Introduction to the Special Issue on AI for Long-Term Autonomy

Lars Kunze[1], Nick Hawes[1], Tom Duckett[2], Marc Hanheide[2]

I. INTRODUCTION

Autonomous systems have a long history in the fields of Artificial Intelligence (AI) and Robotics. However, only through recent advances in technology has it been possible to create autonomous systems capable of operating in long-term, real-world scenarios. Examples include autonomous robots that operate outdoors on land, in air, water, and space; and indoors in offices, care homes, and factories. Designing, developing, and maintaining intelligent autonomous systems that operate in real-world environments over long periods of time, i.e. weeks, months, or years, poses many challenges. This special issue focuses on such challenges and on ways to overcome them using methods from AI.

Long-term autonomy can be viewed as both a *challenge* and an *opportunity*. The challenge of long-term autonomy requires system designers to ensure that an autonomous system can continue operating successfully according to its real-world application demands in unstructured and semi-structured environments. This means addressing issues related to hardware and software robustness (e.g., gluing in screws and profiling for memory leaks), as well as ensuring that all modules and functions of the system can deal with the variation in the environment and tasks that is expected to occur over its operating time. Early research in long-term autonomy for mobile robots focussed extensively on the problem of coping with environment variation, e.g., performing visual localisation over seasonal changes, or SLAM in a dynamic environment. Once such challenges are overcome, the long-term operation of an autonomous system provides the opportunity to specialise that system to its tasks and operating environment, potentially improving both its general robustness and more specific task performance. This specialisation may come through the use of extensive system logs to detect and fix bugs, or through automatic online adaptation via machine learning. There is also the opportunity to aid development through the use of logs to create tests and simulations which are representative of a specific system deployment.

A great many research fields have the potential to contribute to enabling and improving long-term operation of autonomous robots. For example, formal methods can be used to verify robot software or control policies to ensure safe long-term operation, and novel locomotion methods can be used to minimise a robot's long-term energy usage. However, we chose to focus on AI techniques for long-term autonomy since AI provides a wide range of algorithmic approaches for addressing both the challenge of online long-term variation in task and environment (e.g. through task planning or probabilistic inference), and for the opportunity of online adaptation through experience (via machine learning). There are also a variety of AI techniques which have yet to be extensively applied in long-term settings, and we therefore see long-term autonomy as a motivating challenge for our community in the years ahead. This challenge has huge societal and industrial relevance as autonomous systems enter our lives in ever increasing numbers.

[1]Oxford Robotics Institute, Dept. of Engineering Science, University of Oxford, United Kingdom; {lars|nickh}@robots.ox.ac.uk

[2]School of Computer Science, University of Lincoln, United Kingdom; {tduckett|mhanheide}@lincoln.ac.uk

II. GUIDE TO THE SPECIAL ISSUE

The overview of AI for long-term autonomy presented above is greatly expanded in the survey paper included in this special issue [1]. The survey approaches the topic from two complementary angles. It starts by surveying autonomous robot systems that have been fielded in uncontrolled environments for extended periods, grouping these systems across application domains such as space, marine, service and road. It then surveys the AI techniques which cut across these systems and identifies trends in the approaches used to provide long-term behaviour in fielded systems. It concludes by identifying future challenges and opportunities around AI for long-term autonomy. The survey groups AI techniques under research areas including Navigation & Mapping; Perception; Knowledge Representation & Reasoning; Planning; Interaction; and Learning. Three of these areas in particular are represented by the other papers in the Special Issue, described as follows, while the other areas present excellent opportunities for novel future research in AI for long-term autonomy.

A. Interaction

Human-robot interaction research for autonomous long-lived systems has two facets: the modelling and analysis of interaction over long durations; and the support of long-term autonomy through interactive learning.

An example of the latter is the work by Del Duchetto et al. [2], who employ Gaussian Processes in order to interactively learn local navigation recovery behaviours for mobile robots. Using a long-term dataset, they highlight that navigation failures are predictably localised in space-time, and propose a solution that can exploit rare interactive teach-in opportunities from situations where the robot was helped by human bystanders.

Herrero et al. [3] provide a comprehensive analysis of the user base of an autonomous system, which was deployed for long-term operation in a care home, hence studying long-term interaction of an autonomous robotic system. The core contributions of the paper are a processing pipeline to automatically estimate demographics (age, gender) of interacting users, and a model to discriminate between passive interactions (bystanders observing the mobile robot), and active interactions (users directly interacting with the robot). Their work facilitates user analysis of robots deployed "in the wild", i.e. without experimenters on site, in long-term settings.

B. Perception

Pre-trained detectors and classifiers are usually used to allow a robot to perceive the objects in its environment. In a long-term setting it is not typically possible to predict in advance the complete set of objects that a robot needs to perceive, or the contexts in which which they may occur.

Chaudhary et al. [4] address part of this problem by developing a technique to segment generic hand-held objects from RBG-D input. These segmentations can then be used during long-term operation as input into an online object learning approach. The paper presents a methodology based on a deep comparison and segmentation network which significantly outperforms state-of-the-art object segmentation methods from literature.

C. Navigation & Mapping

Maps are needed for wayfinding and other spatial reasoning tasks. In a long-term autonomy setting, maps must support these capabilities under the variation of the environment over the lifetime of the system. The contributed papers in this special issue address some of the core problems in long-term mapping and navigation by mobile robots, including long-term visual place recognition [5], [6], SLAM [7], and semantic mapping [8].

