not considered." [5]. When we consider how often backward passes are a feature of human soccer, it is perhaps a direction worth exploring.

Even with this simplification to team owned regions, we still face an enormous number of possible game states. We can handcraft procedures to deal with common situations, but this immediately constrains us in modes of development. If we take as our goal to attempt to at least automatically tune behaviour to match other teams then we need a method of proceeding from game-states to player actions that at least allows the possibility of tuning.

We take the team ownership of regions as input to the derivation of game states, and then use clustering to organise the states. Once again this includes description of the whole field state through region ownership. This process potentially has a very large number of states, and requires careful organisation to be successful. This process of utilising self-organising maps was also followed by Wunstel et al [6] in the real robot leagues. We employ fuzzy clustering to organise the game states. An important decision is the choice of *metric* in comparing game states. Together with the choice of discretisation this determines the style of play and level of detail considered in the choice of strategy.

Of course there are disadvantages in such an unsupervised approach to game state classification. We collect many states that are incidental to the progress of the play. But it allows us to classify without assumptions about the progress of play, and provides the capability to deal with any style of play both for our team and for opponents.

Figure 4 summarises the process of construction, and the subsequent mode of game play. We use the region ownership, estimated player positions and velocities together with ball position as the game state vector. The game-state acts as a selector for plays. These plays make use of lower level skills are as inherited from other teams [7,8]. At present much of the play determination is hand coded, but we would like to at least provide for the possibility of tunable behaviours in the future. One interesting question is whether this semi-automatic classification of game states can become superior to hard-coded selection of game states.

Early experience with this approach has already resulted in some important changes. When there is crowding of players around the ball it is not feasible or useful to investigate the detailed regional ownership. Instead we simply count the number of players from each team at the ball. The team with the greater number of players is then said to have ownership of the region containing the ball. We have also had to give greater priority to the goal area: obviously ownership of field regions near the goal is of much greater importance than ownership of broad corridors on the field. For the moment we are still hand-crafting play selection, with the objective of automatic selection still to be explored.

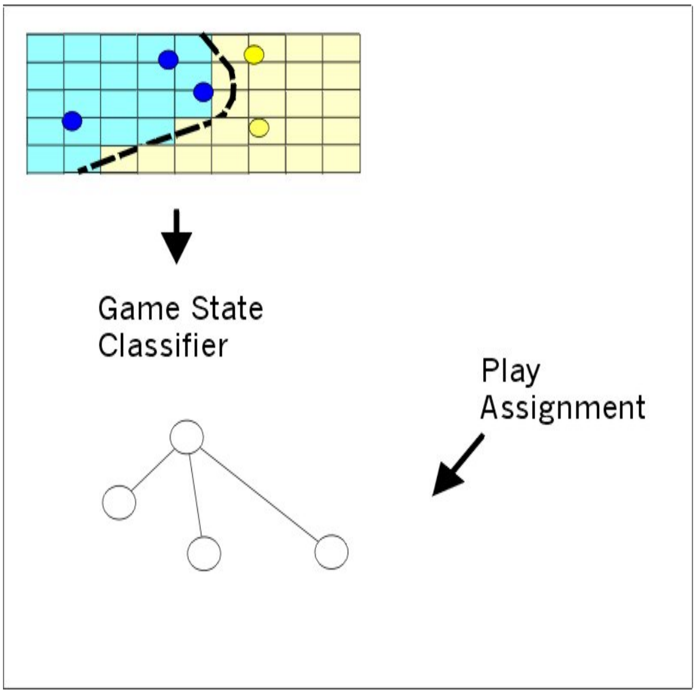

Figure 4 Process of strategy selector construction.

3. Cooperation

It is interesting that much of the motivation for multi-agent research as a research field was to explore methods of cooperation in difficult environments. However if we look at the approaches of the most successful teams (eg. [9,10]) it is clear that there are quite strong constraints on what is possible in the way of agent communication. When we need to act in real time, there is rarely time for an extended exchange of messages. So instead it is imperative to rely on a shared set of assumptions about the environment. Given that the objective is in many cases to create a strong team, there is little incentive to explore issues that will not occur during game play

To make progress we have adopted the prevailing view that we can *infer* player movements from the game state, without exchanging messages. But in the course of development we are exploring the possibility of varying the method of cooperation in the course of play. We are exploring whether it is of advantage to modify the style of cooperation on the basis of the game state. In some situations a small change in a single player position can produce a dramatic change in the strategic picture. For example in Figure 5 the final player movement seals the side of the field and would indicate a switch across field as a desirable movement.

