

UT Austin Villa 2017 Team Description Paper for the Standard Platform League

Katie Genter, Josiah Hanna, Josh Kelle,
Elad Liebman, Jacob Menashe, Sanmit Narvekar,
Ruohan Zhang and Peter Stone

Department of Computer Science
The University of Texas at Austin
2317 Speedway
Austin, Texas 78712

UTAustinVilla-legged@cs.utexas.edu
<http://www.cs.utexas.edu/~AustinVilla>

Abstract. This paper describes the research focus and ideas incorporated in the UT Austin Villa Standard Platform league team entering the RoboCup competition in 2017. UT Austin Villa is a team representing the Department of Computer Science at The University of Texas at Austin.

1 Introduction

The UT Austin Villa Standard Platform Team has participated in every RoboCup competition since RoboCup 2003 in Padua (at which time it was still called the Four-Legged league). The team development began in mid-January of 2003 without any prior familiarity with the robots (AIBOs, at the time). After entering a fairly non-competitive team in RoboCup 2003, the team made several important advances. After reaching the quarterfinals in 2007, the team has consistently performed well. In 2008 the team made the quarter-final of the Nao league and finished 4th in the AIBO league. The team placed 4th at RoboCup 2009, 3rd at RoboCup 2010, and finally captured 1st place at Robocup 2012 in Mexico City. The team took 3rd place at RoboCup 2013 and reached the quarter-final round at RoboCup 2015. The team placed 2nd at RoboCup 2016. Throughout, we have placed extensive focus on identifying and developing the core research contributions from our team.

The technical details of our past Nao and four-legged teams are available in our series of technical reports [26–29, 10, 12, 1, 3], as well as in the inaugural book in the MorganClaypool Synthesis Lecture Series on Artificial Intelligence and Machine Learning [25]. This book presents a roadmap for getting started on *any* vision-based and/or legged-based robot, using the Aibo as a case study. Additionally, the technical details of our 2012 Standard Platform Championship team can be found in our champions paper that was published in the RoboCup-2012: Robot Soccer World Cup XVI book [2].

2 Research Contributions

Our research on the robots, all of which has been based on the UT Austin Villa code base, has led to more than 30 published research papers. Full details are available on our team website: www.cs.utexas.edu/~AustinVilla. This section summarizes some of our recent research contributions using the Nao robots, as well as some interesting older contributions made while using the Aibo robots.

2.1 Drop-in Player Competition

The Standard Platform League (SPL) and the 3D simulation league have both held Drop-in Player Competitions for the last four years. In the Drop-in Player Competitions, each team contributes a player that must play as a team with other players from various teams using a limited communication protocol. Austin Villa has been involved in these competitions both as competitors and as organizers. Members of our team have also worked to document these competitions. We documented the competition across the SPL, 3D simulation, and 2D simulation leagues in 2013 [17]. In 2015, we contributed to a paper that documented the 2014 SPL Drop-in Competition [6]. In 2017, we contributed to a journal article that tracked the progress of the SPL competition over three years [7].

2.2 Grounded Simulation Learning

Simulation is often used in research and industry as a low cost, high efficiency alternative to real model testing. Simulation has also been used to develop and test powerful learning algorithms. However, parameters learned in simulation often do not translate directly to the application, especially because heavy optimization in simulation has been observed to exploit the inevitable simulator simplifications, thus creating a gap between simulation and application that reduces the utility of learning in simulation. This paper [5] introduces Grounded Simulation Learning (GSL), an iterative optimization framework for speeding up robot learning using an imperfect simulator. In GSL, a behavior is developed on a robot and then repeatedly: 1) the behavior is optimized in simulation; 2) the resulting behavior is tested on the real robot and compared to the expected results from simulation, and 3) the simulator is modified, using a machine-learning approach to come closer in line with reality. This approach is fully implemented and validated on the task of learning to walk using an Aldebaran Nao humanoid robot. Starting from a set of stable, hand-coded walk parameters, four iterations of this three-step optimization loop led to more than a 25% increase in the robot's walking speed.

