

Artificial Intelligence for Long-Term Robot Autonomy: A Survey

Lars Kunze¹, Nick Hawes¹, Tom Duckett², Marc Hanheide², Tomáš Krajník³

Abstract—Autonomous systems will play an essential role in many applications across diverse domains including space, marine, air, field, road, and service robotics. They will assist us in our daily routines and perform dangerous, dirty and dull tasks. However, enabling robotic systems to perform autonomously in complex, real-world scenarios over extended time periods (i.e. weeks, months, or years) poses many challenges. Some of these have been investigated by sub-disciplines of Artificial Intelligence (AI) including navigation & mapping, perception, knowledge representation & reasoning, planning, interaction, and learning. The different sub-disciplines have developed techniques that, when re-integrated within an autonomous system, can enable robots to operate effectively in complex, long-term scenarios. In this paper, we survey and discuss AI techniques as ‘enablers’ for long-term robot autonomy, current progress in integrating these techniques within long-running robotic systems, and the future challenges and opportunities for AI in long-term autonomy.

Index Terms—Autonomous Agents, AI-Based Methods, Long-term Autonomy.

I. INTRODUCTION

ROBOT technology has improved tremendously over the last decade. Consequently, autonomous robot systems have been able to operate in increasingly complex environments and for increasingly long periods of time, i.e. weeks, months, or years. When a fully modelled robot is deployed in a completely known, static environment, the challenge of long-term autonomy (LTA) reduces to one of robustness, i.e. enabling the robot to remain operational for as long as possible. Without these simplifying assumptions autonomous robots face a number of interrelated challenges. We roughly characterise these challenges on two dimensions. The first refers to the application requirements, e.g., the robot platform (hardware and software), environment and tasks to be performed. The second dimension describes the long-term nature of these elements, e.g., if and how they change over time, whether their long-term nature can be fully characterised in advance (structured vs. unstructured), and how observable they are. For example, in many long-term applications, the

environment will change over the lifetime of the system. These changes could be short-term (e.g. things moving within the robot’s field of view), medium-term (e.g. furniture moving between visits to a room, parked cars changing positions on roads), or long-term (e.g. seasonal changes, plant growth, wear to surfaces). In addition, parts of the environment may not be fully known before deployment or new objects may appear. In AI terms this means dealing with an *open world*. It is also possible that the end-user may change the tasks or how the robot should accomplish them, or the robot itself may need to adapt to new tools, techniques (e.g. AI algorithms) or knowledge as they become available.

In this paper, we survey systems and approaches that address the challenges of LTA using techniques from AI. We focus on both AI techniques used by robot systems deployed for extended periods in real-world environments (Sec. II), and techniques that align well with future needs of LTA systems, but which have not been convincingly demonstrated within long-term scenarios (Sec. III). We further discuss future challenges and opportunities for AI in LTA (Sec. IV).

By focussing on the above challenges we purposely exclude other applications where robots operate for long periods, but in relatively static, known settings. Specifically this means we do not cover current systems in manufacturing or intralogistics. In both cases the majority of deployed robot systems have demonstrated significant longevity, but this is typically achieved through the creation of environments which are fully known and the dynamics are largely under control of the autonomous systems. Whilst this does not eliminate the need for AI techniques (e.g. long-term localisation [1] and planning [2] for warehouse AGVs), it limits the LTA-specific challenges (e.g. environmental dynamics, lack of structure, open-endedness) seen in other domains.

Since autonomous systems research has a long history, special issues and surveys already exist that relate to LTA. For example, [3], [4] cover AI methods within integrated robot systems. In perception, existing collections cover localisation and mapping in dynamic environments [5], [6], calibration [7] and visual place recognition [8]. However, this is the first survey that focuses specifically on AI techniques for enabling long-term robot autonomy.

II. DOMAINS

Long-term autonomous robots have been deployed in a variety of domains including space, marine, air, field, road, and service. Tab. I provides an overview of these domains and selected systems characterised across common features. In this survey, we adopt the notation of [4] and characterise domains by application features: *environment variability*, *task diversity*, *semantics*, *dynamics*, *partial observability*, *cost & criticality*,

Manuscript received: February 24, 2018; Revised June 5, 2018 ; Accepted July 11, 2018.

This paper was recommended for publication by Editor Dongheui Lee upon evaluation of the Associate Editor and Reviewers’ comments.

This work was supported by EPSRC grant EP/M019918/1 and EP/R02572X/1 (NCNR); EU project No. 732737 (ILIAD); and CZ projects 17-27006Y and CZ.02.1.01/0.0/0.0/16_019/0000765.

