UT Austin Villa 2D Simulation Soccer Team 2013

Patrick MacAlpine, Will Daspit, Kathryn Spradlin,
Chao Zhang and Peter Stone

Department of Computer Science
The University of Texas at Austin
Austin, TX 78712-0233
{patmac, googly, spradlin, czhang, pstone}@cs.utexas.edu

Abstract. This paper describes the research focus and ideas incorporated in the UT Austin Villa 2D Simulation Soccer team entering the RoboCup competitions in 2013.

1 Introduction

In this paper, we describe the agent our team UT Austin Villa is currently developing for participation at the 2013 RoboCup 2D Simulation Soccer competition. While Peter Stone, the advising professor for the UT Austin Villa team, competed in and won the 2D Simulation League during the first few years of RoboCup [13, 11], this is the first time that the UT Austin Villa team will be competing in the 2D Simulation League. We hope to leverage our team's previous experience of competing in other RoboCup leagues, having won both the Standard Platform League [2] and the 3D Simulation League [10] at last year's RoboCup competition in Mexico, in our preparation for this year's competition.

The rest of this paper is organized as follows. Section 2 provides a brief overview of the 2D soccer simulator. In Section 3 we describe the design of the UT Austin Villa agent. We present our current research directions in Section 4, and Section 5 concludes.

2 Brief Overview of 2D Simulation Soccer

As one of the oldest RoboCup leagues, 2D simulation soccer has been well explored, both in competition and in research. The domain consists of two teams of eleven autonomous agents playing soccer against each other on a simulated 2D soccer field. The agents communicate with a central server that controls the game. The agents receive sensory information about the game, including the position of the ball and other agents, from the central server. After processing this information, the agents then communicate to the server what actions they want to take in the world such as dashing, kicking, and turning. 2D soccer abstracts away many of the low-level behaviors required for humanoid robot soccer, including walking and computer vision, and thus affords the chance to focus on higher-level aspects of playing soccer such as multiagent coordination and strategic play.

3 Agent Architecture

We extend the 2010 HELIOS codebase. HELIOS performs well on both offense and defense by using a set of formations dependent on game state and ball location, and additionally uses a short-term lookahead combined with a field evaluator to predict and consider the results of its actions. Both dribbling and passing iteratively consider a large number of possible angles and speeds but filter out the obviously bad ones. This allows the agent to make (within the limitations of the prediction and evaluation framework) nearly optimal action choices, and dribble and pass as quickly as possible while still maintaining control. In practice, this results in players passing the ball up the midfield, then dribbling along the sidelines of the field, completing the play with short passes and a shot from extremely close to the goal.

HELIOS also includes a graphical formations editor that lets a user naturally specify player positions based on ball positions, letting HELIOS-based teams experiment with different configurations easily. The formations consist of a number of ball positions with corresponding player position that are interpolated between to determine appropriate player positions for the ball at any position. Interpolated positions are calculated using one of a variety of different methods including Delaunay Triangulation [1].

4 Research Directions

We are planning to incorporate dynamic role assignment and positioning, based off our last year's world champion 3D Soccer Simulation League team [8]. By porting this behavior, we can begin to build a common code base across our simulation league teams.

We are also working on a more flexible formation system, based off of HELIOS-style formations but allowing for variations to respond to opponent positions, possibly blocking off pass routes or maneuvering around defenders. Furthermore, because passing is so integral for keeping the ball away from the opposing team, we would like to create a network of agents that is consistently available to receive passes.

One such idea that warrants further exploration involves the use of allied agents to cover opposing agents so as to reduce the opponents' efficacy in passing and receiving the ball. The idea is that each allied agent guards a certain region of the field; when an opposing agent enters that region, the ally guarding that region will "mark" and remain in close proximity to the intruding opponent agent, harassing and disrupting potential passes to and from that player. This coverage would continue until the opponent agent leaves the guarded region, with the process repeating for each guarded region that the opponent enters. With enough cleverness, this tactic could also be adapted to create a network of passable agents as well. So long as our agents can remain in front of at least one opposing agent within its guarded region, we hope to show that it will be possible to create a network of passable agents under most circumstances.

In addition, we are testing reinforcement learning techniques applied to various parts of the ball holder's offensive behavior. We learn to predict a score function after an action, related to the forward movement of the ball over a fixed number of timesteps after the action and an approximated utility of terminal states. To do this, we consider a simplified action space such as passing to one of a generated list of reasonable targets or dribbling in a generated direction, and calculate some parameters relevant to each action. We then cluster over the parameter space for each action and build linear models within each cluster. When an agent receives the ball, it chooses a random action with low probability for training, and otherwise selects an action according to the average of the values the linear models predict at the seen parameters, weighted by the inverse of the distance to the linear models' cluster centers.

5 Summary and Conclusion

This paper has presented a high-level view of the architecture, design, and research directions of the UT Austin Villa team being prepared for entry in the RoboCup 2013 2D Simulation Soccer League competition.

