



UNIVERSITY OF LEEDS

This is a repository copy of *The STRANDS Project: Long-Term Autonomy in Everyday Environments*.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/109414/>

Version: Accepted Version

Article:

Hawes, N, Burbridge, C, Jovan, F et al. (30 more authors) (2017) The STRANDS Project: Long-Term Autonomy in Everyday Environments. *IEEE Robotics and Automation Magazine*, 24 (3). pp. 146-156. ISSN 1070-9932

<https://doi.org/10.1109/MRA.2016.2636359>

© 2017 IEEE. This is an author produced version of a paper published in *IEEE Robotics and Automation Magazine*. Personal use of this material is permitted. Permission from IEEE must be obtained for all other users, including reprinting/ republishing this material for advertising or promotional purposes, creating new collective works for resale or redistribution to servers or lists, or reuse of any copyrighted components of this work in other works. Uploaded in accordance with the publisher's self-archiving policy.

Reuse

Unless indicated otherwise, fulltext items are protected by copyright with all rights reserved. The copyright exception in section 29 of the Copyright, Designs and Patents Act 1988 allows the making of a single copy solely for the purpose of non-commercial research or private study within the limits of fair dealing. The publisher or other rights-holder may allow further reproduction and re-use of this version - refer to the White Rose Research Online record for this item. Where records identify the publisher as the copyright holder, users can verify any specific terms of use on the publisher's website.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

The STRANDS Project: Long-Term Autonomy in Everyday Environments

Nick Hawes¹, Chris Burbridge¹, Ferdian Jovan¹, Lars Kunze¹, Bruno Lacerda¹, Lenka Mudrová¹, Jay Young¹, Jeremy Wyatt¹, Denise Hebesberger², Tobias Körtner², Rares Ambrus³, Nils Bore³, John Folkesson³, Patric Jensfelt³, Lucas Beyer⁴, Alexander Hermans⁴, Bastian Leibe⁴, Aitor Aldoma⁵, Thomas Fäulhammer⁵, Michael Zillich⁵, Markus Vincze⁵, Eris Chinellato⁶, Muhannad Al-Omari⁷, Paul Duckworth⁷, Yiannis Gatsoulis⁷, David C. Hogg⁷, Anthony G. Cohn⁷, Christian Dondrup⁸, Jaime Pulido Fentanes⁸, Tomáš Krajník⁸, João M. Santos⁸, Tom Duckett⁸ and Marc Hanheide⁸

¹Intelligent Robotics Lab, School of Computer Science, University of Birmingham, UK

²Akademie Fur Altersforschung Am Haus Der Barmherzigkeit, Austria; and
Donau-Universitaet Krems, Austria

³Centre for Autonomous Systems, KTH Royal Institute of Technology, SE-100 44
Stockholm, Sweden

⁴Rheinisch-Westfälische Technische Hochschule Aachen, Germany

⁵Technische Universität Wien, Austria

⁶Faculty of Science and Technology, Middlesex University London, UK

⁷University of Leeds, UK

⁸LCAS, University of Lincoln, UK

October 17, 2016

1 Introduction

Thanks to the efforts of the robotics and autonomous systems community, robots are becoming ever more capable. There is also an increasing demand from end-users for autonomous service robots that can operate in real environments for extended periods. In the STRANDS project¹ we are tackling this demand head-on by integrating state-of-the-art artificial intelligence and robotics research into mobile service robots, and deploying these systems for long-term installations in security and care environments. Over four deployments, our robots have been operational for a combined duration of 104 days autonomously performing end-user defined tasks, covering 116km in the process. In this article we describe the approach we have used to enable long-term autonomous operation in everyday environments, and how our robots are able to use their long run times to improve their own performance.

