The UT Austin Villa 2007 Simulation 3d Team

Shivaram Kalyanakrishnan and Peter Stone

Department of Computer Sciences
The University of Texas at Austin
1 University Station C0500
Austin, Texas 78712-1188
{shivaram,pstone}@cs.utexas.edu
http://www.cs.utexas.edu/~AustinVilla

Abstract
This document serves as UT Austin Villa's application to participate in the RoboCup 2007 3D Simulation League.

1 Introduction

In this document, please find

- List of RoboCup related publications (not restricted to RoboCup Symposia).
- List of previous results.
- List of Work for the simulation community (Simulator work, committees, etc.)
- Research Proposal

2 Previous Results

- **Best Student Paper Award**, RoboCup Symposium, June 2006.
- **Best Paper Award Nominee**, RoboCup Symposium, June 2006.
- **World Champion** team member in 6 **RoboCup** events:
  - Simulator coach competition, July 2005, July 2003;
  - Simulator competition, August 1999;
  - Simulator and small-size robot competitions, July 1998
  - Small-size robot competition, August 1997.
- 2nd place in simulator coach competition at RoboCup 2006 (June 2006)
- 3rd place in simulator competition at RoboCup 2001 (August 2001)
- 2nd place at american open simulator coach competition (May 2003)
- 3rd place in american open legged robot competition (May 2004, 2005)
- quarterfinalist in legged robot competition at RoboCup 2004, 2005 (July)

3 Simulation Community Service

Peter Stone has been involved in extensive RoboCup service as follows:

- **Member**, RoboCup simulation league organizing committe, July 2007.
- **Chair**, RoboCup US Open simulation league committee, May 2005, April 2004.
- **Associate chair** in charge of simulation events for RoboCup-2001, August 2001.
- **Co-chair**, RoboCup Workshop, August 2000.
- **Co-chair**, RoboCup simulator competition organizing committee, August 1997 – August 1999.
- **Trustee**, RoboCup Federation, July 2003 – present.
- **Advisory committee member**, RoboCup European Championship, May 2000.
- **Executive committee member**, RoboCup Federation, August 1999 – present.
- **RoboCup Symposium PC member**, July 2005, July 2004, July 2003, June 2002, August 2001.
- **Member**, RoboCup US committee, March 1997 – present.

4 RoboCup Related Publications

Peter Stone has been involved in the RoboCup initiative since its inception and has been publishing his research using this domain for the past eight years. The relevant books and journal articles are listed in this section. More than fifty additional relevant book chapters, magazine articles, conference papers, and workshop papers can be found referenced on the principal investigator's CV which is linked to `http://www.cs.utexas.edu/~pstone`. Most of the papers themselves are also available from that page.

Shivaram Kalyanakrishnan has published two RoboCup-related papers with Peter Stone, including the winner of the best student paper award at RoboCup 2006:

[1] **Shivaram Kalyanakrishnan** and Peter Stone. Batch Reinforcement Learning in a Complex Domain. In 6th International Conference on on Autonomous Agents and Multiagent Systems (AAMAS), May 2007. To appear.

[2] **Shivaram Kalyanakrishnan**, Yaxin Liu, and Peter Stone. Half field offense in RoboCup soccer: A multiagent reinforcement learning case study. In Gerhard Lakemeyer, Elizabeth Sklar, Domenico Sorenti, and Tomoichi Takahashi, editors, *RoboCup-2006: Robot Soccer World Cup X*. Springer Verlag, Berlin, 2007. To appear. **Winner of Best Student Paper Award.**

Books

[3] **Peter Stone**, Tucker Balch, and Gerhard Kraetszchmar, editors. *RoboCup-2000: Robot Soccer World Cup IV*. Springer Verlag, Berlin, 2001.

[4] **Peter Stone**. *Layered Learning in Multiagent Systems: A Winning Approach to Robotic Soccer*. MIT Press, 2000.

Journal Articles

[5] Shimon Whiteson, Matthew E. Taylor, and **Peter Stone**. Empirical studies in action selection for reinforcement learning. *Adaptive Behavior*, 15(1), 2007. To appear.

[6] **Peter Stone**, Mohan Sridharan, Daniel Stronger, Gregory Kuhlmann, Nate Kohl, Peggy Fidelman, and Nicholas K. Jong. From pixels to multi-robot decision-making: A study in uncertainty. *Robotics and Autonomous Systems* (RAS), 54(11):933–43, November 2006. Special issue on Planning Under Uncertainty in Robotics.

[7] Daniel Stronger and **Peter Stone**. Towards autonomous sensor and actuator model induction on a mobile robot *Connection Science Journal* (CSJ), 18(2):97–119, June 2006. Special Issue on Developmental Robotics.
Based on "Simultaneous calibration of action and sensor models on a mobile robot." In *IEEE International Conference on Robotics and Automation* (ICRA), April 2005.

