

Team ORIon

2019 DSPL Team Description

Ioannis Havoutis, Nick Hawes, Lars Kunze, Bruno Lacerda, Chia-Man Hung,
Mark Finean, Denis Koksal-Rivet, and Stephen Kyberd

Oxford Robotics Institute, University of Oxford, UK,
ioannis@robots.ox.ac.uk
<http://ori.ox.ac.uk/robocup/>

Abstract. This document outlines the approach *Team ORIon* will take to the RoboCup@Home Domestic Standard Platform League (DSPL). We are a new team with a strong research background and we aspire to compete in this year’s DSPL. Our first experience competing with the Toyota Human Support Robot was at the World Robot Summit (WRS) last October, where we competed in the Partner Robot Challenge. Our research interests are centred around long-term autonomy, mobility, robot learning and knowledge representation. Advances in these directions will enable service robots to interact with humans and complete useful everyday tasks in typical household settings. We aim to demonstrate robust and intelligent autonomous behaviour, that uses experience to learn and refine a growing set of robot skills, on the Toyota HSR.

1 Introduction

Team ORIon is a new team created within the Oxford Robotics Institute (ORI) at the University of Oxford. The team consists of undergraduate and graduate students; robotics researchers; and faculty members of ORI. We come from a strong research institute with seminal work in mobile autonomy and machine learning. ORI has a significant track record in *field robotics* and real-world trials of autonomous systems. It also has a team of professional hardware and software engineers. This experience and support is leveraged to create a RoboCup@Home team capable of delivering across the whole competition.

The Domestic Standard Platform League (DSPL) affords a tangible new domain in which existing ORI research can be applied, and which provides new challenges for the group. We successfully applied for a Toyota Human Support Robot (HSR) last year (2018) and we have received the robot in late July. Despite the short amount of time we had with the HSR, we succeed in competing in the Partner Robot Challenge of this year’s World Robot Summit (WRS) in October (see Figure 1). The HSR allows us to focus on developing the intelligence required for successfully completing the RoboCup@Home tasks, without the added burden of building and maintaining a custom platform. Competing in the 2019 DSPL event in Sydney will allow us to demonstrate the capabilities presented in this document and will provide valuable experience for our future RoboCup@Home participation.



Fig. 1. *Left:* Team ORion at the World Robot Summit (WRS) 2018 in Tokyo, Japan. *Middle:* Setup for the “Bring me” task. *Right:* HSR grasping a plant.




Fig. 2. *Left:* The STRANDS autonomous mobile robot in a real-world office environment. *Right:* View planning for object detection in the office environment.

2 Team Composition

The team is led by Prof. Nick Hawes, who has extensive background in intelligent autonomous robots that can work with or for humans, and Dr. Ioannis Havoutis, an expert in combining motion planning with machine learning. The core of the team for 2019 will be ORI post-doctoral researchers (Lars Kunze, Bruno Lacerda) and PhD students (Chia-Man Hung, Mark Finean, Rowan Border), with more junior members (Denis Koksal-Rivet, undergraduates) being brought in as the team develops.

3 Capabilities and Goals

The capabilities of our DSPL system builds upon the capabilities developed over the last four years within the EU STRANDS Project¹. Key members of this project (Nick Hawes, Lars Kunze, Bruno Lacerda) are part of Team ORion. The STRANDS Project deployed autonomous mobile robots (MetraLabs SC-ITOS A5, see Fig. 2) in a range of human-populated environments for long durations [1]. These robots provided a range of services to real users, similar to the tasks required in the DSPL. The enabling software used on the robots, the ROS-based *STRANDS Core System* (SCS), therefore gives Team ORion an ideal basis for DSPL development. The SCS is open source, and Team ORion

¹ <http://strands-project.eu>



Fig. 3. Snapshots of the Baxter robot performing a water pouring task that is learnt from demonstration [8]. The probabilistic encoding captures the correlation among task variables and produces a controller that generalizes the behaviour.

will contribute to the continued maintenance of this substantial code base which is useful for our entire community. Although originally developed for MetraLabs robots, the SCS has recently been ported to other robots, and Team ORIon will contribute an open port of this software to the Toyota HSR.