2 Long-Term Autonomy in STRANDS

Autonomous robots come in a range of forms, for a range of applications. Across this range, long-term autonomy (LTA) has a variety of meanings. For example, NASA's Opportunity rover has been autonomous for over 10 years on the surface of Mars; wave gliders can autonomously monitor stretches of ocean for months at a time; and autonomous cars have completed journeys of thousands of kilometres. In this article we restrict our contributions to *mobile robots operating in everyday, indoor environments* (e.g. offices, hospitals), capable of performing *a variety of service tasks*. Across all the aforementioned robots there are commonalities in low-level, short-term control algorithms (e.g. closed-loop motor control). Beyond this, the algorithms used to provide long-term, task-specific autonomous capabilities, and the hardware

¹Spatio-Temporal Representations and Activities for Cognitive Control in Long-Term Scenarios, <http://strands-project.eu>.

these algorithms control, varies greatly, according to application and environmental requirements. The challenges that distinguish indoor service robots from the aforementioned examples relate to both their environment and their task capabilities. Indoor task environments are less physically risky than outdoor environments, but have a comparatively higher degree of short- to medium-term physical variability, e.g. people, doors and furniture moving (roads are similar, but traffic movement is generally more predictable and less frequently occluded). In terms of application requirements, multi-purpose service robots must be capable of predictable scheduled behaviour whilst also being retaskable on-demand with high availability, and must be able to navigate in relatively confined, dynamics environments. This is in contrast to the largely restricted-purpose systems mentioned above. Taken together the set of requirements for indoor service robots presents unique challenges, and thus LTA in this context warrants dedicated research.

Given the state of the art, we consider “long-term” for a mobile service robot to be at least multiple weeks of *continuous* operation. In very general terms, such LTA operation requires that a robot’s hardware and software is robust enough to failure to enable such operation. Such robustness can be provided by both design-time and run-time approaches. It is essential that LTA systems are able to actively manage consumable resources (e.g. battery) and that any autonomy-supporting capabilities (e.g. localisation) are not adversely affected by long run times. Whilst this latter point is common sense, and common practice in many other technologies (from operating systems to cars), it has only recently been considered in autonomous robotics.

One reason it is challenging to design a service robot to meet the requirements of LTA is the impossibility of anticipating all the situations in which it may find itself. However, if we can enable robots to run for long periods, then they will have opportunities to learn about the structure and dynamics of such situations. By exploiting the results of such learning, the robots should be able to increase their robustness further, leading to a virtuous cycle of improved performance and greater autonomy. It is this latter point which motivates STRANDS: to go beyond robots which simply survive, to those that can improve their performance in the long term. It is in this context that this article makes its main contribution: a robotic software architecture (the STRANDS Core System) which was designed for LTA service robot applications, and evaluated across four end-user deployments. It contains a mix of common sense and novel elements which have enabled it to support over 100 days of autonomous operation. This article is the first time all of these elements have been presented together, and contains the first presentation of metrics describing performance across deployments. Our approach is inspired by the work of Willow Garage [1] and the CoBot project [2], plus the pioneering work on systems like Rhino and Minerva (e.g. [3]). What distinguishes our work from these is the combination of multiple service capabilities, in a single system capable of weeks or more of continuous autonomous operation, in dynamic indoor environments, whilst using various forms of learning to improve system performance. Many other projects address one or two of these elements, but not all four simultaneously.

3 Application Scenarios

To ensure our research is able meet the demands of end users, our work is evaluated in two application scenarios: security and care. Space does not permit a detailed explanation of the tasks in each scenario. Instead we include citations to further information on the tasks and technology from each scenario.

Our security scenario is developed with G4S Technology. The aim of this scenario is to have a robot monitoring an indoor office environment, generating alerts when it observes prohibited or unusual events. To date we have completed two security deployments in which a mobile robot routinely created models of the environment’s 3D structure [4], objects [5] and people [6]; modelling their changes over time; and using these models to detect anomalous situations and patterns. For example, we have developed robot behaviours to: detect when a human moves through the environment in an unusual manner [6]; build models of the arrangement of objects on desks [7]; and check whether fire exits have been left open. Long-term deployments are essential for these services in order to gather sufficient data to build appropriate models.

Our care scenario is developed with the Akademie für Alterforschung at the Haus der Barmherzigkeit (HdB). In this scenario, the robot supports staff and patients in a large elderly care facility. To date we have completed two care deployments in which a mobile robot: guides visitors; provides information to residents; and assists in walking-based therapies. In the care scenario the robot serves users more directly, and therefore long-term system robustness is crucial, as is adapting to the routines of the facility. For more information on this scenario see [8, 9].



Figure 1: Two of the STRANDS MetraLabs SCITOS A5s in their application environments. On the left is the robot *Bob* at G4S's Challenge House in Tewkesbury, UK. On the right is the robot *Henry* in the reception of Haus der Barmherzigkeit, Vienna.

