TT-UT Austin Villa 2010 Team Description Paper for the Standard Platform League

Samuel Barrett[1], Katie Genter[1], Matthew Hausknecht[1], Todd Hester[1], Piyush Khandelwal[1], Juhyun Lee[1], Adam Setapen[1], Aibo Tian[1], Michael Quinlan[1], Mohan Sridharan[2] and Peter Stone[1]

[1] Department of Computer Science
The University of Texas at Austin
1 University Station C0500
Austin, Texas 78712-1188
[2] Department of Computer Science
Texas Tech University
1302 Pine Street
Abilene, Texas 79601
{sbarrett,katie,mhauskn,todd,piyushk,impjdi,
asetapen,atian,mquinlan,pstone}@cs.utexas.edu
mohan.sridharan@ttu.edu
http://www.cs.utexas.edu/~AustinVilla

Abstract. This paper describes the research focus and ideas incorporated in the TT-UT Austin Villa Standard Platform league team entering the RoboCup competition in 2010. TT-UT Austin Villa is a combinded team representing The University of Texas at Austin (UT) and Texas Tech Univesitry (TT).

1 Introduction

The UT Austin Villa Four-Legged Team has participated in every RoboCup competition since RoboCup 2003 in Padua. The team development began in mid-January of 2003 without any prior familiarity with the Aibos. After entering a fairly non-competitive team in RoboCup 2003, the team made several important advances. By the July 2004 competition that took place in Lisbon, Portugal, it was one of the top few teams, and it has continued to be competitive ever since, including a quarter-final appearance in 2007. In 2008 we made the quarter-final of the Nao league and finised 4th in the AIBO league. In May 2009 we placed 1st at the US Open in the Standard Platform League and we placed 4th in the SPL at RoboCup 2009. In 2010, we repeated as champions at the US Open. Meanwhile, we have placed extensive focus on identifying and developing the core research contributions from our team, which have since spawned several spin-off research efforts using the Aibos.

The technical details of our past Nao and four-legged teams are available in our series of technical reports [14–17, 2, 4], as well as in the inaugural book in the MorganClaypool Synthesis Lecture Series on Artificial Intelligence and Machine Learning [13]. This book presents a roadmap for getting started on *any* vision-based and/or legged-based robot, using the Aibo as a case study.

2 Research Contributions

Our research on the robots, all of which has been based on the TT-UT Austin Villa code base, has led to more than 25 published research papers over the past 5 years. Full details are available on our team website: `www.cs.utexas.edu/~AustinVilla`. This section summarizes our research contributions from the past couple of years.

2.1 Reinforcement Learning on the Nao

Reinforcement learning (RL) algorithms have long been promising methods for enabling an autonomous robot to improve its behavior on sequential decision-making tasks. The obvious enticement is that the robot should be able to improve its own behavior without the need for detailed step-by-step programming. In this work [3], we presented an algorithm, Reinforcement Learning with Decision Trees (RL-DT), that uses decision trees to learn the model by generalizing the relative effect of actions across states. The agent explores the environment until it believes it has a reasonable policy. We tested RL-DT on an Aldebaran Nao humanoid robot scoring goals in a penalty kick scenario. More details and video of the robot learning to kick are available online:`http://www.cs.utexas.edu/~AustinVilla/?p=research/rl_kick`.

2.2 Generalized Planned Color Learning

In previous work [9], we had enabled the robot to learn the colors on the robot soccer field, modeling colors as 3D Gaussians, using a pre-defined motion sequence. In this work, we extended the approach in two significant ways. The color learning works both in the controlled lab setting and in un-engineered indoor corridors by proposing a hybrid color model. We also enabled the robot to plan a motion sequence appropriate for learning colors, using the known model of its color-coded world. The algorithm is described in [10] and detailed experimental results can be found online:`www.cs.utexas.edu/users/AustinVilla/?p=research/gen_color`.

2.3 Adapting to Changing Illumination Conditions

In previous work [8], we had shown that if the robot is provided suitable color maps and image statistics for different illumination conditions, it can transition smoothly between the color maps based on a comparison of the image characteristics. We aim to have the entire color learning algorithm to execute autonomously under changing illumination conditions. We extended our approach by enabling the robot to detect changes in illumination conditions automatically. If an illumination change is detected, the robot automatically adapts to the change by revising its color knowledge be re-learning the colors. Complete details, including the algorithm and experimental results, are available in [11] and supporting images are available for viewing online:`www.cs.utexas.edu/users/AustinVilla/?p=research/illuminvar_colorlearn`.