¹Lars Kunze and Nick Hawes are with Oxford Robotics Institute, Dept. of Engineering Science, University of Oxford, United Kingdom; {lars|nickh}@robots.ox.ac.uk

²Tom Duckett and Marc Hanheide are with School of Computer Science, University of Lincoln, United Kingdom; {tduckett|mhanheide}@lincoln.ac.uk

³Tomáš Krajník is with Dept. of Computer Science, Faculty of Electrical Engineering, Czech Technical University in Prague; tomas.krajnik@fel.cvut.cz

Digital Object Identifier (DOI): see top of this page.

TABLE I
SURVEYED AI-ENABLED LONG-TERM AUTONOMY ROBOT SYSTEMS.

Domain	Application Features								Duration	AI Areas						System
	Environment Variability	Task Diversity	Semantics	Dynamics	Partial Observability	Cost & Criticality	Interaction & Cooperation	Level of Autonomy		Navigation & Mapping	Perception	KR & Reasoning	Planning	Interaction	Learning	
Space	L	L	L	L	H	H	L	M	Years	○	●	–	●	○	–	Opportunity [9], [10]
									Years	○	●	–	●	○	–	IPEX [11]
Marine	M	L	L	M	H	H	L	H	Days	○	●	○	●	–	○	AUVs [12], [13]
									Months	○	○	–	○	–	–	Gliders [14]
Air	M	M	M	H	H	H	M	M	Days	○	●	○	○	–	–	AtlantikFlyer [15]
Field	H	M	L	M	H	M	M	M	Days	●	●	○	–	○	○	VT&R2 [16]
									Years	●	●	○	–	–	○	BearNav [17], [18]
Road									Days	○	●	●	○	–	○	VaMP [19]
	M	L	M	H	M	H	M	L	Days	○	●	○	○	–	○	ARGO [20]
									Months	○	●	○	○	–	○	PANS [21]
									Months	○	●	○	○	–	○	VIAC [22]
									Days	●	○	○	●	●	○	Rhino [23]
Service	H	H	H	L	H	L	H	M	Days	●	○	○	○	●	○	Minerva [24]
									Days	●	○	○	○	●	○	Willow Garage [25]
									Months	●	●	●	●	●	●	STRANDS [26]
									Years	●	●	●	●	●	●	CoBot [27]

Legend: L low, M medium, H high, – not integrated, ○ partially integrated, ● fully integrated

interaction & cooperation, and level of autonomy. As in [4], all features are qualitatively assessed using three levels (*low, medium, high*). Please see the aforementioned paper for a detailed discussion of the features. In this work, we focus on the assessment of deployed robot systems. To this end, we assess them by the duration of their deployment (*days, months, years*) and the level of integration of different AI areas (*not, partially, or fully integrated*). The three qualitative levels of integration distinguish between systems that do *not* employ AI techniques from a particular area; that use them only *partially*, but not for the purpose of LTA; and systems that use them for LTA (*fully integrated*).

Space: Due to extreme communication delays and limited prior access, effective extra-terrestrial exploration requires autonomous systems. NASA’s Opportunity rover has recently passed its 5,000th day operating on the surface of Mars. Its autonomous capabilities come from a mixed-initiative task planner, and an autonomous navigation system. The planner (MAPGEN, [9]) is used to automatically create a daily mission schedule, which is then refined by terrestrial scientists. The navigation system uses stereo cameras to build 3D models for terrain traversability and path planning [10]. LTA has also been a growing part of satellite operations, e.g. the Intelligent Payload Experiment (IPEX) demonstrating over a year of autonomous information gathering using planning and image processing technology [11].

Marine: Due to the limits of communication through water, and the difficulties in fully mapping deployment environments, there are parallels between the requirements for, and benefits of, autonomy in marine and space robotics. Autonomous wave gliders are routinely deployed for long durations, with missions measured in thousands of kilometres and hundreds of days (e.g. 7,400 km in 221 days [14]). Gliders are relatively simple, low-powered robots. More powerful systems have been deployed for days of autonomous operation, e.g. for navigation under ice [12]. The benefits of AI planning have been shown in field trials [13] and controlled settings targeting LTA [28].

Air: The fundamental factor that makes long-term operation of aerial systems difficult is energy. The authors of [15] argue that to achieve perpetual autonomous flight, the UAV has to plan its path according to global and local weather conditions, wind fields, and thermal updrafts. An alternative to perpetual flight is the ability to interrupt the flight to recharge, like the lake monitoring system in [29].

Field: Field robotics deals with unstructured and dynamic environments in diverse domains such as forestry, agriculture, mining, construction, etc. Bechar and Vigneault [30] characterise such domains according to the level of structure present in both the environment and the objects relevant to the robot. The majority of current field robots utilise GPS-based auto-steer systems that follow pre-determined paths

with otherwise limited use of AI capabilities. An alternative approach uses visual ‘teach and repeat’ to enable robust navigation in field environments. In these approaches systems are driven along a training route and then they repeat the route autonomously [31]. Krajnik *et al.* [17] show that their teach-and-repeat method is robust to seasonal appearance changes. Paton *et al.* [16] showed that the integration of multiple experience-based representations [39] results in a system capable of long-term autonomous navigation despite drastic changes to the environment appearance.