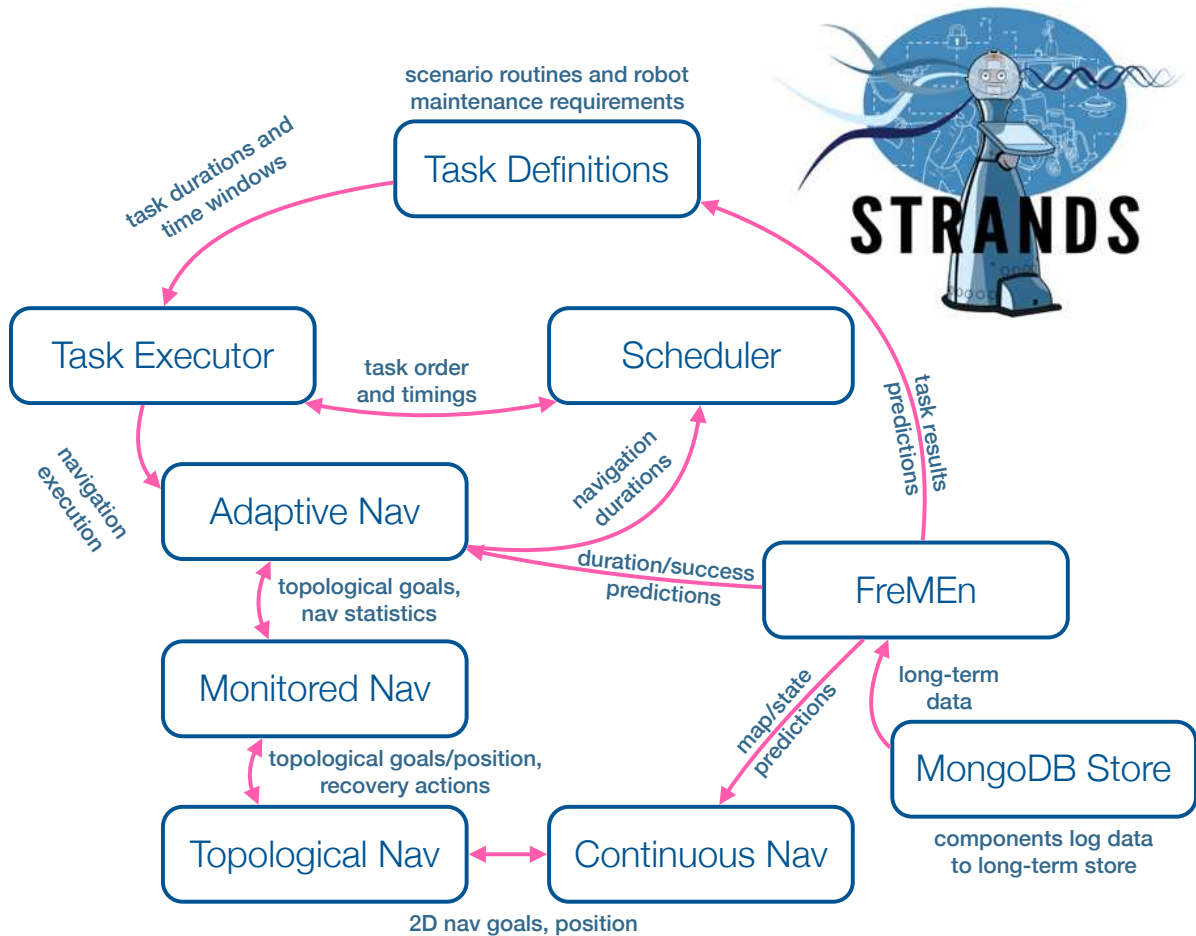


Figure 2: Schematic overview of the Core STRANDS System.

4 Robot Technology

The systems reported in this paper are developed in ROS, available under open source licenses, and binary packaged for Ubuntu LTS.² Whilst the majority of our work is platform neutral, all our deployed systems are based on the MetraLabs SCITOS A5 (see Figure 1). This is an industry-standard mobile robot capable of the long run times (12 hours on one charge) and autonomous charging. Our robots each have SICK S300 lasers in their bases (for localisation, leg detection etc.), and two Asus Xtion PRO RGB-D cameras: one at chest height pointing downward (for obstacle avoidance), the other on a pan-tilt unit (PTU) above the robot’s head. The SCITOS has an embedded Intel Core i7 PC with 8GB RAM to which we have networked two additional PCs each with an i7 and 16 GB RAM.