The simulation of soccer playing agents opens up interesting problems for multiagent systems, optimization, machine learning, and AI. While initially our main emphasis has been on getting familiar with the 2D Simulation Soccer League domain, we are now pushing forward and working on the ideas presented in Section 4. There are numerous areas of research in the 2D Simulation Soccer League we hope to explore; these provide the inspiration and driving force behind UT Austin Villa's desire to participate in this league.

UT Austin Villa has been involved in the past in several research efforts involving RoboCup domains. Kohl and Stone [7] used policy gradient techniques to optimize the gait of an Aibo robot (4-legged league) for speed. Stone *et al.* [12] introduced Keepaway, a subtask in 2D simulation soccer [3, 6], as a test-bed for reinforcement learning, which has subsequently been researched extensively by others (for example, Taylor and Stone [14], Kalyanakrishnan *et al.* [5], and Taylor *et al.* [15]). Most recently the team has used the 3D simulation domain to explore learning walks for bipedal locomotion (MacAlpine *et al.* [9] and Farchy *et al.* [4]). We are keen to continue our research initiative in the 2D Simulation Soccer League.

References

1. H. Akiyama and I. Noda. Multi-agent positioning mechanism in the dynamic environment. In U. Visser, F. Ribeiro, T. Ohashi, and F. Dellaert, editors, *RoboCup 2007: Robot Soccer World Cup XI*, volume 5001 of *Lecture Notes in Computer Science*, pages 377–384. Springer Berlin Heidelberg, 2008.
2. S. Barrett, K. Genter, Y. He, T. Hester, P. Khandelwal, J. Menashe, and P. Stone. Ut austin villa 2012: Standard platform league world champions. In X. Chen, P. Stone, L. E. Sucar, and T. V. der Zant, editors, *RoboCup-2012: Robot Soccer World Cup XVI*, Lecture Notes in Artificial Intelligence. Springer Verlag, Berlin, 2013.
3. M. Chen, E. Foroughi, F. Heintz, Z. Huang, S. Kapetanakis, K. Kostiadis, J. Kummeneje, I. Noda, O. Obst, P. Riley, T. Steffens, Y. Wang, and X. Yin. Users manual: RoboCup soccer server — for soccer server version 7.07 and later. *The RoboCup Federation*, August 2002.
4. A. Farchy, S. Barrett, P. MacAlpine, and P. Stone. Humanoid robots learning to walk faster: From the real world to simulation and back. In *Proc. of 12th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2013.
5. S. Kalyanakrishnan, Y. Liu, and P. Stone. Half field offense in RoboCup soccer: A multiagent reinforcement learning case study. *Proceedings of the RoboCup International Symposium 2006*, June 2006.
6. H. Kitano, M. Asada, Y. Kuniyoshi, I. Noda, E. Osawa, and H. Matsubara. RoboCup: A challenge problem for AI. *AI Magazine*, 18(1):73–85, 1997.
7. N. Kohl and P. Stone. Policy gradient reinforcement learning for fast quadrupedal locomotion. In *Proceedings of the IEEE International Conference on Robotics and Automation*, May 2004.
8. P. MacAlpine, F. Barrera, and P. Stone. Positioning to win: A dynamic role assignment and formation positioning system. In X. Chen, P. Stone, L. E. Sucar, and T. V. der Zant, editors, *RoboCup-2012: Robot Soccer World Cup XVI*, Lecture Notes in Artificial Intelligence. Springer Verlag, Berlin, 2013.
9. P. MacAlpine, S. Barrett, D. Urieli, V. Vu, and P. Stone. Design and optimization of an omnidirectional humanoid walk: A winning approach at the RoboCup 2011 3D simulation competition. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI-12)*, July 2012.
10. P. MacAlpine, N. Collins, A. Lopez-Mobilia, and P. Stone. UT Austin Villa: RoboCup 2012 3D simulation league champion. In X. Chen, P. Stone, L. E. Sucar, and T. V. der Zant, editors, *RoboCup-2012: Robot Soccer World Cup XVI*, Lecture Notes in Artificial Intelligence. Springer Verlag, Berlin, 2013.
11. P. Stone, P. Riley, and M. Veloso. The CMUnited-99 champion simulator team. In M. Veloso, E. Pagello, and H. Kitano, editors, *RoboCup-99: Robot Soccer World Cup III*, volume 1856 of *Lecture Notes in Artificial Intelligence*, pages 35–48. Springer Verlag, Berlin, 2000.
12. P. Stone, R. S. Sutton, and G. Kuhlmann. Reinforcement learning for RoboCup-soccer keepaway. *Adaptive Behavior*, 13(3):165–188, 2005.
13. P. Stone, M. Veloso, and P. Riley. The CMUnited-98 champion simulator team. In M. Asada and H. Kitano, editors, *RoboCup-98: Robot Soccer World Cup II*, volume 1604 of *Lecture Notes in Artificial Intelligence*, pages 61–76. Springer Verlag, 1999.
14. M. E. Taylor and P. Stone. Behavior transfer for value-function-based reinforcement learning. pages 53–59, July 2005.

15. M. E. Taylor, S. Whiteson, and P. Stone. Comparing evolutionary and temporal difference methods for reinforcement learning. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 1321–28, July 2006.