Road: The PANS platform [21] was one of the first autonomous vehicles that drove a long distance (6,000 miles, 98.2% autonomous driving) on public roads over a period of six months using a vision-based driving system. It used a neural network to learn a mapping between road images and appropriate vehicle turn radiuses from human demonstrations. At the same time, the driverless car VaMP [19] drove more than 1,000 miles (95% autonomous driving). The vision-based driving system of the ARGO project [20] achieved a similar result (1,200 miles, 94% autonomous driving). Through learning it was able to adapt to new road conditions (lane width and lane position). More recently, several vehicles covered a distance of 13,000 km from Italy to China using a leader-follower approach in the VisLab Intercontinental Autonomous Challenge (VIAC) (2010) [22].

Service: We characterise service robots as robots that work for, or alongside, humans in environments that are not specially adapted for their presence. Service robots must cope with: dynamic environments (due to people moving, day-night changes, etc.); open worlds (due to people); and changing task requirements. Large-scale research initiatives have deployed mobile robot systems capable of LTA in museums (the seminal Rhino [23] and Minerva [24]), offices (Willow Garage [25], CoBot [27] and STRANDS [26]), stores [32], and care environments (STRANDS [33] and Tangy [34]). All of these robots were deployed for at least multiple weeks, and most around naïve users. Most of these systems were deployed at intervals in the same environment (e.g. daily). The STRANDS and Willow Garage systems also attempted *continuous* autonomous operation, managing a maximum of 28 and 13 days respectively of uninterrupted operation. These research systems have given rise to the current generation of autonomous service robots operating in human-populated spaces. Examples include Bossanova’s stock checking robots in Walmart stores, Knightscope’s security robots, and Savioke’s robot hotel butlers.

Conclusion: With respect to AI areas, *Navigation & Mapping* and *Perception* are the only areas that were present in all surveyed systems. This is no surprise as they provide robots with very fundamental capabilities. *KR & Reasoning* as well as *Planning* were both supported by most systems. However, we hypothesise that work on *KR & Reasoning* in space and marine is limited due to the lack of semantics in these domains. Furthermore, it is interesting to note that systems only partially (if at all) support *Interaction* and *Learning* in most domains (with an exception of the service domain). Although these areas are well researched in general, they haven’t been extensively covered in long-term scenarios. Hence, we believe

that there are many open challenges and research opportunities for both areas (and in their intersection) as we point out in Sec. IV.

In all domains, LTA systems inherently present an *integration* challenge, particularly when different AI abilities need to work together. Over the past years, there has been an increasing trend toward the (re-)integration of AI techniques within robotics. To cope with challenging environments and tasks, robots typically integrate: localisation and navigation; object and/or person perception; plus task planning and/or scheduling. However, although the integration of AI techniques at system-level is an essential part of all research projects, there is no standard solution and little research on how to combine modules from different areas of AI.

Robotic software development [35] and robotic middleware projects such as the Robot Operating System (ROS) [36] provide researchers with common methods to integrate their software components and components of others in a structured way. Some frameworks build on top of such middlewares and integrate particular AI methods in the context of long-term navigation planning and task scheduling (STRANDS [26]), planning and execution (ROSPlan [37]), and knowledge-enabled perception (RoboSherlock [38]). In general, these frameworks make it easier to integrate and use different AI methods. However, overall, there is still a lack of understanding and research in the area of *system-level integration*. Hence, we believe system-level integration of AI methods and their evaluation will continue to be a major challenge in autonomous systems research.

III. AI AREAS

In this section we discuss how different areas of AI can enable autonomous robot systems to perform in real-world environments over extended periods of time. This includes navigation & mapping, perception, knowledge representation & reasoning, planning, interaction, and learning.

A. NAVIGATION & MAPPING

Navigation is an essential ability for purposeful movement by autonomous robots. One approach uses visual ‘teach and repeat’ to enable robust navigation in field environments, where the robot learns a map while being driven along a training route and then repeats the route autonomously [31], [17], [16], as discussed earlier in Sec. II. Recent work [39] demonstrated over 140 km of autonomous driving with an autonomy rate of 99.6%, including driving at night-time.

Over the past 30 years there has been huge interest in autonomous learning of environment models by robots, especially the problem of simultaneous localisation and mapping (SLAM) [6]. However, most approaches assume a static world and do not consider long-term updating of robot maps to reflect environment changes. Here we briefly characterise several complementary strategies to enable long-term mapping and localisation in changing environments, primarily using long-term data sets for their experiments.

