The UT Austin Villa 2003 Legged Robot Team

Peter Stone
With members of th UT Austin Spring 2003 class CS395T: Multi-Robot Systems

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

Abstract. This document serves as the team description paper for the purposes of applying to qualify for the RoboCup 2003 legged robot league competition in Padova, Italy.

1 Research Interests

The primary research interests of the UT Austin Villa legged robot team are in the areas machine learning and multiagent systems within the context of creating complete intelligent agents. That is, we view all aspects of creating robotic agents, including sensation, localization, and locomotion, to be potential opportunities for optimization via machine learning methods and/or leveraging from multi-robot interactions.

We view the RoboCup domain as a wonderful testbed for this purpose. In particular, it *requires* the creation of robust and efficient vision, localization, and locomotion capabilities. Without these capabilities, a team will not stand a chance against quicker, more situated teams. However, once sufficient low-level skills have been achieved, there is lots for room for team improvement at the level of strategy and teamwork, both of which are prime machine learning and multiagent systems challenges.

Given our past and on-going research on multiagent learning in the RoboCup simulator, we also have a deep interest in the extent to which techniques that have proven successful in simulation can be extended to real robots. We have already had one past success in this regard, having developed a flexible team agent architecture in simulation and successfully ported it to a RoboCup small-size robot team (Stone&Veloso, 1999). A current open research question is whether temporal difference learning techniques can be applied as successfully on real robots as they can be in simulation. We are also particularly interested in exploring the extent to which our *layered learning* machine learning paradigm (Stone, 2000) can be applied in a real robotic domain.

2 Approach

Our approach to the RoboCup challenge with legged robots is essentially a bottom-up approach. As a new team that began our concerted effort in mid-January of 2003, our over-arching foci thus far has been on developing low-level sensory and locomotive skills and on "closing the loop:" creating our first fully-autonomous soccer-playing behavior. We are committed to doing so *from scratch*, i.e. without using any pre-existing code bases, so as to contribute a completely new code-base to the community.

As we proceed, we are collecting research challenges to return to. And we are always aiming to create general, flexible solutions that will easily extend to, for instance, less controlled lighting environments and environments with different and/or fewer beacons.

We view our current efforts as largely the infrastructure-building phase of our long-term research program using legged robots. The American Open and RoboCup 2003 competitions are serving as concrete deadlines by which time we must have fully working systems. While we aim to be competitive at these events, we are mainly looking forward to being able to study our working solutions with controlled experiments after the competitions, and to being able to extend them with novel techniques.

With regards to specific components, we are currently pursuing many approaches in parallel, including the following:

Vision: For the purposes of color segmentation, we are using a nearest neighbor approach in the YCrCb color space based on a small number of hand-segmented images. For the purposes of speeding up segmentation, we have developed a tool that allows us to capture images, "paint" them, and load the resulting color table onto a continually running robot.

Localization: Our localization routine is designed to produce distributions of possible current locations. The distributions result from a sampling-based method connected to the uncertainty in beacon sizes calculated from the pixel-length of the beacons in the images. Inspired by the German Team's presentation on the Open-R website, we are currently pursuing edge-detection methods for the purposes of localizing based entirely on the field-wall boundaries and the while lines on the field.

Locomotion: We are actively developing and comparing several different approaches to walking including one based on tracing out simple shapes with the aid of a new, efficient solution to the inverse kinematics problem; one based on a hill-climbing approach in joint space with respect to an internal simulated model of the robot's physical structure; and one based on a decomposition of an arbitrary trajectory into four equivalent component steps. Each of these approaches has yielded a parameterized walk with different properties. We are in the process of optimizing the parameters and exploring their relative strengths. We view these optimization problems as some of the prime candidates for our first machine learning experiments.

3 Background of Principal Investigator

Dr. Peter Stone is an Assistant Professor in the Department of Computer Sciences at the University of Texas at Austin. He received his Ph.D. in 1998 and his M.S. in 1995 from Carnegie Mellon University, both in Computer Science. He received his B.S. in Mathematics from the University of Chicago in 1993. From 1999 to 2002 he was a Senior Technical Staff Member in the Artificial Intelligence Principles Research Department at AT&T Labs - Research.

Prof. Stone's research interests include planning and machine learning, particularly in multiagent systems. His doctoral thesis research contributed a flexible multiagent team structure and multiagent machine learning techniques for teams operating in real-time noisy environments in the presence of both teammates and adversaries. His long-term research goal is to create complete, robust, autonomous agents that can learn to interact with other intelligent agents in a wide range of complex, dynamic environments.

Prof. Stone is currently continuing his investigation of machine learning and multiagent learning at UT Austin. Application domains include robotic soccer, autonomous bidding agents for auctions, and social agents. Within the robotic soccer domain, he is studying multiagent techniques in reinforcement learning, specifically temporal difference learning, for learning successful policies by a team of cooperating agents. In the context of auctions, he is investigating adaptive bidding policies that are applicable for simultaneous multi-round auctions involving interacting goods. In a chatroom style domain, he has helped create a social agent that learns when and how to initiate conversations with human users.