2.3 Grounded Action Transformation for Robot Learning in Simulation

Robot learning in simulation is a promising alternative to the prohibitive sample cost of learning in the physical world. Unfortunately, policies learned in simulation often perform worse than hand-coded policies when applied on the physical robot. Grounded

simulation learning (GSL) promises to address this issue by altering the simulator to better match the real world. We contributed a paper that proposes a new algorithm for GSL – Grounded Action Transformation – and applies it to learning of humanoid bipedal locomotion [8]. Our approach results in a 43.27% improvement in forward walk velocity compared to a state-of-the-art hand-coded walk. We further evaluate our methodology in controlled experiments using a second, higher-fidelity simulator in place of the real world. The results presented in our paper contribute to a deeper understanding of grounded simulation learning and demonstrate its effectiveness for learning robot control policies.

2.4 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 [11], 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.

2.5 Ground Truth Detection System

Ground truth detection systems can be a crucial step in evaluating and improving algorithms for self-localization on mobile robots. Selecting a ground truth system depends on its cost, as well as on the detail and accuracy of the information it provides. In this work [15], we present a low cost, portable and real-time solution constructed using the Microsoft Kinect RGB-D Sensor. We use this system to find the location of robots and the orange ball in the SPL environment in the RoboCup competition. This system is fairly easy to calibrate, and does not require any special identifiers on the robots. We also provide a detailed experimental analysis to measure the accuracy of the data provided by this system. Although presented for the SPL, this system can be adapted for use with any indoor structured environment where ground truth information is required. Details on using this system are available online: <http://www.cs.utexas.edu/~AustinVilla/?p=research/kinect>.

2.6 Using Gaussian Fitness Scores for Vision Improvements

In RoboCup, although the fields are standardized and color coded, the area outside the fields often contains many objects of various colors. Sometimes objects off the field may look very similar to balls, robots, or other objects normally found on the soccer field. Robots must detect all of these objects, and then differentiate between the true positives and false positives. This paper [18] presents a new method using Gaussian fitness scores

to differentiate between true positives and false positives for balls, robots, and penalty crosses. We also present some other improvements in our code base following our 2012 championship, such as our usage of a virtual base for forward kinematics calculations, our ability to flexibly transition player roles given dynamic numbers of teammates, and our ability to quickly integrate new kicks of varying speeds into our strategy. With these improvements, our UT Austin Villa team finished third in the Standard Platform League at RoboCup 2013.

2.7 Controlled Kicking under Uncertainty

In RoboCup, robots must make quick decisions under uncertainty. To this end, we developed a new approach to enable humanoid soccer robots to execute kicks quickly and ensure that they move the ball down field. We developed a kick engine capable of kicking at a variety of distances and angles and then a novel kick decision method for selecting from among a large set of possible kicks. This method prunes and orders the kicks according to a metric and then chooses the first possible kick that ensures that our field position is improved [4].

2.8 Vision Calibration and Processing

The Aldebaran Nao, has two cameras for visual input, of which only one has been typically used. The integration of both cameras presents a new opportunity but also a challenge. While it is possible to obtain better information using both cameras, more cameras require more work to calibrate. We developed a novel camera calibration algorithm which automatically tuned a camera such that its color perceptions matched those of another camera. Additionally, recent vision challenges introduced in RoboCup have necessitated the use of higher resolution images. We built on existing work in color based segmentation and presented novel extensions to facilitate the move to higher resolution images, including memory optimizations, fast line and curve detection, and differentiation via robot pose based transformations [14].

2.9 Generalized Planned Color Learning

In previous work [21], 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 [22] and detailed experimental results can be found online: www.cs.utexas.edu/users/AustinVilla/?p=research/gen_color.

2.10 Adapting to Changing Illumination Conditions

In previous work [20], 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 by re-learning the colors. Complete details, including the algorithm and experimental results, are available in [23] and supporting images are available for viewing online: www.cs.utexas.edu/users/AustinVilla/?p=research/illumivar_colorlearn.

2.11 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 [16]. In this work, we presented an implementation of the policy gradient machine learning algorithm that searches for a parametrized walk while optimizing for both speed and stability [19]. 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.

2.12 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 [19]. 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 parametrized kick capable of kicking the ball a specified distance. The associated paper [9] and additional resources can be found at http://www.cs.utexas.edu/~AustinVilla/?p=research/aibo_kick.

2.13 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 [31]. 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.

2.14 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) [30]. 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.

2.15 Negative Information and Line Observations for Monte Carlo Localization

In previous work [24], 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 [13].

