

Austin Villa@Home 2019

DSPL Team Description Paper

Justin W. Hart¹, Yuqian Jiang¹, Raymond J. Mooney¹, Scott Niekum¹, Luis Sentis³, Peter Stone¹, and Andrea Thomaz²

¹Department of Computer Science

²Department of Electrical and Computer Engineering

³Department of Aerospace Engineering

University of Texas at Austin

<http://www.cs.utexas.edu/AustinVilla/?p=athome>

Abstract. UT Austin Villa@Home has now participated in three RoboCup@Home competitions, performing respectably in each. What is more exciting, however, is that we have begun a strong program of research that has been in part inspired by our efforts in this competition. It is our intention to build a comprehensive service robot system which is used in our laboratories, in a real-world deployment in our CS Department, and to compete in RoboCup@Home. In this Team Description Paper, you will find the highlights of our efforts and our plans for the 2019 competition.

1 Introduction

Using the RoboCup@Home team as a focal point for inter-department and inter-laboratory collaboration, UT Austin Villa@Home has pursued an ambitious research program towards the goal of the development of a comprehensive service robot system. We want to enter RoboCup@Home not with a suite of different programs for each round, but with a single program which is capable of competing and winning. In basically the same format, we also want this system to support our Building-Wide Intelligence (BWI) project, which is a longitudinal real-world deployment of a service robot in our computer science department.

UT Austin Villa@Home is a collaboration of five research groups at the University of Texas at Austin, spanning three departments: the Machine Learning Research Group, headed by Raymond Mooney; the Personal Autonomous Robotics Lab (PeARL), headed by Scott Niekum; the Human Centered Robotics Laboratory (HCRL), headed by Luis Sentis; the Learning Agents Research Group (LARG), headed by Peter Stone; and the Socially Intelligent Machines Lab, headed by Andrea Thomaz. These labs span the Aerospace Engineering, Computer Science, and Electrical and Computer Engineering departments.

We have competed in three RoboCup@Home events. In 2007, we took second place. In 2017, we entered into the newly-formed Domestic Standard Platform League (DSPL) and took third place, having received our robot only a couple of months before the competition. We had a big year in 2018. We constructed



Fig. 1: The RoboCup@Home Arena at UT Austin.

a RoboCup@Home Arena in the Anna Hiss Gymnasium in which our labs collaborate and where we are able to work in a true-to-life arena, Figure 1. The team developed a design intended to allow us to develop a single system which would enter into all of the stages of the competition, encompassing knowledge representation, mapping, and architectural aspects. Our efforts resulted in five papers submitted for publication, three of which have already been accepted and presented [1–3], with more in the planning phases. We also have been porting our RoboCup@Home software back to our BWIBot platform, in order to support general purpose service robot projects at UT Austin.

Going into 2019, we intend to complete some efforts which began last year and to continue pushing toward a comprehensive service robot system. Major efforts include further development on projects inspired by last year’s efforts. In particular, heavy development is currently underway on extensions of PRISM [1], a semantic mapping tool developed as part of our effort to port technologies developed for RoboCup@Home back to the BWI project, which would allow the robot to autonomously semantically map the arena both in real-time during the competition and during setup days, in lieu of human map annotation effort.

2 Software Overview

In 2018, our major effort was to revamp the core system software to work towards our goal of a single system that wins RoboCup@Home.

2.1 Core Components

The core system consists of three components, one of which was only beginning development in time for the 2018 competition and so was not used. These are:

LAAIR A Layered Architecture for Autonomous Interactive Robots [2]

Knowledge Representation Subsystem A general purpose knowledge representation subsystem for our robot.

PRISM Pose Registration for Integrated Semantic Mapping [1]. A system for integrating semantics into map data, such as room markup and item locations.

In 2017, we leveraged BWI code and other code from UT Austin’s laboratories to quickly field a system. Since the 2018 RoboCup@Home competition, we have been porting these components back to our BWI Infrastructure. The cleaned-up versions are released as open source packages, which can be found under the in the `utexas-bwi` organization on `github.com`. <https://github.com/utexas-bwi/> By accomplishing the goals of both BWI and RoboCup@Home, we hope to develop a comprehensive service robot system. Details of this effort can be found in [3]. Going into 2019, we plan significant development on our PRISM module, and are in the planning stages for enhanced HRI capabilities.

