Taking the Long View: A Report on Two Recent Workshops on Long-Term Autonomy

By Jonathan Kelly, Gabe Sibley, Tim Barfoot, and Paul Newman

The topics of long-term autonomy and lifelong learning are attracting an increasing amount of attention in the robotics research community. Robots now routinely operate without human intervention for short periods of time, and several systems have demonstrated operation over longer durations. However, in the future, the majority of robots will need to function autonomously, outside the laboratory or the factory, on time scales ranging from days to years. Developing adaptive and flexible machines capable of this level of independence will require significant technical advances.

In this short report, we summarize the results of two recent workshops on long-term autonomy and lifelong learning: 1) the International Conference on Robotics and Automation (ICRA) 2011 Workshop on Long-Term Autonomy and 2) the Robotics Science and Systems (RSS) 2011 Workshop on Autonomous Long-Term Operation in Novel Environments (ALONE). These workshops brought together prominent researchers from a diverse range of subdisciplines to establish the key issues involved in building robotic systems that are able to operate in increasingly large-scale environments and over long periods of time. Many of the problems identified are difficult. At present, we can offer numerous questions but few answers, and this, in a way, is the defining characteristic of a new research area. Based on the success of the initial discussions, long-term autonomy clearly warrants an extended dialogue, with the next step being a follow-up workshop at ICRA 2012.

Challenges for Persistent Robotic Systems

State-of-the-art results in mapping and teach-and-repeat systems have convincingly shown that robots are able to operate over substantial spatial and temporal extents. The Defense Advanced Research Projects Agency (DARPA) Grand Challenge, the Mars Exploration Rovers, Willow Garage PR2 demos, and Atlantic Ocean-crossing autonomous underwater vehicles (AUVs) are only some examples of the increasing maturity of the field as a whole. However, with these successes, new challenges related to robust long-term operation and lifelong learning are also emerging.

In many respects, the problem domain is unlike any other. Particularly, long-term experiments often require a level of systems integration substantially above that of experiments that test only one or two primary technologies. This integration bar must be surmounted before the actual research can begin. There is an additional complication of academic verification—it can be very difficult to repeat another group's work to replicate their results. Simply replaying logged data is straightforward, but online decision-making and reactive control are substantially more complex to implement.

Both the ICRA and RSS workshops dealt with a broad range of subjects related to the construction of reliable, persistent autonomous systems. This encompassed long-term planning and exploration strategies; long-term adaptation; estimation in dynamic environments; perceptual representations incorporating space and time; exploitation of semantic information; semiautonomy and human-in-the-loop control; fault

Bowler WildCat autonomous vehicle platform. (Photo courtesy of the University of Oxford Mobile Robotics Group.)

Digital Object Identifier 10.1109/MRA.2011.2181792
Date of publication: 26 March 2012

Field robot on Devon Island in 2009 testing long-range autonomous driving in planetary analog domain. (Photo courtesy of Tim Barfoot.)

tolerance and failure prediction; novel sensors, actuators, and power sources; resource-constrained operation; robust systems engineering; and online calibration and identification.

One immediate task is to determine how to best ensure the future success of the subfield. Can we ascertain where researchers should focus their efforts now? Can we ascertain where researchers should focus their efforts now, to further facilitate the development and deployment of persistent systems? The community has started to consider ways to extend safe and reliable operation for longer durations. What valuable lessons have we learned from large-scale robotic experiments? What are the fundamental obstacles that need to be overcome to ensure robust and continual operations? What is the most suitable way to move current research from the laboratory into long-term service roles? These problems will be addressed as part of the ongoing discourse that has been initiated by the workshops.

The topics of long-term autonomy and lifelong learning are attracting an increasing amount of attention in the robotics research community.

Outcomes

Several common themes emerged from the workshop discussions, beyond simply establishing that the problems are especially hard in general. The themes included, e.g., the importance of redundant sensing for robustness, the need for novel map representations to improve large-scale navigation, the use of semiautonomy as a stepping stone to speed progress, and the importance of long-duration field experiments to uncover system weaknesses.

Another outcome was a list of pertinent questions, whose answers will help to guide the evolution of the subfield. Here, we outline these questions and give some related context.

- *What is long-term autonomy?* The phrase "long-term autonomy" is inherently broad and can often mean different things to different people. A concise definition and appropriate scoping may be necessary to secure funding from government agencies and other groups.
- *What characterizes successful long-term autonomy?* Related to the previous question, how does one characterize the success of these systems? What metrics are available? Are metrics specific to individual projects, or can we say something more generally? How can we evaluate our progress?
- *How far will sensing technology get us? What new sensors are required?* Many problems that are algorithmically difficult can be substantially mitigated or solved through improvements in sensing. Indeed, the release of products, such as the Velodyne laser scanner and the Microsoft Kinect RGB-D sensor, has largely revolutionized simultaneous localization and mapping (SLAM) research. What sensing advances can we depend on? And what types of new sensors would be most useful?
- *What is certification? What does it mean to certify these systems as safe?* What constitutes a safe system? What process will be required to certify a system as safe for long-term operation? What steps will be necessary for public acceptance and to ensure public trust?
- *Are these problems appropriate for the academic world?* Is long-term autonomy primarily an engineering problem? Or are there deep academic questions to be answered? What role should academia play in developing these systems?
- *What is an acceptable level of human intervention or human-in-the-loop control?* Some systems will not be fully autonomous but will instead rely on a human in the loop. What amount of human involvement is acceptable? Can some degree of human-in-the-loop control reduce the risks involved in development and deployment? Can human involvement be used as a stepping stone to get systems up and running more quickly with dependence then reduced over time?
- *Is it too soon to ask these questions?* Early in the proposal phase of both workshops, a point was made that it may be too soon to ask these questions. Has robotics advanced to the stage where the fielding of

such systems is practical? Are the relevant technologies at least partially ready, or are we getting ahead of ourselves?