References

1. Samuel Barrett, Katie Genter, Matthew Hausknecht, Todd Hester, Piyush Khandelwal, Juhyun Lee, Michael Quinlan, Aibo Tian, Peter Stone, and Mohan Sridharan. Austin Villa 2010 Standard Platform team report. Technical Report UT-AI-TR-11-01, The University of Texas at Austin, Department of Computer Sciences, AI Laboratory, January 2011.

2. Samuel Barrett, Katie Genter, Yuchen He, Todd Hester, Piyush Khandelwal, Jacob Menashe, and Peter Stone. UT Austin Villa 2012: Standard Platform League world champions. In Xiaoping Chen, Peter Stone, Luis Enrique Sucar, and Tijn Van der Zant, editors, *RoboCup-2012: Robot Soccer World Cup XVI*, Lecture Notes in Artificial Intelligence. Springer Verlag, Berlin, 2013.
3. Samuel Barrett, Katie Genter, Todd Hester, Piyush Khandelwal, Michael Quinlan, Peter Stone, and Mohan Sridharan. Austin Villa 2011: Sharing is caring: Better awareness through information sharing. Technical Report UT-AI-TR-12-01, The University of Texas at Austin, Department of Computer Sciences, AI Laboratory, January 2012.
4. Samuel Barrett, Katie Genter, Todd Hester, Michael Quinlan, and Peter Stone. Controlled kicking under uncertainty. In *The Fifth Workshop on Humanoid Soccer Robots at Humanoids 2010*, December 2010.
5. Alon Farchy, Samuel Barrett, Patrick MacAlpine, and Peter 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.
6. Katie Genter, Tim Laue, and Peter Stone. Benchmarking robot cooperation without pre-coordination in the robocup standard platform league drop-in player competition. In *Proceedings of the 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-15)*, September 2015.
7. Katie Genter, Tim Laue, and Peter Stone. Three years of the robocup standard platform league drop-in player competition: Creating and maintaining a large scale ad hoc teamwork robotics competition. *Autonomous Agents and Multi-Agent Systems (JAAMAS)*, pages 1–31, 2016.
8. Josiah Hanna and Peter Stone. Grounded action transformation for robot learning in simulation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, February 2017.
9. 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.
10. 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.
11. 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.
12. 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.
13. Todd Hester and Peter Stone. Negative information and line observations for monte carlo localization. In *IEEE International Conference on Robotics and Automation*, May 2008.
14. Piyush Khandelwal, Matthew Hausknecht, Juhyun Lee, Aibo Tian, and Peter Stone. Vision calibration and processing on a humanoid soccer robot. In *The Fifth Workshop on Humanoid Soccer Robots at Humanoids 2010*, December 2010.
15. Piyush Khandelwal and Peter Stone. A low cost ground truth detection system using the kinect. In Thomas Roefer, Norbert Michael Mayer, Jesus Savage, and Uluc Saranlı, editors, *RoboCup-2011: Robot Soccer World Cup XV*, Lecture Notes in Artificial Intelligence. Springer Verlag, Berlin, 2012.
16. Nate Kohl and Peter Stone. Machine learning for fast quadrupedal locomotion. In *The Nineteenth National Conference on Artificial Intelligence*, pages 611–616, July 2004.

17. Patrick MacAlpine, Katie Genter, Samuel Barrett, and Peter Stone. The RoboCup 2013 drop-in player challenges: Experiments in ad hoc teamwork. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, September 2014.
18. Jacob Menashe, Katie Genter, Samuel Barrett, and Peter Stone. Ut austin villa 2013: Advances in vision, kinematics, and strategy. In *The Eighth Workshop on Humanoid Soccer Robots at Humanoids 2013*, October 2013.
19. 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.
20. M. Sridharan and P. Stone. Towards illumination invariance in the legged league. In *The International RoboCup Symposium*, 2004.
21. M. Sridharan and P. Stone. Autonomous color learning on a mobile robot. In *The Twentieth National Conference on Artificial Intelligence (AAAI)*, 2005.
22. 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.
23. 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.
24. 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.
25. Peter Stone. *Intelligent Autonomous Robotics: A Robot Soccer Case Study*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers, 2007.
26. Peter Stone, Kurt Dresner, Selim T. Erdođan, 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.
27. 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.
28. 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.
29. 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.
30. 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.
31. 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.