LAAIR: A Layered Architecture for Autonomous Interactive Robots

The core architecture of our system is composed as three layers (from top to bottom): a "reactive layer," a "deliberative layer," and a layer we call "skills." [2] The reactive layer implements state machines in SMACH¹, both to provide a quick, responsive interface and to allow simple scripting of the top-level structure of each round ("pick up each item on the table and shelve in the cupboard," "answer questions in speech and person recognition"). The deliberative layer interfaces to our knowledge base and performs planning tasks using Answer Set Programming (ASP) [4] in the CLINGO solver [5]. The skills layer is designed as an interface to ROS, allowing us to interface all of this to be integrated into current and future ROS infrastructure.

Knowledge Representation Subsystem Our KR subsystem stores knowledge in a MySQL database in order to allow for fast access and easy querying. Queries can be formed using a custom C++ library. The library can be interfaced through a simple predicate logic form which can be then imported into ASP for planning. Core to our KR subsystem is the ability to reason about hypothetical objects, as can be seen in our team’s demonstration video. In the video, when the person asks for a piece of fruit, the system is both able to reason about incomplete information (the non-specificity of "fruit") and to reason about a hypothetical apple that it has not yet witnessed. It is also able to address the potential inaccuracy of this information (there is no apple), and to report this error to the user.

PRISM Efforts on PRISM [1] began as we started to consider porting back techniques from RoboCup@Home to the BWI project. BWI relies, in part,

¹ <http://wiki.ros.org/smach>

on large, annotated maps, similar to those constructed during setup days in RoboCup@Home. PRISM is intended to enable our system to automatically construct these maps. PRISM adds semantic information to maps constructed using SLAM algorithms by adding a classifier, pose estimator, and semantic extractor to the system. These are used to identify objects, their poses in the environment, and to annotate the semantics associated with them. PRISM started with building signage, but we are in the process of incorporating additional semantics. It is our hope that our maps for RoboCup@Home this year will be constructed simply by turning the robot on in the arena during setup.

2.2 Other Important Components

Outside of these core components, our team relies on a number of pieces of software which have been previously described in detail in our 2018 Team Description Paper, which can be found at our team website.² Included in this paper are descriptions of current manipulation code, our audio pipeline and sound localization, and optimizations in our natural language pipeline.

Fast Perception Presented during last year’s open challenge was a system that we call Fast Perception, which improves perception speed in part by bypassing ROS calls where possible. The demo featured a partially-completed system, which we plan to have fully-operation and offered as open source software in time for the 2019 competition.

3 Research Focus

Core research from our laboratories includes work on reinforcement learning, learning from demonstration, natural language understanding, controls and human-robot interaction. Significant research efforts in our group that grew out of RoboCup@Home in 2018 include work on semantic mapping (PRISM) and the development of our robot’s core architecture. We believe that building integrated, autonomous systems is a worthwhile research endeavor in itself, and so have developed several complete robot platforms in the pursuit of our research, several of which are depicted in Figure 2.

4 Innovative Technology

Our group has made significant contributions in areas spanning robotics, artificial intelligence, and human-robot interaction, specifically in areas that both impact and are impacted by our RoboCup@Home efforts. Here is a brief overview.

² <http://www.cs.utexas.edu/AustinVilla/?p=athome>



(a) Two BWIBots from the Building-Wide Intelligence Project.



(b) SegBot with mounted Kinova Mico arm.



(c) Custom robot used in Niekum's laboratory.



(d) Dreamer is a compliant humanoid robot used in Sentis's laboratory.

Fig. 2: Platforms in use by the UT Austin Villa RoboCup@Home Team.

4.1 Natural Language Understanding

The approach that we use in natural language understanding is able recognize more complex language than keyword-based approaches. For many experiments we use a semantic parser which is able to resolve lambdas such as possessives, and to format commands in semantic forms suitable for input to our planning technologies. As such, it can resolve the intentions of previously unheard commands through what it reconstructs of their semantics. The BWIBot uses a dialog agent which is capable of functions such as clarification dialogs to better understand user intent.