Prof. Stone is a member of the international RoboCup executive committee and was a co-chair of RoboCup-2001 at IJCAI-01. He has developed teams of robotic soccer agents in that have won RoboCup championships in the simulation (1998, 1999) and in the small-wheeled robot (1997, 1998) leagues. He led tutorials on robotic soccer at AAAI-99, Agents-99, and IJCAI-99. He has also developed agents that have won auction trading agents competitions (2000, 2001). Peter has served on various program committees and has co-chaired workshops on learning agents (at Agents-2000, Agents-2001, and the AAAI Spring Symposium in 2002) and on RoboCup (at RoboCup-2000).

Prof. Stone is the author of *Layered Learning in Multiagent Systems: A Winning Approach to Robotic Soccer* (MIT Press, 2000) and editor of *RoboCup 2000: Robot Soccer World Cup IV* (Springer Verlag, 2001) as well as several technical papers in conferences and journals. Prof. Stone won the best-paper award at the Agents-2001 conference and was awarded the Allen Newell Medal for Excellence in Research in 1997.

4 Team Organization

The UT Austin Villa 2003 legged robot team is the result of a focused class effort during the Spring semester of 2003 at the University of Texas at Austin. Nineteen

graduate students and one undergraduate are enrolled in the course CS395T: *Multi-Robot Systems: Robotic Soccer with Legged Robots*. The class website is available at `http://www.cs.utexas.edu/~pstone/Courses/395Tspring03`.

Students in the class have been studying past approaches, both as described in the literature and as reflected in publicly available source code. However they are developing the entire code base *from scratch* with the goal of contributing a completely new code base to the community.

Although most students are interested in several aspects of the overall challenge, each is focusing primarily on one of the following subtasks:

- Vision
- Localization
- Fall detection
- Movement
- Kicking
- Communication
- General architecture
- Coordination
- Auxiliary tools

Class sessions are devoted to students educating each other about their findings and progress, as well as coordinating the integration of everybody's code. Just nine weeks after their initial introduction to the robots, the students already have preliminary working solutions to vision, localization, fast walking, kicking, and communication.

The concrete goal of the course is to have a completely new working solution by the end of April so that we can participate in the American Open competition, which happens to fall during the last week of the class. After that point, we expect to have five to seven students continue working towards RoboCup 2003 in Padova.

5 Relevant Publications

The principal investigator 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 thirty 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.

Books

1. **Peter Stone**, Tucker Balch, and Gerhard Kraetszchmar, editors. *RoboCup-2000: Robot Soccer World Cup IV*. Springer Verlag, Berlin, 2001.
2. **Peter Stone**. *Layered Learning in Multiagent Systems: A Winning Approach to Robotic Soccer*. MIT Press, 2000.

Journal Articles

1. 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.
2. Michael Bowling Manuela Veloso and **Peter Stone**. The CMUnited-98 champion small-robot team. *Advanced Robotics*, 13(8):753 – 766, 2000.
3. Manuela Veloso, **Peter Stone**, and Kwun Han. The CMUnited-97 robotic soccer team: Perception and multi-agent control. *Robotics and Automated Systems*, 2000. Also in *Proceedings of the Second International Conference on Autonomous Agents*, May 1998.
4. **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.
5. **Peter Stone** and Manuela Veloso. A layered approach to learning client behaviors in the RoboCup soccer server. *Applied Artificial Intelligence*, 12:165–188, 1998.
6. 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.
7. **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.

6 Statement of Commitment

If we qualify for the competition, UT Austin Villa will enter RoboCup-2003 in Padova. We will be responsible for securing sufficient funding to travel to Padova and register for RoboCup.

UT Austin Villa participates in the Domestic Standard Platform League (DSPL), in which all teams use the Toyota HSR robot (Yamamoto et al. 2019). Figure 1 shows the HSR, which has a holonomic omnidirec-tional base, a 4 DoF arm, a 1 DoF torso lift, a head RGBD sensor, and an RGB sheye hand camera. For fast neural net-work computation, an external laptop with an Nvidia RTX 2070 GPU is mounted to the robot.Â Since the LIDAR has a wider angle than head camera, the leg tracker results are preferred. When the targetâ€™s legs are no longer visible, the head turns around to reidentify the target. For each person detected by OpenPose, a similarity score is computed using the color histogram of the personâ€™s body region and the original photo of the target. in RoboCup-2003: Robot Soccer World Cup VII. Add To MetaCart. Tools.Â " For a robot to learn to improve its performance based entirely on real-world environmental feedback, the robot's behavior specification and learning algorithm must be constructed so as to enable data-efficient learning. Building upon previous work enabling a quadrupedal robot to learn a fast wa ..." The UT Austin Villa RoboCup 2003 Four-Legged Team was a new entry in the ongoing series of RoboCup legged league ompetitions. The team development began in mid-January of 2003, at whi h time none of the team members had any familiarity with the Aibos. Without using any RoboCup-related ode from other teams, we entered a team in the Ameri an Open ompetition at the end of April, and met with some su ess at the annual RoboCup ompetition that took pla e in Padova, Italy at the beginning of July. In this report, we des ribe both our development pro ess and the te hni al details of its end result, th...