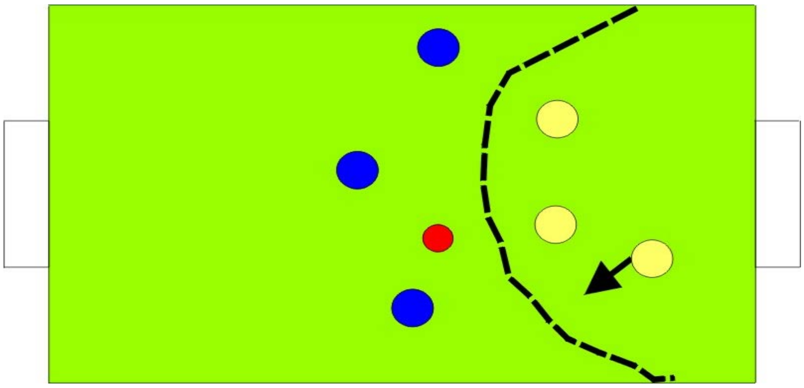

Figure 5. Switch indication from a region closure.

3. Learning

Unfortunately there is a bias towards learning as a superior method of construction, which perhaps owes much to misconceptions about the human approach [11]. We should aim to use learning where it is a feasible method of achieving our objectives, not as a substitute for other methods. In the history of computer chess [12] it was only when grandmaster play became possible that tuning of play became possible. Our long term aim is to have the capability to adapt our strategy based on a series of games played against a better opponent.

The first obstacle to learning is the sheer volume of information in a game trace. Stone [1] shows the necessity for levels of abstraction in learning: it is neither possible nor desirable to learn to dribble and to learn strategic plays at the same time. Our focus is solely on learning at the strategic level. The world model representation we have described above may have advantages for learning. In the example of Figure 5 a single player movement has significant impact on region ownership: so small movements can sometimes be highly significant. On the other hand, broad movements in the outer field may have limited impact on region ownership. If this representation is of value, then it will give higher weighting to player moves that have a greater impact on the conduct of the game.

In our framework based on field ownership we aim first to tune the selection of plays from game-states. The choice of a particular play is based on the degree of match to a game-state in the classifier of states. We attempt to tune the selection based on success of play. At present we can only tune in set plays, but we plan to extend the learning mechanism to provide for more open play.

Of course the principal difficulty here is the *credit-assignment problem.* A successful game may result from many moves that are incidental to the success, and similarly many poor moves in the game may not be desirable features to preserve. We attempt to follow a similar path to chess evaluation tuning, through comparison of moves with expert human moves. In the simplest case, if a move results in an improvement in region ownership then it can be judged successful. This is a very difficult process, and we are only at the beginning of this journey. One important early test for our approach is to examine whether our representation is of advantage for automatic strategy selection.

Summary

We can characterise the goals of Team Yarra as pursuing two objectives. Firstly we attempt to create a new description of the game state that lends itself to automatic tuning of behaviour. We are not so ambitious that we attempt to synthesise new behaviours on the basis of game state exploration, but we regard the near term objective of automatic tuning as important. Secondly, through this representation we aim to play a game that plays more attention to the global state of the field of play, that is more strategic in nature. It remains to be seen whether these directions result in useful new possibilities.

References

1. Peter Stone and Manuela Veloso "Layered Learning" Technical Report, Carnegie-Mellon University, 1999
2. Manuela Veloso, Post-Competition Review, Robocup98.
3. Tsuyoshi Taki, Jun-ichi Hasegawa & Teruo Fukuruma "Development of a Motion Analysis System for Quantitative Evaluation of Teamwork in Soccer Games" *ICIP International Conference on Image Processing*, pp. 815-818, Lausanne 1996
4. H.Edelsbrunner *Algorithms in Computational Geometry* Springer 1987
5. Luis Paulo Reis & Nuno Lai "FC Portugal Team Description: Robocup 2000 Simulation Champion" *Robocup 2000*, Springer Verlag 2001
6. Michael Wunstel, Daniel Polani, Thomas Uthman & Jurgen Perl "Behaviour Classification with Self-Organising Maps" *Fourth International Robocup Workshop*, pp. 179-188, Melbourne 2000.

7. Jelle Kok, Remco de Boer, Nikos Vlassis & Frans Groen "UvA TriLearn 2002 Team Description" *Robocup International Symposium*, Fukuoka 2002
8. Yao Jinyi and Cai Yunpeng "Team Description of TsinghuAeolus" *Robocup International Symposium*, Seattle 2001.
9. Jelle R. Kok and Nikos Vlassis "Mutual Modelling of Teammate Behaviour" *Intelligent Autonomous Systems Technical Report*, 2002
10. Cai Yunpeng, Chen Jiang, Yao Jinyi & Li Shi "Global Planning from Local Eyeshot: An Implementation of Observation Based Plan Coordination in Robocup Simulation Games" *Robocup International Symposium*, Fukuoka 1992
11. Steven Pinker *The Blank Slate: The Modern Denial of Human Nature* Viking, NY 2002
12. Feng-hsiung Hsu, Thomas Anantharman, Murray Campbell and Andreas Nowatzyk "A Grandmaster Chess Machine" *Scientific American*, Vol. 263, pp. 18-23, 1990