4.2 Grounded multimodal language learning

The language grounding problem is one of pairing percepts to language symbols. For instance, “bring me the red bottle” would require the words “me,” “red,” and “bottle” to resolve to real-world interpretations. Our group has pursued significant work in the form of robots playing language games such as “I spy,” as a child might, in order to learn object groundings. The robot performs object explorations such as picking up or dropping objects in order to learn object properties paired with words [6].

4.3 Learning multi-step tasks from demonstration

We developed a series of LfD algorithms that, for the first time, allowed a robot to learn the structure of complex tasks such as IKEA furniture assembly from a small number of demonstrations [7]. This research led to state-of-the-art Bayesian nonparametric and control techniques that were able to automatically identify an appropriate number of skills from task demonstrations [8], infer the goals of each skill [9], construct controllers to accomplish these goals, and intelligently sequence these controllers based on perceptual feedback [7].

4.4 Robot-centric human activity recognition

Our robot uses its autonomous navigation capability in a large, unstructured, and human-inhabited environment. The activities learned by our robot were performed spontaneously by many different people who interacted with (or were observed by) the robot, as opposed to the standard methodology of asking study participants to perform certain actions. In contrast to classic computer vision approaches, our system uses both visual and non-visual cues when recognizing the activities of humans that it interacts with [10].

For space, we are limiting the length this list of contributions. Please see our team website and research group pages for more information.

5 Re-usability

We will continue our tradition of releasing components of our system to support their broad use within the research community.

We maintain public repositories and release software that comprises a fully integrated system for autonomous service robots for the every-day, human-inhabited office environment. Our system has been built on top of ROS, and is available open-source at <https://github.com/utexas-bwi/> under the BSD license. Included capabilities: Multi-robot control, navigation, floor-switching when the robot takes an elevator, symbolic/logical navigation integrated with a symbolic planner, planning and reasoning for high-level tasks.

The research platforms of Prof. Thomaz and Prof. Niekum are both a part of the HLP-R software system and share much of their code and development through this project. Included capabilities: kinesthetic teaching pipeline, primitives and tools for manipulation, and extensive, modular perception library. This software also supports research in Prof. Sonia Chernova laboratory at Georgia Tech. The HLP-R packages are available at <https://github.com/HLP-R> under the BSD license.

Our TEXPLORE code provides an open-source package for reinforcement learning on real robots.

Our `ar_track_alvar` ROS package has become a community standard for tag based perception.

Our ROS implementation of Dynamic Movement Primitives has become a popular tool for learning from demonstration.

We have also made research code available for: Bayesian changepoint detection, active articulated model estimation, and Bayesian nonparametric skill learning from demonstration.

6 Applicability to the Real World

In 2018 we established a permanent RoboCup@Home arena for our team's use at the UT Austin Campus, supplementing our existing laboratories which have simulated living room, dining, and kitchen environments. Since the 2018 competition, we have been porting our RoboCup@Home software back to our BWIBot platform, whose intended purpose is a real-world longitudinal deployment in the UT Austin Computer Science Department, performing tasks for the building's occupants. We view our work on BWI, on RoboCup@Home, and in the laboratory, as synergistic; producing real-world comprehensive service robot software that is suited to all of these environments, and inspiring research that helps us in each of these domains.

Prof. Thomaz's also serves as the CEO of Diligent Robotics, a start-up aiming to bring service robotics to hospitals. This venture is seeing the direct application of her research, robotics platform, Poli, and its supporting software.

7 Conclusion

Austin Villa@Home has been a strong competitor and has a tradition of synergistic research our RoboCup@Home team and our other research efforts. RoboCup@Home has become a driving force in robotics research at UT Austin. We look forward to seeing everyone in Sydney this summer.