2.4 Learning a More Stable Walk

A fast gait is an essential component of any successful team in the RoboCup 4-legged league. However, quickly moving quadruped robots, including those with learned gaits, often move in such a way so as to cause unsteady camera motions which degrade the robot's visual capabilities. In previous research, we presented a method for automatically learning a *fast* gait [6]. In this work, we presented an implementation of the policy gradient machine learning algorithm that searches for a parameterized walk while optimizing for both speed and stability [7]. To the best of our knowledge, previous learned walks have all focused exclusively on speed. Our method is fully implemented and tested on the Sony Aibo ERS-7 robot platform. The resulting gait is reasonably fast and considerably more stable compared to our previous fast gaits. We demonstrate that this stability can significantly improve the robot's visual object recognition. Videos are available on-line at `www.cs.utexas.edu/~AustinVilla/?p=research/learned_walk`.

2.5 Learning Powerful Kicks

Coordinating complex motion sequences remains a challenging task for robotics. Machine Learning has aided this process, successfully improving motion sequences such as walking and grasping [7]. However, to the best of our knowledge, outside of simulation, learning has never been applied to the task of kicking the ball. We apply machine learning methods to optimize kick power entirely on a real robot. The resulting learned kick is significantly more powerful than the most powerful hand-coded kick of one of the most successful RoboCup four-legged league teams, and is learned in a principled manner which requires very little engineering of the parameter space. Finally, model inversion is applied to the problem of creating a parameterized kick capable of kicking the ball a specified distance. The associated paper [1] and additional resources can be found at `http://www.cs.utexas.edu/~AustinVilla/?p=research/aibo_kick`.

2.6 Selective Visual Attention for Object Detection

Autonomous robots can use a variety of sensors, such as sonar, laser range finders, and bump sensors, to sense their environments. Visual information from an on-board camera can provide particularly rich sensor data. However, processing all the pixels in every image, even with simple operations, can be computationally taxing for robots equipped with cameras of reasonable resolution and frame rate. We present a novel method for a legged robot equipped with a camera to use selective visual attention to efficiently recognize objects in its environment [19]. The resulting attention-based approach is fully implemented and validated on an Aibo ERS-7. It effectively processes incoming images 50 times faster than a baseline approach, with no significant difference in the efficacy of its object detection. More information and a video is available on-line at `www.cs.utexas.edu/~AustinVilla/?p=research/model-based_vision`.

2.7 Autonomous Sensor and Actuator Model Induction

We presented a novel methodology for a robot to autonomously induce models of its actions and sensors called ASAMI (Autonomous Sensor and Actuator Model Induction) [18]. While previous approaches to model learning rely on an independent source of training data, we show how a robot can induce action and sensor models without any well-calibrated feedback. Specifically, the only inputs to the ASAMI learning process are the data the robot would naturally have access to: its raw sensations and knowledge of its own action selections. From the perspective of developmental robotics, our robot's goal is to obtain self-consistent internal models, rather than to perform any externally defined tasks. Furthermore, the target function of each model-learning process comes from within the system, namely the most current version of another internal system model. Concretely realizing this model-learning methodology presents a number of challenges, and we introduce a broad class of settings in which solutions to these challenges are presented. ASAMI is fully implemented and tested, and empirical results validate our approach in a robotic test-bed domain using a Sony Aibo ERS-7 robot. Videos of the learning process are available on-line at `www.cs.utexas.edu/~AustinVilla/?p=research/learned_walk`.

2.8 Negative Information and Line Observations for Monte Carlo Localization

In previous work [12], we had developed a robust Monte Carlo Localization algorithm for use on vision-based legged robots. In this work, we improved upon that algorithm by incorporating negative information and line observations into our algorithm. Particles are updating using negative information anytime a landmark is expected but not seen. In an environment with few landmarks, updating with negative information can be very useful. Our new algorithm also makes use of observations of field lines, incorporating them into the algorithm using the distance and heading to the nearest point on the line. The algorithm has been fully implemented and tested both on a Sony Aibo ERS-7 robot as well as in simulation. The algorithm and the results are described in [5].

3 Conclusion

We look forward to continuing and expanding our above research on the new Nao robot.

References

1. Matthew Hausknecht and Peter Stone. Learning powerful kicks on the aibo ers-7: The quest for a striker. In *Proceedings of the RoboCup International Symposium 2010*. Springer Verlag, 2010.
2. Todd Hester, Michael Quinlan, and Peter Stone. UT Austin Villa 2008: Standing on Two Legs. Technical Report UT-AI-TR-08-8, The University of Texas at Austin, Department of Computer Sciences, AI Laboratory, November 2008.
3. Todd Hester, Michael Quinlan, and Peter Stone. Generalized model learning for reinforcement learning on a humanoid robot. In *IEEE International Conference on Robotics and Automation (ICRA)*, May 2010.