5 The Core STRANDS System

The *STRANDS Core System* (Figure 2) is an application-neutral architecture for LTA in mobile robots. It is a mix of widely-used components, plus components designed specifically for LTA. As mentioned above, hardware and software robustness is essential for LTA. Hardware robustness is beyond the scope of our research, thus we assume our software is running on an appropriate robot and computational platform. We address software component robustness through a mix of strategies. During development we encourage components to be designed in a way that makes the minimum assumptions about the existence of other components and services (e.g. by checking service existence before running). We also pay particular to error handling to ensure component-local errors and exceptions do not propagate unnecessarily. This allows components, and whole subsystems, to be brought up and down automatically.

²See <http://strands-project.eu/software.html>.

At run-time we use built-in ROS functionality to automatically relaunch crashed components, and try to run most subsystems only when required (saving CPU and power, and reducing opportunities for errors). We also use run-time topic monitoring to detect problems (e.g. low publish rates) and trigger component restarts. Finally, we run a continuous integration server that tests components and the whole system in isolation, on recorded data, and in simulation. The rest of this section summarises the STRANDS Core System, and provides references to additional technical details.

The overall performance of a mobile robot is constrained by its localisation and navigation systems, so we use widely-adopted ROS packages to provide state-of-the-art performance. When deploying we build a fixed map from laser, localise in it using adaptive Monte Carlo localisation, and navigate using the dynamic window approach (DWA) over 3D obstacle information.³ Whilst our use of a fixed map appears at odds with LTA in a dynamic environment, our environments are dominated by static features (e.g. walls), which prevent the robot’s localisation performance from degrading. We also take advantage of the the robot regularly docking with a charging station by resetting the robot’s position to this known location whilst docked. This limits localisation drift to that which can occur during time away from the dock.

We manually build a topological map on top of the fixed continuous map. We place topological nodes at key places in the environment for navigation (e.g. either side of a door) or for tasks (e.g. by a desk to observe). The topological map from our 2015 security deployment is in Figure 6. Edges in the topological map are parametrised by the action required to move along them. In addition to DWA navigation, our system can perform door passing, docking on to a charging station, and adaptive navigation near humans [10].

In our experience, navigation performance is major determiner of the autonomous run time of a mobile robot. This is because navigation failures (e.g. getting stuck near obstacles) can result in the robot being unable to return to its charging station. Thus the aforementioned elements of the STRANDS Core System support LTA in the following ways. First, by constraining the robot’s movements to the topological map we are able to restrict navigation to known good areas of the environment. We additionally restrict movement by marking areas of the static map as ‘no go’ zones which cannot be planned through. Despite these restrictions, navigation failures still occur due to environmental dynamics (e.g. people walking in front of the robot). Therefore edge traversals in the topological map are executed by a *monitored navigation* layer that can perform a range of recovery actions in the event of failure (see Section 7). Also, topological route planning and execution is one place where our core system adapts to long-term experience, as described in Section 8.

The main unit of behaviour in our system is a *task*. Tasks represent something the robot can do (e.g. check whether a fire door is open, serve information via a GUI), and have an associated topological location, a maximum duration, and a time window for execution. Our executive framework [11] schedules tasks to be executed within their time windows, and manages task-directed navigation then execution. To prevent task failures from interfering with long-term operation, our framework detects task time-outs and failures, then stops or restarts robot behaviours as necessary. Maintenance actions such as charging, batch learning and database backups are all handled as tasks, allowing the executive framework control of most of the robot’s behaviour. This is essential for LTA as it enables the system to actively manage its limited resources. A plot of tasks from the 2015 security deployment can be seen in Figure 3.

Our system relies on separate pipelines for perceiving different elements of its environment: real-time multi-person RGB-D detection and tracking [12]; visual object instance and category modelling and recognition [13]; and 3D spatio-temporal mapping [4]. This article does not cover our work on perceptually challenging tasks. Instead we refer readers to other papers where we have exploited these perception pipelines, e.g. [10, 5, 7].