Multiple representations: Long-running robots need to consider environment mapping as a never-ending process, and thus make decisions on what to remember and what to forget. However, deleting information from a map is risky since an observed change may only be temporary and the environment may yet revert to a previously observed state. One approach is thus to maintain multiple environment representations [8], then to select the most relevant model at the current time for localisation and planning. Early work [40] developed dynamic maps that handle changes through use of robust statistics and multiple local maps at different timescales, where the map that best explains the current sensor data is used for localisation. The short-term maps are updated online while the longer-term maps are adapted offline based on long-term experience. Stachniss & Burgard [41] cluster local grid maps created at different times and learn distinct configurations of these locations. A related approach for pose-graph SLAM [42] maintains multiple view-based representations of mapped locations, while discarding obsolete views, thus limiting overall map size. Similarly, Churchill & Newman [43] propose to integrate similar observations at the same spatial locations into ‘experiences’ which are then associated with a given place. For localisation, they select the experience that best matches the visual input of the robot. An alternative approach is to keep the data from all mapping sessions and integrate them offline into a single, high fidelity representation [44].

Robustness to appearance change: A parallel strategy attempts to select the representation which is most stable in time. Valgren & Lilienthal [45] demonstrated the robustness of local image features for localisation across seasons. SeqSLAM [46] attempts to match sequences of images rather than individual images, achieving robust place recognition across seasons. A method for learning long-term stable features is described by Dayoub et al. [47], where image features detected across mapping sessions are first stored in a short-term memory, which is used to filter out spurious observations, before being admitted to long-term memory. A further approach involves learning to predict appearance changes across seasons [48], by learning the expected translation between a vocabulary of superpixels for different seasons and using this to generate predicted images for localisation at run-time. Recent research showed that season-specific images can also be predicted using generative adversarial networks [49], [50]. Lowry & Milford [51] compare a similar appearance prediction technique with a change removal method and conclude that change removal is more robust and less data-intensive to train. Related work on laser-based localisation [52] uses long-term experience to learn error distributions for individual points in 3D point-cloud maps, which are then used during localisation to suppress the observations corresponding to map points with high errors.

Learning about dynamics: While the above approaches are mainly concerned with learning the persistent elements of the scene, another strategy attempts to model the dynamics. Tipaldi et al. [53] use dynamic occupancy grids, which model the occupancy of each cell as a two-state Markov process, and showed that their approach improves localisation robustness in a car park environment. Kucner et al. [54] learn conditional probabilities of neighbouring cells in an occupancy grid to

model typical motion patterns in dynamic environments. Krajinik et al. [55] proposed to represent rhythmic or periodic processes in the environment using Fourier analysis, and showed that the resulting spectral models obtained from long-term experience enable prediction of future environment states, improving localisation and navigation in human-populated environments.

Notable applications of long-term mapping include visual survey of natural environments by an autonomous surface vessel surveyed a lake shore over a 14-month period [56], and a 4D reconstruction approach to crop monitoring over time [57]. The latter comprises a 3D SLAM pipeline, data association to find correspondences between crop rows and sessions, and optimisation of the full 4D reconstruction.

Finally, complementary work on topological and semantic mapping may further enhance long-term robustness to change, by abstracting away from the finer details of metric and feature-based representations, although a detailed review is beyond the scope of this paper. Current trends suggest that future work on long-term navigation and mapping will include more application-specific developments across all domains, as long-running systems continue to be deployed in practice, and development of richer environment representations including especially more semantics and integration of more perceptual and contextual cues.

B. PERCEPTION

In addition to perception algorithms for navigation and mapping, autonomous robots need general perception routines for object recognition and scene understanding. Indeed, early approaches to mapping of dynamic environments were object-centric. These methods identify moving objects and remove them from the maps [58] or use them as moving landmarks for self-localisation [59]. However, not all dynamic objects actually move at the moment of mapping, meaning that their identification requires long-term observations.

To address this challenge, Ambrus et al. [60] processed several 3D point clouds of the same environment recorded over several weeks to identify and separate movable objects, and refine the static environment structure at the same time. Biswas & Veloso [61] proposed an approach for long-term localisation based on explicit reasoning about object categories including mapped objects, unmapped static objects and unmapped dynamic objects. Bore et al. [62] detect and localise objects in large environments, where objects can change locations between observations by the robot, while assuming a closed world to ensure computational tractability.

Other approaches enable open-ended learning of new object categories during long-term operation, e.g. using spatial context information to query possible category labels from semantic knowledge on the web. Recent work includes an embodied system for open-ended learning and manipulation of new object categories, based on human-robot interaction [63], and a lifelong learning framework in which a human user can direct a robot to capture domain-relevant data for training classifiers of household objects [64].