Han et al. [5] introduce a novel representation for visual place recognition that simultaneously integrates semantic landmarks and holistic information to achieve long-term operation. Evaluations on public benchmarks, including the Nordland and CMU-VL datasets, demonstrate the ability to recognise places in long-term scenarios, while outperforming state-of-the-art approaches for visual recognition.

Chen et al. [6] propose a visual attention model for place recognition, which learns to automatically focus on regions that are most discriminative in defining a place, using context information extracted from a deep neural network. Evaluation against state-of-the-art approaches, including FABMap and SeqSLAM, on several benchmarking datasets demonstrated superior performance in place recognition against strong viewpoint and condition variations.

Bescos et al. [7] present a dynamic SLAM algorithm, DynaSLAM, for handling dynamic objects in RGB-D, stereo and monocular SLAM, by adding a new front-end for motion segmentation to the existing ORBSLAM2 system. Evaluations on the TUM Dynamic Objects and KITTI datasets show that DynaSLAM outperforms the accuracy of standard visual SLAM baselines in highly dynamic scenarios.

Sun et al. [8] propose an approach for long-term semantic mapping, using a deep-learning approach to fuse the semantic features from 3D lidar data, where the dynamic state of each cell in a 3D OctoMap is modelled by a recurrent neural network. Evaluations on the ETH Parking-Lot dataset demonstrate the ability to learn the semantic state transitions while outperforming conventional Bayesian updates in long-term scenarios.

ACKNOWLEDGMENTS

The work of formulating and editing this special issue has been supported by EPSRC/UK Research and Innovation grants EP/M019918/1 (Mobile Robotics Programme Grant), EP/R02572X/1 (NCNR), EP/R026084/1 (RAIN), EP/M015777/1 (ALOOF); and EU projects No. 732737 (ILIAD) and 600623 (STRANDS).

REFERENCES

[1] L. Kunze, N. Hawes, T. Duckett, M. Hanheide, and T. Krajnk, "Artificial intelligence for long-term robot autonomy: A survey," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 4023–4030, Oct 2018.

[2] F. D. Duchetto, A. Kucukyilmaz, L. Iocchi, and M. Hanheide, "Do not make the same mistakes again and again: Learning local recovery policies for navigation from human demonstrations," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 4084–4091, Oct 2018.

[3] R. P. Herrero, J. P. Fentanes, and M. Hanheide, "Getting to know your robot customers: Automated analysis of user identity and demographics for robots in the wild," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3733–3740, Oct 2018.

[4] K. Chaudhary, K. Wada, X. Chen, K. Kimura, K. Okada, and M. Inaba, "Learning to segment generic handheld objects using class-agnostic deep comparison and segmentation network," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3844–3851, Oct 2018.

[5] F. Han, S. E. Beleidy, H. Wang, C. Ye, and H. Zhang, "Learning of holism-landmark graph embedding for place recognition in long-term autonomy," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3669–3676, Oct 2018.

[6] L. Sun, Z. Yan, A. Zaganidis, C. Zhao, and T. Duckett, "Recurrent-OctoMap: Learning state-based map refinement for long-term semantic mapping with 3-d-lidar data," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3749–3756, Oct 2018.

[7] B. Bescos, J. M. Fcil, J. Civera, and J. Neira, "DynaSLAM: Tracking, mapping, and inpainting in dynamic scenes," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 4076–4083, Oct 2018.

[8] Z. Chen, L. Liu, I. Sa, Z. Ge, and M. Chli, "Learning context flexible attention model for long-term visual place recognition," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 4015–4022, Oct 2018.

Lars Kunze is a Postdoctoral Research Assistant in Mobile Robotics in the Oxford Robotics Institute, within the Department of Engineering Science at the University of Oxford, UK, and a Stipendiary Lecturer in Computer Science at Keble College. His research concerns the design and development of fundamental AI techniques for autonomous robot systems that operate in complex, real-world environments. In particular he focuses on semantic scene understanding using logical as well as statistical-relational models. Within the Oxford Robotics Institute he works on transparent and interpretable models that can provide detailed explanations in the context of autonomous driving.

Nick Hawes is an Associate Professor in the Oxford Robotics Institute, within the Department of Engineering Science at the University of Oxford, UK, and a Tutorial Fellow at Pembroke College. Within the Oxford Robotics Institute he leads the GOALS Lab which performs research around the problems of behaviour generation for autonomous systems. In particular he focuses on long-term autonomy, and task and mission planning, for mobile robots which must operate for extended periods (days, weeks or months) in dynamic, uncertain environments. He led the EU STRANDS project which deployed mobile robots in end-user environments for over a year of autonomous behaviour.

Tom Duckett is a Professor of Robotics & Autonomous Systems at the University of Lincoln, UK, where he also leads the Lincoln Centre for Autonomous Systems. His research interests include autonomous robots, artificial intelligence, machine perception and robotic mapping, with applications including agri-food and service robotics. He has been an investigator in many national and international research projects, funded by the European Commission, the UK Research Councils and Innovate UK, including ILIAD, STRANDS and various projects in precision agriculture.

Marc Hanheide is a Professor of Intelligent Robotics & Interactive Systems in the School of Computer Science at the University of Lincoln, UK. Marc Hanheide is a Principal Investigator in many national and international research projects, funded by H2020, EPSRC, Innovate UK, and industry partners, including STRANDS, ILIAD, RASberry, NCNR and others. In his research he focuses on long-term autonomy for robots, robot navigation, human-robot interaction, and long-term learning and adaptation, with applications in logistics, agricultural robotics, care, nuclear decommissioning, and service robotics in the wider sense.