The data observed and generated (e.g. as inter-component communication) by an LTA system is crucial for both learning, and for monitoring and debugging the system. We therefore use tools based on MongoDB⁴ to save ROS messages to a document-oriented database. Database contents (e.g. observations of doors being opened or closed) can then be interpreted by the Frequency Map Enhancement (FreME) component [14], which integrates sparse and irregular observations into spatio-temporal models representing (pseudo-)periodic environment variations. These can be used to predict future environment states (see Section 8).

³See <http://wiki.ros.org/navigation> for details on these techniques.

⁴http://wiki.ros.org/mongodb_store




Figure 3: A plot of the tasks performed by the robot during the 2015 security deployment. White space indicates that the robot is not performing any tasks. This indicates that the robot is charging or a failure has occurred.

6 Metrics

So far we have performed two evaluation deployments for each of the security and care scenarios. For each deployment we monitored overall system performance against two metrics: *total system lifetime* (TSL), and *autonomy percentage* (A%). TSL measures how long the system is available for autonomous operation, and is reset if the system experiences an unrecoverable failure, or needs an unrequested expert intervention (i.e. something which cannot easily be done by an end-user on site). A% measures the duration the system was actively performing tasks as a proportion of the time it was allowed to operate autonomously (which in our case is typically restricted to office hours). The motivation of A% is that it is trivial to achieve a long TSL if the system does nothing. However, neither TSL nor A% measure the quality of the services being provided. As this article focuses on LTA we restrict our presentation to the aforementioned, task-neutral but LTA-specific metrics. End-user evaluations of our systems’ task-specific performance are ongoing, and will be published in the future (see [9, 8] for early evaluations from the care scenario).

| | Care 2014 | Security 2014 | Care 2015 | Security 2015 | Total |
|----------------------------|-----------|---------------|-----------|---------------|----------|
| Total Distance Travelled | 27.94km | 20.64km | 23.41km | 44.25km | 116.24km |
| Total Tasks Completed | 1985 | 963 | 865 | 4631 | 8444 |
| Max TSL | 7d 3h | 6d 19h | 15d 6h | 28d 0h | |
| Cumulative TSL | 20d 19h | 21d 0h | 27d 8h | 35d 3h | 104d 7h |
| Individual Continuous Runs | 18 | 18 | 5 | 2 | 43 |
| Autonomy Percentage (A%) | 38.80% | 18.27% | 53.51% | 51.10% | |

Table 1: LTA metrics from the first four STRANDS system deployments.

Table 1 presents our systems’ LTA performance so far. In 2014 we aimed for 15 days TSL. However the longest run we achieved was seven days. Most of our system failures were caused by the lack of robustness of our initial software, leading to unrecoverable component behaviour (crashes or deadlock states). This was fixed for our 2015 deployments by following the development approaches outlined in Section 5. In 2015 we targeted 30 days TSL, coming close with 28 days in the security deployment. This long run was terminated by the robot’s motors not responding to commands, an issue which has since been fixed by a firmware update. In the 2015 deployments, most failures were due to computer-related issues beyond the direct contributions of the project (e.g. USB drivers, power cables, network problems etc.). Of the seven runs in 2015, one run was ended due to user intervention (a decorator powered off the robot), two due to bugs in our software, and the remaining four due to faults in software or hardware beyond our components.

The variations across deployments in terms of number of tasks completed and distance travelled were largely down to the different types of tasks performed by the robots, and the different environments they were deployed in. For example, information serving tasks may take tens of minutes with very little travel, but door checking tasks will be brief and will also require the robot to travel both before and during the task.

Systems in the literature have delivered more autonomous time and distance *cumulatively* (i.e. accumulated across multiple robots and/or system runs), but we believe the 28 day run is the longest a single continuous autonomous run of an indoor mobile service robot capable of multiple tasks. The most relevant comparison we can make is to the CoBots. The CoBot analysis in [2] reports a total of 1,279.5 hours of autonomy time, traversing 1,006.1km. This was achieved by four robots in 3,199 separate continuous autonomous runs over three years, at an average of 0.31km, 23 minutes per run. They do not report the longest single continuous run (either in time or distance), but even an extremely long run for a CoBot would only be measured in hours, not days (as they don’t have autonomous charging capabilities). In contrast the STRANDS systems performed a total of 43 separate continuous runs, yielding a total of 2,545 hours and 116 km over the four deployments, at an average of 2.7km and 58:12 hours per run. How the duration of individual runs varied can be seen in Figure 4. Note that we use this data to provide a point of comparison. The two projects are targeting different metrics (total distance for CoBots, single run duration for STRANDS) thus the systems naturally have different performance characteristics.