Future service robots would also benefit from techniques to improve their perception of people over time,

e.g. by integrating long-term experience in tracking-learning-detection [65] and tracking-learning-classification [66] approaches, and learning the long-term activity patterns of people [67]. This would in turn enable robots to adapt to and move more harmoniously with the expected flow of humans. Long-term applications involving interaction with specific people also require algorithms for person re-identification at different temporal scales (very short term, same day, different day). For these cases, different assumptions can be made based on the persistence of supporting cues (e.g., position at dinner table, clothing, size/stature, hair colour, facial features). Recent work integrates person re-identification with multi-target multi-camera tracking [68], however, adaptation of person-specific appearance models over long time-periods remains an open challenge for autonomous robots. Direct parallels may be drawn for other related challenges such as recognition of human activities, where long-term experience can be leveraged to improve performance over time [69]. In general, most prior work on perception considers only the initial training phase prior to deployment of the robot, but not the ongoing adaptation of learned models during operation.

C. KNOWLEDGE REPRESENTATION & REASONING

Knowledge representation (KR) is concerned with representing the world (in particular domains with rich semantics, see Tab. I), and is closely linked to many other AI areas including perception, planning, and learning. The previous sections already discussed important aspects of representations in the context of navigation and perception.

In general, KR goes hand and hand with reasoning as both decision making and inference are tightly coupled with the way the knowledge is represented. Long-term, autonomous robot systems that are deployed in a real-world environments require KR and reasoning capabilities to represent various aspects of the world and reason about them, in particular when they change over time. Therefore, AI areas such as spatio-temporal reasoning, non-monotonic reasoning, and belief revision are of great importance in long-term scenarios.

Several works investigated models that infer locations of entities in space and time. Mason *et al.* [70] proposed an object-based semantic world model for long-term change detection in dynamic environments, and [71] modelled the temporal persistence of objects. Similarly, Krajník *et al.* [72] proposed frequency-based spatio-temporal models for reasoning about the location of people. Such spatio-temporal information is essential in long-term scenarios as it can inform AI planners (cf. Sec. III-D) about non-stationary costs and/or rewards. Santos *et al.* [73] presented a first lifelong information-driven approach to spatio-temporal exploration that incrementally completes and refines environment maps.

In LTA it is of great importance that robots can access and learn from their own experience. OPEN-EASE [74] is a KR infrastructure that makes experience data from robots and human manipulation episodes semantically accessible. Users can retrieve experiences and query what the robot perceived, reasoned and did. Balint-Benczedi *et al.* [75] propose a more specialised framework for storing and retrieving perceptual

memories for long-term manipulation tasks. Similarly, [76] propose a long-term knowledge acquisition framework using contextual information in a memory-inspired robot architecture. The framework allows robots to memorise their perceptions and to recall them, e.g. in a manipulation task.

To cope with the challenges of open worlds, novelty and anomaly detection is of great importance. To this end, [77] proposed a framework for anomaly reasoning which includes the recognition and interpretation of unfamiliar and familiar objects appearing in unexpected contexts. This aspect of KR and reasoning is strongly linked to work in adaptation and learning (Sec. III-F) as it can trigger learning in LTA systems.

D. PLANNING

AI planning and scheduling technologies, which determine the sequence of actions necessary to achieve a task, are often used to adapt the robot's behaviour online to account for environment or task dynamics [4]. We have seen planning systems deployed on almost all LTA systems. For example, planning approaches were used to produce daily task lists and the associated action sequences for the Opportunity rover [9], and the STRANDS [26], CoBot [78] and Tangy [34] service robots, allowing these robots to adapt their behaviour to the needs of their users. AUVs used planning to deal with changing environmental conditions and resources [13], while logistics systems used planning to enable large numbers of robots to cope with variety in customer orders [2].

Planning approaches vary in their ability to represent critical elements of a system's long-term experience. The aforementioned systems vary in terms of whether or not they model the effects/dependencies of a robot's actions on time or resources (such as battery), or under uncertainty. They also vary in how they handle oversubscription (choosing between multiple goals, a key issue in integrating exploration). In general, planning algorithms in LTA robots are embedded in a wider integrated system which handles the parts omitted from the planning model (e.g. replanning on failed on actions, or reactively triggering charging on low battery, or managing goal choices). More generally, an executive control system which manages tasks, and responds to opportunities and failures, is an essential part of a robot architecture for long-term autonomy [28], [13], [26], [79], [23]. Such a system prevents the robot getting stuck in behavioural loops, and provides recovery mechanisms to address autonomy-hindering failures. This behaviour can be seen in simple yet effective form in the finite state controller of Willow Garage's office marathon system [79], through the planner-based executives of Rhino [23] and Minerva [24], through to the T-Rex executive used on fielded AUV teams [13].