The SCS builds upon standard ROS components to provide the following capabilities, all of which have been tested in long-term deployments in real user environments: topological, human-aware robust navigation; object detection, identification and classification; autonomous online object learning; human detection, skeleton tracking, and activity analysis; basic human-robot interaction via speech and screen; and goal management and task planning. Our capabilities for person tracking [2] can be seen online² and formed the basis of many interactive behaviours, including social navigation, and activity learning and recognition [3], which are relevant to DSPL.

3.1 Manipulation – Learning new skills

In the context of manipulation, the robot will require a number of key skills to successfully perform a wide variety of tasks that involve interaction with the environment or other agents (Figure 3), eg. pushing buttons, tuning handles, grasping and passing objects, etc. The robots used in the STRANDS Project did not have manipulation capabilities, therefore the SCS does not provide software to support this. To deliver these capabilities we have started from Team ORIon member Lars Kunze’s previous experience of knowledge-enabled manipulation [4]. This previous work resulted in a system which could grasp an egg (<https://www.youtube.com/watch?v=jLz87H4q3hU>) and make a pancake (<https://www.youtube.com/watch?v=YQs5gRei8k4>). To augment this we aim to build in our framework the ability to learn and refine new skills as tasks change or as new tasks need to be added the task repertoire. Such capability will be based on Team ORIon member Ioannis Havoutis’ background in learning, synthesis and control of complex motions [5]. Skill representations are learnt from demonstrations—allowing also the use by non-experts—using a probabilistic generative encoding [6]. Motion generation is formulated as an optimal control problem that adapts to changing task configurations on-line [7,8] (<https://youtu.be/NiRPEOegymk>).

² <https://youtu.be/zdnvhQU1YNo>

3.2 Navigation

To enable robust navigation in all settings we take a hierarchical approach to navigation. The hierarchy is structured around a topological map in which discrete locations are connected by directed edges [9]. Edges correspond to navigation actions the robot can perform to transition between locations. These may be standard `move_base` actions, social navigation, closed loop controllers (such as wall following or door passing), or teach-and-repeat paths. Choices between the actions are made by a Markov decision process-based planner which jointly optimises for success probability and completion time, using probabilistic models learnt online through experience [10]. To ensure the robot does not get stuck we employ a monitored navigation layer which monitors the execution of the low level edge actions and performs recovery behaviours (e.g. backtracking, HRI) to correct observed problems [1]. This collection of techniques drove the STRANDS robots for over 360km of autonomous navigation in human-populated environments. For Team ORION we will extend the framework to enable integration of ORI’s visual teach-and-repeat paradigm, to enable the robot to navigate in areas where laser-based localisation is likely to result in imprecise navigation. We will also look to integrate some of ORI’s previous 3D mapping work (e.g. [11]) to increase the accuracy of the robot’s environment representation.

We have already integrated most parts of the hierarchical topological navigation to our software stack. We have used this to robustly navigate the arena of the Partner Robot Challenge arena at WRS, where our HSR needed to reach different parts of the simulated house environment and stop before a closed door is opened.

3.3 Semantic Vision

Learning and recognising objects during operation is a key task for a mobile service robot in human environments. Team ORION will exploit the work done in the STRANDS Project in terms of autonomous object learning plus the recognition and modelling of previously unseen objects. The work is based on the *meta-room* approach which builds dense RGB-D reconstructions of regions around locations in the robot’s topological map. Objects are found through inspecting or differencing meta-rooms. Surfaces and possible objects found in meta-rooms become either targets for more detailed view planning [12] (see Figure 2) leading to recognition, or for autonomous object learning [13]. Our recognition pipeline mixes top-down semantic reasoning with bottom-up appearance-based processing for scene understanding [14], jointly estimating object locations and categories based on qualitative spatial models [15]. The object learning process can build detailed 3D models entire without supervision [13]. Previously unknown objects are processed with a mix of deep vision and semantic web technologies to provide the robot with an initial estimate of their identity [16].

3.4 Robot System Integration

A substantial effort will be needed to develop behaviours that are robust to changes in the environment and to noise typical of real-world scenarios. In this respect we will exploit our experience from the STANDS project [1] and build upon the tested SCS. Some of this software has already been integrated to our stack for the needs of the Partner Robot Challenge of WRS. Given the similarity of the HSR to the SCITOS platform, and other platforms the team are familiar with, we predict that porting the SCS to the Toyota HSR will be an easy task. SCS was developed to be expandable and is built with standard ROS components which are also supported by the Toyota HSR. Additionally, we are planning to build a mock-up of the “house” arena to allow us to run live robot trials and limit simulation use to the development phase. We aim to schedule recurring trials as our framework is developed, to ensure that robot behaviours are successful and to collect data on the success probabilities of tasks and sequences of tasks.

