The International Journal of Robotics Research

http://ijr.sagepub.com/

Special Issue on Long-Term Autonomy

Tim Barfoot, Jonathan Kelly and Gabe Sibley

The International Journal of Robotics Research 2013 32: 1609
DOI: 10.1177/0278364913511182

The International Journal of
Robotics Research
32(14) 1609–1610

DOI: 10.1177/0278364913511182
ijr.sagepub.com

Special Issue on Long-Term Autonomy

This special issue of *The International Journal of Robotics Research (IJRR)* focuses on the theme of Long-Term Autonomy and is concerned with the leap from laboratories to long-term service roles. Naturally, there are big questions to which we would like answers. What are the challenges we must overcome in order to get robots out of lab and into the world? What assumptions should we relax to move forward? What direction should the community go in as a whole? These are difficult questions and this special issue will merely scratch the surface of the answers, but in opening a dialogue we hope to help focus the community moving forwards.

This special issue is another chapter in a dialogue that began a few years ago through a series of workshops:

- ICRA 2011 workshop on long-term autonomy;
- RSS 2011 workshop on autonomous long-term operation in novel environments (ALONE);
- ICRA 2012 workshop on long-term autonomy;
- RSS 2012 workshop on long-term operation of autonomous robotic systems in changing environments;
- ICRA 2013 workshop on long-term autonomy.

These workshops have grown in size and are often standing-room only; we hope this trend continues and we encourage others to attend and contribute.

We are delighted to have accepted nine papers to this special issue on a variety of topics that can be broadly classified into two main long-term autonomy themes: (i) localization and mapping for dynamic environments, and (ii) lifelong learning for adaptive behavior. These have also been recurring themes in the workshop series. One of our papers is a *data paper*, a special type of IJRR paper that formalizes the process of providing a user-friendly dataset to the community to facilitate comparison and simplification of large-scale experiments; this data paper certainly lives up to the moniker "long-term" as it provides 42 km of data, the full distance of a running marathon. Additional data papers can be found at: http://ijr.sagepub.com/cgi/collection/datapapers.

As a community, robotics has focused intently on the problem of localization and mapping since the mid-1980s. Much progress has been made under the banner of simultaneous localization and mapping (SLAM). However, for long-term autonomous operations in real-world environments, localization and mapping must be viewed as tasks that continue over the lifetime of a robot. Real environments change in their shape and appearance over time, both gradually and acutely. Our authors have provided a number of exciting new approaches to handling environmental change to support lifelong localization and mapping.

Non-linear, optimization-based SLAM has become a core tool in recent years. State-of-the-art systems typically divide the task of building a pose graph into a "back-end", which has the task of book-keeping all of the uncertainties in the system and appropriately relaxing the graph into a common frame, and a "front-end", which recognizes revisited places and adds appropriate loop-closure constraints for use by the back-end. Latif, Cadena, and Neira's paper, "Robust loop closing over time for pose graph SLAM", provides a novel method of maintaining the front-end/loop-closure decisions over time. Their method makes use of all available information to detect and remove past incorrect loop closures. The method can be executed in an incremental manner, with the result that more accurate maps can be constructed and maintained during robot operations. They validate their work on a variety of real-world datasets.

Saarinen, Andreasson, Stoyanov, and Lilienthal's paper, "3D normal distributions transform occupancy maps: An efficient representation for mapping in dynamic environments", combines two popular map representations, Normal Distribution Transforms (NDT) and Occupancy Grids, into a framework that allows them to build consistent laser-based maps online for use in dynamic environments. They demonstrate their approach using data from a milk-factory robot, whose environment is constantly changing as cargo is moved about.

Churchill and Newman's paper, "Experience-based navigation for long-term localisation", attacks the problem of outdoor, vision-based localization of a vehicle along a route. They demonstrate that despite large changes in lighting and weather across many route traversals, they can build a map consisting of a finite number of experiences that are sufficient to "span" the appearance of the route. They deliberately avoid fusing these experiences and show that they can use them in parallel to keep a vehicle well localized. A key advantage of their approach is the ability to selectively remember different appearances of the route that may recur at future times. They demonstrate their method on over 37 km of data from a car driving in an urban environment.