References

1. Justin W. Hart, Rishi Shah, Sean Kirmani, Nick Walker, Kathryn Baldauf, Nathan John, and Peter Stone. Prism: Pose registration for integrated semantic mapping. In *Proceedings of the 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Madrid, Spain, October 2018.
2. Yuqian Jiang, Nick Walker, Minkyu Kim, Nicolas Brissonneau, Daniel S. Brown, Justin W. Hart, Scott Niekum, Luis Sentis, and Peter Stone. Laair: A layered architecture for autonomous interactive robots. In *Proceedings of the AAAI Fall Symposium on Reasoning and Learning in Real-World Systems for Long-Term Autonomy*, Washington, DC, USA, October 2018.
3. Justin W. Hart, Harel Yedidsion, Yuqian Jiang, Nick Walker, Rishi Shah, Jesse Thomason, Aishwarya Padmakumar, Rolando Fernandez, Jivko Sinapov, Raymond Mooney, and Peter Stone. Interaction and autonomy in robocup@home and building-wide intelligence. In *Proceedings of the AAAI Fall Symposium on Interactive Learning in Artificial Intelligence for Human-Robot Interaction*, Washington, DC, USA, October 2018.
4. Vladimir Lifschitz. Answer set programming and plan generation. *Artificial Intelligence*, 138(1-2):39–54, 2002.
5. Martin Gebser, Benjamin Kaufmann, Roland Kaminski, Max Ostrowski, Torsten Schaub, and Marius Schneider. Potassco: The potsdam answer set solving collection. *AI Communications*, 24(2):107–124, 2011.
6. Jesse Thomason, Aishwarya Padmakumar, Jivko Sinapov, Justin Hart, Peter Stone, and Raymond J. Mooney. Opportunistic active learning for grounding natural language descriptions. In Sergey Levine, Vincent Vanhoucke, and Ken Goldberg, editors, *Proceedings of the 1st Annual Conference on Robot Learning*, volume 78 of *Proceedings of Machine Learning Research*, pages 67–76. PMLR, 13–15 Nov 2017.
7. Scott Niekum, Sarah Osentoski, George Konidaris, Sachin Chitta, Bhaskara Marthi, and Andrew G. Barto. Learning grounded finite-state representations from unstructured demonstrations. *The International Journal of Robotics Research*, 34(2):131–157, 2015.
8. S. Niekum, S. Osentoski, G. Konidaris, and A. G. Barto. Learning and generalization of complex tasks from unstructured demonstrations. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5239–5246, Oct 2012.
9. S. Niekum, S. Osentoski, C. G. Atkeson, and A. G. Barto. Online bayesian change-point detection for articulated motion models. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1468–1475, May 2015.
10. Ilaria Gori, Jivko Sinapov, Priyanka Khante, Peter Stone, and J. K. Aggarwal. Robot-centric activity recognition ‘in the wild’. In Adriana Tapus, Elisabeth André, Jean-Claude Martin, François Ferland, and Mehdi Ammi, editors, *Social Robotics*, pages 224–234, Cham, 2015. Springer International Publishing.

Addendum

Team Name UT Austin Villa

Team Members

Professors

- Justin Hart, Raymond Mooney, Scott Niekum, Luis Sentis, Peter Stone, and Andrea Thomaz

Students

- Nicolas Brissonneau, Daniel Brown, Yuqian Jiang, Minkyu Kim, Rishi Shah, Abrar Anwar, Kathryn Baldauf, Ethan Brown, Arnau Canyadell, David Chen, Jamin Goo, David Lai, Ashay Lokhande, Marcus Mao, Gilbert Martinez, Aaditya Murthy, Yajie Niu, Mayuri Raja, Stone Tejeda, Brian Wang, Thomas Wei, Jason Cox, Christine Tu

Contact Information

Dr. Justin Hart (*Team Leader*) hart@cs.utexas.edu

Prof. Raymond Mooney mooney@cs.utexas.edu

Prof. Scott Niekum sniekum@cs.utexas.edu

Prof. Luis Sentis lsentis@austin.utexas.edu

Prof. Peter Stone pstone@cs.utexas.edu

Prof. Andrea Thomaz athomaz@ece.utexas.edu

Website URL <http://www.cs.utexas.edu/~AustinVilla/?p=athome>

Hardware For Domestic Standard Platform League we will be using a Toyota Human Support Robot (HSR) to be furnished by Toyota.

External Devices

- Google Speech API
- Alienware Alpha R2 (Computer)
- Alienware 17 Laptop (Backpack)

Third-Party Software

- Clingo
- MoveIt
- Robot Operating System (ROS)

Robot HSR Hardware Description

Specifications are as follows:

- Name: Human Support Robot (HSR)
- Footprint: 43 cm
- Robot dimensions: height: 1.35m (max)
- Robot weight: 37 kg.