Team ORIon benefits from the many years of experience of the team of creating integrated robot systems. Team members contributed to the first public demonstration of a self-driving car in the UK³, and all members have contributed to integrated robot systems demonstrated at science museums, public engagement events and trade shows across Europe. All of these systems integrate perception, planning and action in non-trivial ways. Such integration is central to producing a functional and reliable system, but can be incredibly challenging when trying to produce novel capabilities for robots in task environments which you are only able to experience a short time before a deadline. We had our first competition just last October in the Partner Robot Challenge of WRS. Most of the team members were part of this experience and now have demonstrated our team’s ability to develop, integrate and test solutions on a very short time schedule; we achieved in competing given only 3 weeks of preparation. The team’s joint experience of bringing diverse robot capabilities together for successful demos will enable the team to start working effectively very quickly, and to deal with common team and system teething problems smoothly. Our experience on systems which span the capability spectrum from low-level sensing to high-level cognition means that the diverse capabilities described above will be successfully integrated to produce a competitive entry in the 2019 DSPL Sydney.

4 Applicability and Re-Usability of the System

The continued maintenance and development of the SCS will provide a well-tested software framework for mobile service robots. Continuing the practice started by the STRANDS Project, we will make extensions to the SCS available as open source software. This will enable our core framework to be reusable by other groups. The validity of this approach has already been demonstrated by the reuse of the SCS at labs including the Intelligent Robots and Systems group at the Institute for Systems and Robotics, Lisbon, and at Honda Research Institute

³ <https://www.epsrc.ac.uk/newsevents/news/lutzpathfinder/>

Europe. The aforementioned use of the majority of our technology within systems which have already been successfully demonstrated in real challenging service robot environments shows that our approach is applicable in the real world.

5 Current software stack

- **Topological navigation** - The topological navigation system from the STRANDS project was integrated into the HSR as the main navigation system. The system allows for the placement of predefined nodes in our map and for the creation of edges between any of the nodes. The system also allows for the specification of actions on each edges and nodes, allowing the user to define any number of actions to take upon reaching a node, or whilst navigating to a node; this allows us to use the move_base configurations specified by Toyota, or to define our own navigation with optimized parameters.
- **Human-Robot Interaction via Speech** - For interacting with the robot using Natural Language (via speech commands) we have integrated Julius, an open-source large vocabulary continuous speech recognition engine⁴. To this end, we have defined a simple grammar that allows an operator to communicate with the robot in predefined dialogues. This allows an operator to issue commands (“tidy up, bring me *something*”, “move to *location*”) and/or confirm requests made by the robots (“*Please open the door and confirm when you are done*”).
- **Object detection** - For object detection in 2D we use a trained deep neural network, *YOLO*. We decided to start with YOLO because of the simplicity of the framework and fast detection speeds with deep models reported in the literature. Image coordinates of detected objects are transformed into 3D and into object markers, using simple processing and ROS’ *tf*.
- **Manipulation and grasping** - In our current software stack, grasping is realised via the ROS motion planning framework *MoveIt*⁵. A target frame is provided by our object detection system, while MoveIt is used for motion planning and grasping.

6 Conclusion

Team ORIon is a new DSPL team that builds on the strong research background of its team members and on extensive real-world robot operating experience, in a range of service tasks, similar to the DSPL.

Competing in the 2019 DSPL event in Sydney will allow us to demonstrate the presented capabilities on the HSR and will provide valuable experience for our newly-born team.