Sections 7 and 8 describe novel elements of our system that have enabled such long run times everyday environments. These are followed by examples of tasks that exploit these long run times to improve robot service performance.




Figure 4: A histogram of individual continuous run lengths over the 4 STRANDS deployments.




Figure 5: Per-recovery counts for our 2015 security (left) and care (right) deployments.

7 Monitored Navigation

Given the huge variety of situations an LTA service robot will encounter, it is impossible to develop a navigation algorithm to successfully deal with all of them. We therefore developed a framework that executes topological navigation actions and monitors them for failure. If a failure is detected, then the framework iterates through a list of recovery behaviours until either the navigation action completes successfully, or the list is exhausted (in which case failure is reported back to the calling component). Failure types can be mapped to specific lists of recoveries. When the robot’s bumper is pressed, a hardware cut-off prevents it driving, therefore in this case the robot must ask to be pushed away from obstructions by nearby humans. If the local DWA planner fails to find a path, then simply clearing the navigation costmap (to remove transient obstacles) may suffice. We also developed a backtrack behaviour which uses the PTU-mounted depth camera to sense backwards whilst reversing along the path it took to the failure location. This is triggered when navigation fails, but clearing the costmap does not overcome the failure.

| Failure | Recoveries | Successful | Unsuccessful | Total |
|--|---|------------|--------------|-------|
| Bumper pressed | Request help via screen and voice. Repeated until recovered. | 177 | 148 | 325 |
| Navigation failure (no valid local or global path) | Sleep then retry; backtrack to last good pose; request help via screen and voice. Repeated request until recovered. | 707 | 993 | 1700 |
| Stuck on carpet | Increased velocities commanded to motors | 16 | 247 | 263 |

Table 2: Classes of navigation failure, their associated recoveries, and the overall counts of successful and unsuccessful recoveries from these failures. Per-recovery counts are show in Figure 5

Table 2 presents the recovery behaviours used in our 2015 deployments. Successful recoveries are those which are not followed by another failure within one minute or one metre of travel, otherwise they are unsuccessful. A successful recovery may be preceded by any number of unsuccessful recoveries. A sequence of unsuccessful recoveries can come from the monitored navigation system as it attempts recoveries that then fail, or from the task execution framework unsuccessfully trying to navigate the robot to another task after a previous failure. Figure 6 shows where all the successful recoveries from our 2015 security deployment occurred. They are largely clustered around areas where it was difficult to navigate, such as near doors, and close to desks. This novel approach significantly contributed to the LTA performance of our systems, as each recovered failure could have potentially caused the end of a continuous run.

8 Adaptive Topological Navigation

Whilst monitored navigation helps the robot recover from navigation failures, it does not help it avoid them. To do this we aggregate the robot’s navigation experience into a Markov Decision Process (MDP)




Figure 6: The map of the deployment area in Challenge House, Tewkesbury with the topological map superimposed. Also displayed are the locations where the robot successfully recovered from a navigation failure. Locations where the bumper was triggered are red. The robot asked humans for help at these locations. It also did this at locations marked with green (for non-bumper fails). Places where recoveries were performed by reversing along the previous path are marked in yellow, and by simply retrying in blue.

automatically built from the topological map [15]. Using an MDP allows the system to model uncertainty over the success of the robot traversing an edge in the map and its expected duration. By learning models for these success probabilities and durations online, the robot is able to continually adapt its behaviour to the environment it is deployed in. Every time the robot navigates an edge, the duration and success of the traversal is logged to the robot’s database. These logs are processed by FreMEn (see Section 5) to produce a temporal predictive model that allows the actions of the MDP to be assigned probabilities and travel durations appropriate for the time of execution [11]. This MDP is then solved for a target location to produce a policy for topological navigation that prefers low duration edges with high success probabilities (see [15] for details). This improves the system’s robustness by making it avoid areas where it previously encountered navigation failures. This is only possible in an LTA setting where the robot runs repeatedly in the same environment.