The planned behaviour of the STRANDS [26] and CoBot [78] robots was generated using models learnt during execution: STRANDS robots created optimal task and navigation plans from learnt MDP models of environmental dynamics [80], [81]; CoBot robots learnt and planned with models which predicted when humans would be available to help complete a task [82]. These robots were therefore able to adapt their task and navigation plans over the long term.

As robots become more adept at navigation and manipulation in less structured environments, we expect to see planning playing an even more prominent role in logistics, since a greater degree of variation will need to be managed autonomously over the life of the robot. We also expect to see overlap in learning and planning/optimisation processes in autonomous transport systems [83], where system performance will need to be managed over variation in demand and execution resources. Current trends also include the augmentation of plans or policies created by mission or task planning systems with richer execution knowledge [84]. This hybrid approach allows mission planning to address long-term or large-scale problems with abstract, computationally tractable models, but at execution time behaviour is guided by richer models which allow appropriate responses to dynamic events.

E. INTERACTION

Some of the most challenging application domains for long-running robots involve interacting autonomously with a diverse range of users, offering opportunities for the robots not only to learn and adapt from this experience, but also facilitating longitudinal studies to gain a better understanding of long-term engagement of humans with autonomous robots. In general, humans and other actors introduce a level of dynamics and non-predictability into any application scenario, and hence pose dedicated challenges for LTA systems, also indicated in Tab. I, where *environment variability* is considered high in domains with a high level of *interaction and cooperation*. However, long-term Human-Robot Interaction (HRI) studies with truly autonomous *social* robots are still a rarity today, as many researchers resort to Wizard-of-Oz settings [85], where subjects in studies are deceived into believing a robot is acting autonomously while it is in fact remote-controlled by a human operator. Among the most explored domains for long-term autonomous systems with an emphasis on interaction are museums [24], [86], care [33], domestic [87], retail [32], [88], hospitality [89], and educational environments [90], [91], [92].

A recent survey [93] identified key domains for long-term interactive robotic systems including ‘Health Care and Therapy’, ‘Education’, ‘Work Environments and Public Spaces’, and ‘At Home’, discussing a total of 45 different studies in these fields. From their analysis, the key conclusion drawn regarding autonomy is the lack of but also need for more learning and adaptation. Indeed several systems mentioned above (e.g. [88], [92]) develop personalised models to maintain an interaction context. Such individualised user profiles are one of the key abilities required to enable interaction in long-running autonomous systems [94].

Hence, interaction in the context of long-term autonomy must not only be seen as a challenge, but also as an opportunity, where representations can be learned or adapted in an *in-situ* fashion to improve a system’s autonomous behaviour from exploiting long-term interactions with users. [82] propose a model enabling the robot to predict when humans are most likely available to help a robot, while [33] follow similar ideas, learning spatio-temporal usage patterns to maximise the utility of the mobile robot.

F. LEARNING

Machine learning plays a role in many of the above areas, and is clearly a key enabling technology for LTA. Beyond this component role, a cluster of learning types are specifically suited to LTA. In general we see techniques that allow a robot to learn *during operation* (rather than during a design phase) as crucial to success in LTA applications. Long-term deployment in open/dynamic worlds means that any knowledge or experience the robot starts with is unlikely to be sufficient to cover the behaviour required of it during operation. Thus learning during operation is essential to delivering good performance. We have seen this from the relatively low level of estimating cost and probability models for planning [26], up to learning new object [95] and activity models for service robotics tasks [96]. Since it is hard to receive supervision signals during long-term autonomous operation, the majority of the online learning techniques employed by LTA systems are unsupervised.

By definition, a robot is restricted to a fixed viewpoint in space and time. This means that it is limited in the experiences, and thus training data, it can generate to facilitate online learning. Therefore many LTA systems also include an *exploration* component which drives the gathering of new experiences. For example, CoBot robots were able to choose navigation routes which provided updated observations for environmental models [97], and the STRANDS robots balanced exploration and exploitation to maximise interactions with humans during an information provision task [33] and to build 3D maps for object discovery [60], [73].

Given the richness and diversity of techniques in machine learning, many approaches could influence the ability of LTA systems to learn on the job in the future. Techniques which allow robots to continually learn from experience such as reinforcement learning, or focus on particular experiences (e.g. failures, novelty) such as learning from demonstration should allow online improvement of capabilities. Problems due to limited training in a particular domain (or open worlds) can be addressed by transfer learning, and supported by work from the exploration and active learning communities. Ongoing research is also investigating deep learning methods for long-term autonomy, including recent work on prediction of human trajectories from long-term observations [98].