⁴ <https://github.com/julius-speech/julius>

⁵ <https://moveit.ros.org>

References

1. Nick Hawes, Chris Burbridge, Ferdian Jovan, Lars Kunze, Bruno Lacerda, Lenka Mudrová, Jay Young, Jeremy L. Wyatt, Denise Hebesberger, Tobias Körtner, Rares Ambrus, Nils Bore, John Folkesson, Patric Jensfelt, Lucas Beyer, Alexander Hermans, Bastian Leibe, Aitor Aldoma, Thomas Faulhammer, Michael Zillich, Markus Vincze, Muhannad Al-Omari, Eris Chinellato, Paul Duckworth, Yiannis Gatsoulis, David C. Hogg, Anthony G. Cohn, Christian Dondrup, Jaime Pulido Fentanes, Tomás Krajník, João M. Santos, Tom Duckett, and Marc Hanheide. The STRANDS project: Long-term autonomy in everyday environments. *IEEE Robotics & Automation Magazine (RAM)*, 24(3), September 2017.
2. Christian Dondrup, Nicola Bellotto, Ferdian Jovan, and Marc Hanheide. Real-time multisensor people tracking for human-robot spatial interaction. In *International Conference on Robotics and Automation (ICRA) - Workshop on Machine Learning for Social Robotics*, 2015.
3. P. Duckworth, Y. Gatsoulis, F. Jovan, N. Hawes, D. C. Hogg, and A. G. Cohn. Unsupervised learning of qualitative motion behaviours by a mobile robot. In *Proc. of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016)*, Singapore, 2016.
4. Lars Kunze and Michael Beetz. Envisioning the qualitative effects of robot manipulation actions using simulation-based projections. *Artificial Intelligence, Special Issue on AI and Robotic*, 2015.
5. I. Havoutis and S. Calinon. Learning assistive teleoperation behaviors from demonstration. In *IEEE International Symposium on Safety, Security and Rescue Robotics (SSRR)*, Lausanne, Switzerland, October 2016.
6. I. Havoutis and S. Calinon. Supervisory teleoperation with online learning and optimal control. In *Proc. IEEE Intl Conf. on Robotics and Automation (ICRA)*, pages 1534–1540, Singapore, May-June 2017.
7. M. J. A. Zeestraten, I. Havoutis, S. Calinon, and D. G. Caldwell. Learning task-space synergies using riemannian geometry. In *Proc. IEEE/RSJ Intl Conf. on Intelligent Robots and Systems (IROS)*, pages 73–78, Vancouver, Canada, September 2017.
8. M. J. A. Zeestraten, I. Havoutis, J. Silvério, S. Calinon, and D. G. Caldwell. An approach for imitation learning on Riemannian manifolds. *IEEE Robotics and Automation Letters (RA-L)*, 2(3):1240–1247, June 2017.
9. Jaime Pulido Fentanes, Bruno Lacerda, Tomáš Krajník, Nick Hawes, and Marc Hanheide. Now or later? predicting and maximising success of navigation actions from long-term experience. In *International Conference on Robotics and Automation (ICRA)*, 2015.
10. B. Lacerda, D. Parker, and N. Hawes. Optimal and dynamic planning for markov decision processes with co-safe ltl specifications. In *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'14)*. IEEE, September 2014.
11. Paul Amayo, Pedro Piniés, Lina Maria Paz, and Paul Newman. A unified representation for application of architectural constraints in large-scale mapping. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Stockholm, Sweden, May 2016.
12. Lars Kunze, Keerthi Kumar Doreswamy, and Nick Hawes. Using qualitative spatial relations for indirect object search. In *IEEE International Conference on Robotics and Automation (ICRA 2014)*, Hong Kong, China, May 31 - June 7 2014.

13. T. Faeulhammer, R. Ambrus, C. Burbridge, M. Zillich, J. Folkesson, N. Hawes, P. Jensfelt, and M. Vincze. Autonomous learning of object models on a mobile robot. *IEEE Robotics and Automation Letters (RA-L)*, 2(1):26 – 33, 2016.
14. Lars Kunze, Chris Burbridge, Marina Alberti, Akshaya Tippur, John Folkesson, Patric Jensfelt, and Nick Hawes. Combining top-down spatial reasoning and bottom-up object class recognition for scene understanding. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Chicago, Illinois, US, September, 14–18 2014.
15. Lars Kunze, Chris Burbridge, and Nick Hawes. Bootstrapping probabilistic models of qualitative spatial relations for active visual object search. In *AAAI Spring Symposium 2014 on Qualitative Representations for Robots*, Stanford University in Palo Alto, California, US, March, 24–26 2014.
16. Jay Young, Lars Kunze, Valerio Basile, Elena Cabrio, and Nick Hawes. Semantic web-mining and deep vision for lifelong object discovery. In *IEEE International Conference on Robotics and Automation (ICRA)*, Singapore, May 29–June 3 2017.