9 Predicting Human-Robot Interaction

In the HdB care facility our robot acts as an information terminal, using its touch screen to present the date, daily menu, news etc., to staff and to residents with potentially severe dementia. This behaviour is scheduled as a task at different topological nodes in the care home. As we did not know in advance the locations and times people would prefer to interact with the robot, we allowed it to adapt its routine based on long-term experience. To achieve this, each node in the topological map is associated with a FreMEn model that represents the probability of someone interacting with the robot’s screen at a given time. This is built from logs of screen interactions stored in MongoDB. These FreMEn models are used to predict the likelihood of interactions at given times and locations. These predictions are used by the robot to schedule where and when it should provide information during the day.

The schedule has to satisfy two contradicting objectives common to many online, active learning tasks: exploration (to create and maintain the spatio-temporal models), and exploitation (using the model to maximise the chance of interacting with people). Exploration requires the robot to visit locations at times when the chance of obtaining an interaction is uncertain. Exploitation requires scheduling visits to maximise the chance of obtaining interactions. To tackle this trade off, the schedule is generated using Monte Carlo sampling from the location/time pairs according their FreMEn-predicted interaction probability (exploitation) and entropy (exploration). For more details see [16].

Figure 7 shows that by using this approach the robot was able to increase the number of successful interactions (i.e. when information was offered and someone interacted with the screen) on average per day over the course of its deployment. Although we have no control group to compare against, our on-site observations indicate that the robot’s choices are having a positive effect. This demonstrates the ability of the system to improve its application-specific behaviour from long-term experience.




Figure 7: The results of the robot selecting interaction times and locations using FreMEEn models learnt during the 2015 care deployment.




Figure 8: Top: The manually-created semantic map from the 2015 security deployment. Bottom: Example human trajectories with length close to the average trajectory length of $2.44m$. Also pictured are the manually annotated room regions we used for task planning.

10 Activity Learning

In our security scenario, the robot should learn models of normal human activity then raise an alert if an observation deviates from this. We have explored activity learning using walking trajectories (see Figure 8). Over the 2015 security deployment, the robot detected 42,850 individual trajectories. As described in [6], we use Qualitative Spatio-Temporal Activity Graphs (QSTAGs) to generalise from individual trajectories to spatial and temporal relations between trajectories and landmarks in a semantic map (see Figure 8). A QSTAG ignores minor quantitative variations across trajectories, but captures larger, qualitative changes. Every night the robot created QSTAGs for a subset of all trajectories (based on their displacement ratio) observed during the day. It then clustered these to create classes of movement activities. Some examples of the results can be seen in Figure 9.

During the day, an observation of a trajectory sufficiently far from any cluster centre triggered a task to approach the tracked human and request confirmation of their identity using a card reader. To enable a fast response it is important that the robot can accurately match the start of the trajectory to a cluster. Table 3 shows how the accuracy of predicting the cluster of a trajectory from an initial segment (20%) improves as more data is gathered over the robot’s lifetime. This provides another example of how a robot can improve its application-specific performance once it can operate over long periods.




Figure 9: Trajectories belonging to three learned clusters in the region at the bottom left of Figure 8 (direction of motion is red to green). These can be interpreted as two clusters of a desk approaching activity, and one of desk leaving.

| Training weeks ($\#traj.$) | K | recall | prec | $F1$ |
|------------------------------|-----|--------|------|------|
| week 0 (342) | 9 | 0.24 | 0.72 | 0.29 |
| weeks 0-1 (511) | 12 | 0.43 | 0.54 | 0.44 |
| weeks 0-2 (707) | 12 | 0.43 | 0.56 | 0.43 |
| weeks 0-3 (811) | 10 | 0.43 | 0.71 | 0.49 |
| weeks 0-4 (1016) | 14 | 0.48 | 0.63 | 0.53 |

Table 3: Accuracy of activity cluster prediction on week 5 data, from partial input trajectories.

11 Conclusions and Future Work

The STRANDS Core System features a mix of design- and run-time approaches which allow it to deliver LTA in everyday environments. A key strategy for delivering long-term robustness is the monitoring of system behaviour, from the individual component level up to navigation and task behaviours, plus the ability to restart system elements on demand. This allows the system to cope with unexpected situations both internally and in the external environment. Our aim is also to use the long-term experience of failures to learn to avoid these failures in the future. We presented our approach for doing this for navigation (Section 8), and hope to generalise this to other parts of the system. Whilst these features provide a fundamental ability to operate autonomously for long durations in everyday environments, our robots currently have no way to manage failures which are more catastrophic, harder to predict, or both. For example, our systems have suffered from PC component failure and subtle networking issues. In the future we would like to look at the use of redundancy and online reconfiguration (e.g. substituting a failing software or hardware component), coupled with more general failure detection approaches (both have which have been extensive researched in robots and other systems).