[8] **Peter Stone**, Richard S. Sutton, and Gregory Kuhlmann. Reinforcement learning for RoboCup-soccer keepaway. *Adaptive Behavior* (AB), 13(2):165–188, 2005.
Based on "Scaling reinforcement learning toward RoboCup soccer." In *Proceedings of the Eighteenth International Conference on Machine Learning* (ICML), 2001.

[9] Shimon Whiteson, Nate Kohl, Risto Miikkulainen, and **Peter Stone**. Evolving keepaway soccer players through task decomposition. *Machine Learning* (MLJ), 59(1):5–30, May 2005.
Based on earlier version in *Proceedings of the Genetic and Evolutionary Computation Conference* (GECCO), July 2003.

[10] Elizabeth Sklar, Simon Parsons, and **Peter Stone**. Using RoboCup in university-level computer science education *Journal of Educational Resources in Computing* (JERIC), 4:2, June 2004. Special Issue on Robotics in Undergraduate Education, Part 1.

[11] Itsuki Noda and **Peter Stone**. The RoboCup soccer server and CMUnited clients: Implemented infrastructure for MAS research. *Autonomous Agents and Multi-Agent Systems*, 2002. To appear.

[12] Michael Bowling Manuela Veloso and **Peter Stone**. The CMUnited-98 champion small-robot team. *Advanced Robotics*, 13(8):753 – 766, 2000.

[13] Manuela Veloso, **Peter Stone**, and Kwun Han. The CMUnited-97 robotic soccer team: Perception and multiagent control. *Robotics and Automated Systems*, 2000. Also in *Proceedings of the Second International Conference on Autonomous Agents*, May 1998.

[14] **Peter Stone** and Manuela Veloso. Task decomposition, dynamic role assignment, and low-bandwidth communication for real-time strategic teamwork. *Artificial Intelligence*, 110(2):241–273, June 1999.

[15] **Peter Stone** and Manuela Veloso. A layered approach to learning client behaviors in the RoboCup soccer server. *Applied Artificial Intelligence*, 12:165–188, 1998.

[16] Minoru Asada, Yasuo Kuniyoshi, Alexis Drogoul, Hajime Asama, Maja Mataric, Dominique Duhaut, **Peter Stone**, and Hiroaki Kitano. The RoboCup physical agent challenge: Phase-I. *Applied Artificial Intelligence*, 12.2:251–263, March 1998.

[17] **Peter Stone** and Manuela Veloso. Towards collaborative and adversarial learning: A case study in robotic soccer. *International Journal of Human-Computer Studies*, 48(1):83–104, January 1998.

[18] Nate Kohl and Peter Stone. Machine learning for fast quadrupedal locomotion. In *The Nineteenth National Conference on Artificial Intelligence*, pages 611–616, July 2004.

5 Research Proposal

The research in the learning agents research group (LARG) at UT Austin is focussed on machine learning in multiagent systems. As such, the RoboCup simulation league has been, and continues to be, a valuable research platform for us.

As the main challenge in 2007 will be low-level control of the humanoid simulated robot, the main focus of our proposed research is machine learning for low-level control. In particular, we aim to leverage both the policy search reinforcement learning algorithms that we have successfully developed in the context of the Four-Legged league for the purpose of Aibo walking (18), and the model-based reinforcement learning approach to 3 vs. 2 keepaway that we have developed in the 2D simulator. A paper the model-based reinforcement learning approach is currently under review. The details of our research methodology appear in these two full-length conference papers.

6 Agent Description

For stage 1 of qualification, the UT Austin Villa humanoid agent is programmed to exhibit a simple (and somewhat unconventional!) mode of locomotion. Starting from a standing position, the agent bends his knees to fall to the ground. Once flat on the ground, he raises his arms and beats them downward, as a result of which he is propelled forward. By repeating this action over and over again, the agent is able to move forward.

The mechanism describing the agent's behavior is as follows. We define the state of the agent to be a vector comprised of all the joint angles. This specification is incomplete, as orientation information (roll, pitch, and yaw w.r.t. field) is not available; nonetheless, we find it sufficient for the purpose of this demonstration. Begining from the start state, a set of rules describes the next state to be reached from any given state. For instance, the state ARMS_IN_FRONT (shoulder angles 120 degrees in front, all other angles 0) is followed by the state ARMS_BEHIND (shouder angles 70 degrees behind, all other angles 0), which is again followed by ARMS_IN_FRONT. The transition between states is achieved through PID control of all the effectors to achieve the target configuration.

At this point, our agent does not process visual information. If, by chance, the agent falls on the ball, he is quite unaware of it, and continues to beat his hands as usual. In such a situation, it may take a while for him to slide back on to the ground and start moving. Our primary focus in further development of the agent would be to design a stable (possibly bipedal) locomotion mechanism, which will also rely on orientation (in the next release?) and visual information. Subsequently, we will extend this to actions like passing and intercepting, before integrating them into team behavior.

N.B.: The locomotion behavior described above works without problems in our installation. We expect it to work just as well when the organizers execute it, provided there aren't timing issues. One check is in the length of received messages (main.cpp), which we print out: these should be 1500.