Tipaldi, Meyer-Delius, and Burgard's paper, "Lifelong localization in changing environments", also addresses the problem of localizing a robot in a dynamic environment, in their case a robot with a horizontal laser scanner driving in indoor office environments as well as outdoor environments such as parking lots. In contrast to the data-driven approach of the previous paper, these authors take a model-driven approach based on a dynamic occupancy grid, which not only learns a map's occupancy but also its dynamics based on multiple traversals of an environment. They demonstrate that it is possible to navigate a busy outdoor parking lot throughout a full day despite the fact that cars are continually arriving and leaving.

Biswas and Veloso's paper, "Localization and navigation of the CoBots over long-term deployments", presents the results of a long-term project to develop indoor service robots that can be deployed indefinitely in an environment undergoing continual change. Their approach seeks to find high-level line features of the environment that do not change in order to robustly localize over long periods of time. They demonstrate their method on datasets consisting of over 10 million laser scans from over 130 km of indoor driving gathered over a year and a half.

Fallon, Johannsson, Kaess, and Leonard's data paper, "The MIT Stata Center dataset", provides a large-scale dataset consisting of vision (stereo and RGB-D), laser, and proprioceptive data gathered using a PR2 robot over 42 km and 10 stories of the famous Stata Center. Given the variety and quantity of data, along with good ground-truth, high-quality comparisons of long-term indoor localization and mapping, and other techniques, should be possible.

The second theme in our special issue is lifelong learning for adaptive behavior. We cannot hope to equip robots with everything they need to know about the world prior to deployment. Some things will have to be learned along the way. Introspection, the ability for a robot to gauge its own performance, appears to be a promising notion to help in behavior adaptation for long-term autonomy.

Koos, Cully, and Mouret's paper, "Fast damage recovery in robotics with the T-resilience algorithm", considers the problem of a damaged robot that needs to adapt its controller to maintain performance. Here, introspection comes in the form of comparing the behavior of an internal model (of the undamaged robot) to that of the actual robot, in order to detect problematic situations and adapt. The method is demonstrated using a legged robot whose limbs are removed, broken, and rendered inoperable. They show the ability to continue to walk by adapting behavior.

Ott and Ramos' paper, "Unsupervised online learning for long-term autonomy", examines the problem of learning to predict collisions with obstacles from sensor data, simply by interacting with the environment (unsupervised), in an introspective manner. Through an online clustering method based on affinity propagation, they are able to learn models of obstacles from high-dimensional sensor streams, using onboard sensors such as bumpers for labeling. They demonstrate this on two problems: learning to predict collisions from laser data, and learning to classify obstacles from vision data.

Paul and Newman's paper, "Self-help: Seeking out perplexing images for ever improving topological mapping", nicely fits into both of our special issue themes. They consider the problem of having a robot perpetually improve its ability to navigate through introspection and selective data retrieval. They demonstrate this concept in the context of place recognition (often used as the front-end in pose graph SLAM), and in particular the popular FABMAP algorithm, which uses a training set of images to learn a statistical model of the appearance of a robot's environment. In the current paper, they allow the robot to download and incorporate new training images from an off-board repository to better explain perplexing images that it has gathered from the world, rather than relying on a static training set. They show the utility of their method on over 32 km of outdoor image data.

Together, we believe these nine papers represent some of the most interesting and promising avenues of research towards building autonomous robots that can operate for long periods of time without human intervention. We hope readers of this special issue will enjoy these contributions and that this will be the first of several such issues in the coming years. We would like to take this opportunity to thank the authors for their many wonderful submissions, the reviewers for their insights and suggestions, John Hollerbach for his ongoing support, and above all Jennet Batten who kept us going in the right direction all along. Happy reading!

Guest Editors:

Tim Barfoot
University of Toronto, Canada
tim.barfoot@utoronto.ca

Jonathan Kelly
University of Toronto, Canada

Gabe Sibley
University of Washington, USA