IV. FUTURE CHALLENGES

This paper discussed the LTA-related challenges in different areas of AI and the importance of system-level integration for unlocking the potential of AI technologies. In addition, we see the following major future challenges for LTA systems in real-world environments:

Human-in-the-Loop Systems: How can LTA systems leverage human knowledge in unforeseen situations within long-term scenarios? As LTA systems have to deal with *open worlds*, they will certainly require additional information when facing situations that were not foreseen at design time. This additional input might be given by end-users, maintainers, and/or domain experts. It might also be provided through direct control (i.e. teleoperation), natural interaction (e.g. via

language or gestures) or labelled examples and/or data sets (e.g. via crowd-sourcing). To this end, LTA systems require mechanisms to integrate new, but potentially conflicting and/or untrustworthy, information in their KR about the world. This also requires that representations have some kind of semantic abstraction that can be linked to human knowledge. The fan-out [99], or number of robots a human can control simultaneously, will help drive the mixture of human supervisors and robots in such a paradigm.

Knowledge Transfer between LTA Systems: What information about the world should be exchanged by LTA systems, and when? As more and more LTA robot systems get deployed, they can exchange important information to help bootstrap other systems and/or to improve their performance. As it is not realistic that all logged information is exchanged, it is important to investigate what information should be exchanged, and when. This also opens up privacy and security concerns. In this context, we believe that cloud-enabled knowledge bases [100] and other approaches in cloud robotics [101] will play an important role.

Systems Integration: Building robotic systems capable of long-term operations is inherently also a software engineering challenge, as they require the close integration of different AI techniques, a challenge also highlighted by [3]. While ROS has established itself as a de-facto standard framework for building integrated robotic systems, it provides only few instruments to ensure reliable and robust system architectures. Here, model-based approaches [102] might pave the way towards more dependable and verifiable integrated systems in the future.

More Domain Specialisation: Alongside the development of general principles of AI for long-term autonomy, there will be many interdependencies and synergies from solving the application-specific challenges in parallel. For example, in precision agriculture the accuracy of relative positioning and navigation, e.g. with respect to crop rows, is more important than that of absolute navigation and position as provided by RTK GPS. Therefore, any improvements in recognition of crops would in turn improve the robustness and accuracy of navigation in crop care and harvesting tasks.

Verification and Evaluation of LTA Systems: How can the behaviour and the performance of LTA systems be verified and evaluated when robot system (including its models), task specification, and environment are constantly changing (at different timescales)? This requires LTA systems to keep a record of all their internal models that were used at a given time. Furthermore, it requires novel ways to provide formal guarantees under the assumption that parts of the environments might change (with some probability) [103].

Conclusion: Further to these technological challenges, we also see ethical, social, and legal issues when realising LTA systems, though these are beyond the scope of this paper. Overall, we believe strongly that AI methods can provide LTA systems with many of the capabilities needed to overcome these challenges. In turn, rather than merely extending the lifetime of existing AI-enabled robots, AI approaches may actually help to solve some of the really tough open problems in robotics, e.g. perception-based mobile manipulation in real-world settings, by leveraging long-term experience. However,

we recognise that, despite recent progress, there are still many exciting open challenges.