Our robots are able to learn online from lengths of experiences that no other robots to date have access to. The results above demonstrate what we have always known from machine learning: more data improves performance. However, the novel element here is that a robot must be able to operate for longer in order to gather additional data, and can make active choices about what data is gathered.

In the future we will also focus on the robot’s ability to understand human activities (the major causes of environment dynamics at most scales) and to actively close the gaps in its understanding it has already obtained from weeks of autonomous runtime.

12 Acknowledgements

We would like to acknowledge the contribution our project reviewers and project officers have made to our research: Luc De Raedt, James Ferryman, Horst-Michael Gross, Olivier Da Costa and Juha Heikkilä. The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement No 600623, STRANDS.

References

- [1] E. Marder-Eppstein, E. Berger, T. Foote, B. P. Gerkey, and K. Konolige, “The office marathon: Robust navigation in an indoor office environment,” in *ICRA ’10*, 2010.

- [2] J. Biswas and M. Veloso, “The 1,000-km challenge: Insights and quantitative and qualitative results,” *IEEE Intelligent Systems*, vol. 31, no. 3, pp. 86–96, May 2016.
- [3] S. Thrun, M. Bennewitz, W. Burgard, A. B. Cremers, F. Dellaert, D. Fox, D. Haehnel, C. Rosenberg, N. Roy, J. Schulte, and D. Schulz, “Minerva: A second-generation museum tour-guide robot,” in *ICRA '99*, 1999.
- [4] R. Ambrus, J. Ekekrantz, J. Folkesson, and P. Jensfelt, “Unsupervised learning of spatial-temporal models of objects in a long-term autonomy scenario,” in *IROS '15*, 2015.
- [5] 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 RA-L*, vol. 2, no. 1, pp. 26 – 33, 2016.
- [6] 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 *AAMAS '16*, 2016.
- [7] L. Kunze, C. Burbridge, M. Alberti, A. Tippur, J. Folkesson, P. Jensfelt, and N. Hawes, “Combining top-down spatial reasoning and bottom-up object class recognition for scene understanding,” in *IROS '14*, 2014.
- [8] D. Hebesberger, C. Dondrup, T. Körtner, C. Gisinger, and J. Pripfl, “Lessons learned from the deployment of a long-term autonomous robot as companion in physical therapy for older adults with dementia - A Mixed Methods Study,” in *HRI '16*, 2016.
- [9] D. Hebesberger, T. Körtner, J. Pripfl, and M. Hanheide, “What do staff in eldercare want a robot for? An assessment of potential tasks and user requirements for a long-term deployment.” in *IROS Workshop on "Bridging user needs to deployed applications of service robots"*, 2015.
- [10] C. Dondrup, N. Bellotto, M. Hanheide, K. Eder, and U. Leonards, “A computational model of human-robot spatial interactions based on a qualitative trajectory calculus,” *Robotics*, vol. 4, no. 1, pp. 63–102, 2015.
- [11] L. Mudrová, B. Lacerda, and N. Hawes, “An integrated control framework for long-term autonomy in mobile service robots,” in *ECMR '15*, 2015.
- [12] O. H. Jaffari, D. Mitzel, and B. Leibe, “Real-Time RGB-D based People Detection and Tracking for Mobile Robots and Head-Worn Cameras,” in *ICRA '14*, 2014.
- [13] J. Prankl, A. Aldoma, A. Svejda, and M. Vincze, “Rgb-d object modelling for object recognition and tracking,” in *IROS '15*, 2015.
- [14] T. Krajník, J. P. Fentanes, G. Cielniak, C. Dondrup, and T. Duckett, “Spectral analysis for long-term robotic mapping,” in *ICRA '14*, 2014.
- [15] B. Lacerda, D. Parker, and N. Hawes, “Optimal and dynamic planning for Markov decision processes with co-safe LTL specifications,” in *IROS '14*, 2014.
- [16] J. M. Santos, T. Krajník, J. P. Fentanes, and T. Duckett, “Lifelong information-driven exploration to complete and refine 4-d spatio-temporal maps,” *IEEE RA-L*, vol. 1, no. 2, pp. 684–691, 2016.