In fact, the overall consensus was that now it is exactly the right time to begin exploring these issues. We expect that, as researchers continue to push the boundaries of long-term operation, many of the answers will begin to emerge.

The Workshops

ICRA 2011 Workshop on Long-Term Autonomy

The ICRA 2011 Workshop, organized by Gabe Sibley, Brian Gerkey, and Paul Newman, featured 12 presentations, including:

- Thierry Peynot (Australian Centre for Field Robotics), "Persistent Perception for Long-Term Autonomy of Ground Vehicles"
- Matthias Hentschel (University of Hannover), "An Adaptive Memory Model for Long-Term Navigation of Autonomous Mobile Robots"
- Cesar Cadena (University of Zaragoza), "A Learning Algorithm for Place Recognition"
- Wim Meeussen (Willow Garage), "Long-Term Autonomy in Office Environments"
- Michael J. Milford (Queensland University of Technology), "A Two-Week Persistent Navigation and Mapping Experiment Using RatSLAM: Insights and Current Developments"
- Tim Barfoot (University of Toronto), "Exploiting Reusable Paths in Mobile Robotics: Benefits and Challenges for Long-Term Autonomy"
- Nathaniel Fairfield (Google), "Self-Driving Vehicles at Google"
- Bradford Neuman (Carnegie Mellon University), "Anytime Online Novelty and Change Detection for Mobile Robots"
- Larry Matthies (NASA Jet Propulsion Laboratory), "Seven Years on Solar Power: Experience with Mars Rover Operations and Autonomy"
- Edwin Olson (University of Michigan), "The MAGIC Robotics Competition and the Challenges of 50 Robot-Hours of Supervised Autonomy"
- Rohan Paul (University of Oxford), "Self-Help: Seeking Out Perplexing Images for Ever Improving Navigation"
- Michael Kaess (Massachusetts Institute of Technology), "Towards Life-long Mapping and Navigation."

The workshop was well attended (standing room only), indicating that there is substantial interest in these topics among researchers and practitioners.

RSS 2011 ALONE Workshop

The 2011 ALONE Workshop was organized by Jonathan Kelly, Paul Newman, and Sebastian Thrun as part of the annual RSS conference. The speakers and talks were as follows:

- Tim Barfoot (University of Toronto), "Extensions to the Visual Odometry Pipeline for Long-Range Operations in Planetary Environments"
- Greg Dudek (McGill University), "Mission Planning and Endurance for Underwater and Harsh Terrain Missions with Medium- and Long-Term Duration"
- Brian Gerkey (Willow Garage), "Long-Term Autonomy in Office Environments"
- Leo Hartman (Canadian Space Agency), "Scalable Fault Management for Long-Lived Operations"
- Lionel Ott (Australian Centre for Field Robotics), "Unsupervised Incremental Learning for Object Discovery and Long-Term Autonomy"
- Salah Sukkarieh (Australian Centre for Field Robotics), "Autonomous Soaring: Exploration and Exploitation of the Wind Field for Persistent Flight"
- Chris Urmson (Carnegie Mellon University/Google), "Experiences in Field Testing: From Deserts to Downtown."

The workshop also incorporated a debate session to promote interaction between attendees. Feedback from the participants indicated that many found the session to be a valuable opportunity to exchange ideas.

Future Directions

In the near term, planning is underway for a special issue of a major robotics journal that focuses on long-term autonomy and related topics (coordinated by Gabe Sibley). A follow-up workshop (the Workshop on Long-Term Autonomy II) is being organized by Paul Furgale, Gabe Sibley, and Tim Barfoot, which is to be held at ICRA 2012 in Minneapolis, Minnesota.

Feedback from participants indicated that many found the session to be a valuable opportunity to exchange ideas.

Based on the response so far, the topics of long-term autonomy and lifelong learning appear to resonate with roboticists across a spectrum of application domains. The road ahead will be challenging, but the potential rewards are great. We look forward to continue and broaden the dialogue within the community in the years to come.

Jonathan Kelly, Robust Robotics Group, Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, USA. E-mail: jskelly@csail.mit.edu.

Gabe Sibley, Autonomous Robotics and Perception Group, George Washington University, D.C., USA. E-mail: gsibley@gwu.edu.

Tim Barfoot, Autonomous Space Robotics Lab, University of Toronto Institute for Aerospace Studies, Canada. E-mail: tim.barfoot@utoronto.ca.

Paul Newman, Mobile Robotics Group, Department of Engineering Science, University of Oxford, U.K. E-mail: pnewman@robots.ox.ac.uk.