4. Todd Hester, Michael Quinlan, Peter Stone, and Mohan Sridharan. TT-UT Austin Villa 2009: Naos across Texas. Technical Report UT-AI-TR-09-08, The University of Texas at Austin, Department of Computer Science, AI Laboratory, December 2009.
5. Todd Hester and Peter Stone. Negative information and line observations for monte carlo localization. In *IEEE International Conference on Robotics and Automation*, May 2008.
6. Nate Kohl and Peter Stone. Machine learning for fast quadrupedal locomotion. In *The Nineteenth National Conference on Artificial Intelligence*, pages 611–616, July 2004.
7. Manish Saggar, Thomas D'Silva, Nate Kohl, and Peter Stone. Autonomous learning of stable quadruped locomotion. In Gerhard Lakemeyer, Elizabeth Sklar, Domenico Sorenti, and Tomoichi Takahashi, editors, *RoboCup-2006: Robot Soccer World Cup X*. Springer Verlag, Berlin, 2007. To appear.
8. M. Sridharan and P. Stone. Towards illumination invariance in the legged league. In *The International RoboCup Symposium*, 2004.
9. M. Sridharan and P. Stone. Autonomous color learning on a mobile robot. In *The Twentieth National Conference on Artificial Intelligence (AAAI)*, 2005.
10. M. Sridharan and P. Stone. Autonomous planned color learning on a mobile robot without labeled data. In *The Ninth IEEE International Conference on Control, Automation, Robotics and Vision (ICARCV)*, December 2006.
11. M. Sridharan and P. Stone. Color learning on a mobile robot: Towards full autonomy under changing illumination. In *The International Joint Conference on Artificial Intelligence (IJCAI)*, Hyderabad, India, January 2007.
12. Mohan Sridharan, Gregory Kuhlmann, and Peter Stone. Practical vision-based monte carlo localization on a legged robot. In *IEEE International Conference on Robotics and Automation*, April 2005.
13. Peter Stone. *Intelligent Autonomous Robotics: A Robot Soccer Case Study*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers, 2007.
14. Peter Stone, Kurt Dresner, Selim T. Erdoğan, Peggy Fidelman, Nicholas K. Jong, Nate Kohl, Gregory Kuhlmann, Ellie Lin, Mohan Sridharan, Daniel Stronger, and Gurushyam Hariharan. UT Austin Villa 2003: A new RoboCup four-legged team. Technical Report UT-AI-TR-03-304, The University of Texas at Austin, Department of Computer Sciences, AI Laboratory, 2003.
15. Peter Stone, Kurt Dresner, Peggy Fidelman, Nicholas K. Jong, Nate Kohl, Gregory Kuhlmann, Mohan Sridharan, and Daniel Stronger. The UT Austin Villa 2004 RoboCup four-legged team: Coming of age. Technical Report UT-AI-TR-04-313, The University of Texas at Austin, Department of Computer Sciences, AI Laboratory, October 2004.
16. Peter Stone, Kurt Dresner, Peggy Fidelman, Nate Kohl, Gregory Kuhlmann, Mohan Sridharan, and Daniel Stronger. The UT Austin Villa 2005 RoboCup four-legged team. Technical Report UT-AI-TR-05-325, The University of Texas at Austin, Department of Computer Sciences, AI Laboratory, November 2005.
17. Peter Stone, Peggy Fidelman, Nate Kohl, Gregory Kuhlmann, Tekin Mericli, Mohan Sridharan, and Shao en Yu. The UT Austin Villa 2006 RoboCup four-legged team. Technical Report UT-AI-TR-06-337, The University of Texas at Austin, Department of Computer Sciences, AI Laboratory, December 2006.
18. Daniel Stronger and Peter Stone. Towards autonomous sensor and actuator model induction on a mobile robot. *Connection Science*, 18(2):1–23, 2006. Special Issue on Developmental Robotics.
19. Daniel Stronger and Peter Stone. Selective visual attention for object detection on a legged robot. In Gerhard Lakemeyer, Elizabeth Sklar, Domenico Sorenti, and Tomoichi Takahashi, editors, *RoboCup-2006: Robot Soccer World Cup X*. Springer Verlag, Berlin, 2007. To appear.