The UT Austin Villa 2004 Simulator Coach

Gregory Kuhlmann and Peter Stone

Department of Computer Sciences
The University of Texas at Austin
Austin, Texas 78712-1188
{kuhlmann,pstone}@cs.utexas.edu
http://www.cs.utexas.edu/~{kuhlmann,pstone}

Abstract. The UT Austin Villa 2004 simulated online soccer coach is based on our previous year's champion entry. The main research focus continues to be placed on treating advice-giving as a machine learning problem. In this paper, we review the multi-faceted learning strategy that our coach uses and discuss this year's innovations.

1 Background

The UT Austin Villa[1] 2004 coach is based on our previous year's entry [1], and operates as follows. Prior to a match, the coach examines the provided logfiles of games played by the *fixed opponent*. The coach collects data about players on the fixed opponent team as well as the players on the team the fixed opponent is playing against.

The data collected during logfile analysis are fed into a group of learning algorithms that generate player models for both teams. The models are then used to produce three different kinds of advice: formational, offensive, and defensive. The learned advice is combined with a few hand-coded rules and sent to the coachable team at the beginning of the match.

While the coach is best able to reason about players in terms of their roles, CLang requires players to be specified by their uniform numbers. For this reason, the coach maintains a mapping between roles and uniform numbers for each player on both teams. If during the course of the game players change roles, the affected rules are sent again with the updated player numbers.

2 Learning

For both offensive and defensive advice, the product of our learning algorithm is a classifier that is able predict the next high-level event to occur, given the current state of the game. To encode the simulator's state, we used a large set of features including the ball location, player locations, and relative distances.

The game analysis module outputs instances of these variables labeled with the actual action taken by the modeled team in that state. These examples are used to train a series of decision trees, one for each modeled player. We used the

[1] http:/www.cs.utexas.edu/~AustinVilla

J48 decision tree algorithm implemented in the Weka machine learning software package [2]. Because the structure of a decision tree is easily understandable, it is fairly straightforward to convert a tree into CLang advice.

For offensive advice, the coach builds a classifier for each player that tries to predict what that player will do with the ball in any given situation. Once we have trained a decision tree for a given player, we can convert it into advice to be given to our own player.

To generate defensive advice, we model the behavior of the opponent and attempt to foil its predicted strategy. For defensive advice, we predict only passes. Because we are interested in predicting a pass before it is made, we record the 10 cycles prior to when the ball is kicked. We then use a heuristic model to convert the learned predictions regarding opponent behaviors to defensive actions that can *prevent* that action. For instance, to prevent a pass, it is a good idea to position a defender along a passing lane closer to the intended receiver than to the passer.

Finally, our approach to learning a team formation was similar to our approach to learning offensive advice. The coach observes a team that can beat the opponent and then attempts to mimic that team's behavior. We model the formation as a home position and ball attraction vector for each player. The position values are calculated as the average position of the observed player during the course of the game and the ball attraction values were handpicked through brief experimentation.

3 Innovations

In examining the defensive advice more closely, we have decomposed the problem into the subtasks of first predicting the future actions of the opponent, and second, reacting to these predictions. In the latter case, we tested the team's ability to improve its performance given prescient knowledge of the opponents' intended actions. While these are both difficult problems, we plan to make strides toward each in our 2004 coach.

In our research we are continuing to enhance the learned defensive and offensive advice, and we plan to increase the degree to which each type of learned advice is opponent-specific.

Meanwhile, because of the apparent importance of formational advice, we continue working on finding ways to learn better, more adaptive formations. Our previous technique of mimicking a successful formation was certainly effective. Our 2004 coach improves upon this technique by adapting the formation to the opponent online.

In addition, we explore various methods for generating set play advice. One such approach involves simply selecting from a library of set plays the one with the highest probability of success. We hope to develop this technique to the point of being adaptive online.

Overall, our focus this year is on improving upon the learning method of our previous coach by making it more accurate, more adaptive and more flexible.

Acknowledgments

This research is supported in part by NSF CAREER award IIS-0237699.

References

1. Gregory Kuhlmann, Peter Stone, and Justin Lallinger. The champion UT Austin Villa 2003 simulator online coach team. In Daniel Polani, Brett Browning, Andrea Bonarini, and Kazuo Yoshida, editors, *RoboCup-2003: Robot Soccer World Cup VII*. Springer Verlag, Berlin, 2004. To appear.
2. Ian H. Witten and Eibe Frank. *Data Mining: Practical Machine Learning Tools and Techniques with Java Implementations*. Morgan Kaufmann, October 1999. `http://www.cs.waikato.ac.nz/ml/weka/`.