REFERENCES

- [1] M. Gadd and P. Newman, "A framework for infrastructure-free warehouse navigation," in *ICRA*. IEEE, 2015, pp. 3271–3278.
- [2] P. R. Wurman, R. D'Andrea, and M. Moutz, "Coordinating hundreds of cooperative, autonomous vehicles in warehouses," *AIM*, 2008.
- [3] K. Rajan and A. Saffiotti, "Towards a science of integrated AI and robotics," *Artificial Intelligence*, vol. 247, pp. 1–9, 2017.
- [4] F. Ingrand and M. Ghallab, "Deliberation for autonomous robots: A survey," *Artificial Intelligence*, vol. 247, pp. 10–44, 2017.
- [5] T. Barfoot, J. Kelly, and G. Sibley, "Special issue on long-term autonomy," *IJRR*, vol. 32, no. 14, pp. 1609–1610, 2013.
- [6] C. Cadena *et al.*, "Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age," *IEEE T-RO*, 2016.
- [7] P. Furgale, C. Pradalier, and T. D. Barfoot, "Editorial: Special issue on calibration for field robotics," *Journal of Field Robotics*, 2015.
- [8] S. Lowry *et al.*, "Visual place recognition: A survey," *IEEE T-RO*, 2016.
- [9] J. L. Bresina, A. K. Jónsson, P. H. Morris, and K. Rajan, "Activity planning for the mars exploration rovers." in *ICAPS*, 2005, pp. 40–49.
- [10] S. Goldberg, M. Maimone, and L. Matthies, "Stereo vision and rover navigation software for planetary exploration," in *Aerosp. Conf.*, 2002.
- [11] S. Chien *et al.*, "Onboard autonomy on the intelligent payload experiment cubesat mission," *J. of Aero. Inf. Sys.*, 2016.
- [12] C. Kunz *et al.*, "Toward extraplanetary under-ice exploration: Robotic steps in the arctic," *J. Field. Rob.*, no. 4, pp. 411–429, 2009.
- [13] L. Chrupa *et al.*, "On mixed-initiative planning and control for autonomous underwater vehicles," in *IROS*, 2015.
- [14] C. P. Jones, "Slocum glider persistent oceanography," in *Autonomous Underwater Vehicles (AUV)*. IEEE/OES, 2012, pp. 1–6.
- [15] P. Oetershagen *et al.*, "Design of small hand-launched solar-powered UAVs: From concept study to a multi-day world endurance record flight," *J. Field. Rob.*, 2017.
- [16] M. Paton *et al.*, "I can see for miles and miles: An extended field test of visual teach and repeat 2.0," in *Field and Serv. Rob.*, 2018.
- [17] T. Krajník *et al.*, "Simple yet stable bearing-only navigation," *J. Field. Rob.*, 2010.
- [18] T. Krajník *et al.*, "Image features for visual teach-and-repeat navigation in changing environments," *RAS*, 2016.
- [19] E. D. Dickmanns, *Dynamic Vision for Perception and Control of Motion*, Secaucus, NJ, USA, 2007.
- [20] A. Broggi *et al.*, "The ARGO autonomous vehicles vision and control systems," *IJICS*, vol. 3, no. 4, pp. 409–441, 1999.
- [21] T. Jochem, D. Pomerleau, B. Kumar, and J. Armstrong, "PANS: a portable navigation platform," in *IVS*, 1995.
- [22] A. Broggi *et al.*, "The vislab intercontinental autonomous challenge: an extensive test for a platoon of intelligent vehicles," *IJVAS*, 2012.
- [23] W. Burgard *et al.*, "The Interactive Museum Tour-Guide Robotics," in *AAAI*, 1998.
- [24] S. Thrun *et al.*, "Minerva: A second-generation museum tour-guide robot," in *ICRA*, 1999.
- [25] W. Meeussen, E. Marder-Eppstein, K. Watts, and B. Gerkey, "Long Term Autonomy in Office Environments," in *ICRA LTAW*, 2011.
- [26] N. Hawes *et al.*, "The STRANDS Project: Long-Term Autonomy in Everyday Environments," *IEEE RAM*, 2017.
- [27] J. Biswas and M. Veloso, "The 1,000-km challenge: Insights and quantitative and qualitative results," *IEEE Intelligent Systems*, 2016.
- [28] N. Palomeras *et al.*, "Toward persistent autonomous intervention in a subsea panel," *Auton. Rob.*, vol. 40, no. 7, 2016.
- [29] R.-A. Pélouquin, D. Thibault, and A. L. Desbiens, "Design of a passive vertical takeoff and landing aquatic uav," *IEEE RAL*, 2017.
- [30] A. Bechar and C. Vigneault, "Agricultural robots for field operations: Concepts and components," *Bio. Eng.*, vol. 149, pp. 94–111, 2016.
- [31] P. Furgale and T. D. Barfoot, "Visual teach and repeat for long-range rover autonomy," *J. Field. Rob.*, vol. 27, no. 5, pp. 534–560, 2010.
- [32] H. Gross *et al.*, "TOOMAS: Interactive Shopping Guide robots in everyday use-final implementation and experiences from long-term field trials," in *IROS*, 2009.
- [33] M. Hanheide *et al.*, "The When, Where, and How: An Adaptive Robotic Info-Terminal for Care Home Residents," in *HRI*, 2017.
- [34] T. Tran *et al.*, "Robots in retirement homes: Applying off-the-shelf planning and scheduling to a team of assistive robots," *IJRR*, 2017.
- [35] D. Brugalí and E. Prassler, "Software engineering for robotics [from the guest editors]," *IEEE RAM*, vol. 16, no. 1, pp. 9–15, 2009.

- [36] M. Quigley *et al.*, “Ros: an open-source robot operating system,” in *ICRA Workshop on Open Source Software*, 2009.
- [37] M. Cashmore *et al.*, “ROSPlan: Planning in the robot operating system,” in *ICAPS*, 2015.
- [38] M. Beetz *et al.*, “RoboSherlock: Unstructured Information Processing for Robot Perception,” in *ICRA*, 2015.
- [39] M. Paton *et al.*, “Expanding the limits of vision-based localization for long-term route-following autonomy,” *IJRR*, 2017.
- [40] P. Biber and T. Duckett, “Dynamic maps for long-term operation of mobile service robots,” in *RSS*, 2005.
- [41] C. Stachniss and W. Burgard, “Mobile robot mapping and localization in non-static environments,” in *Conf. on Artificial Intelligence*, 2005.
- [42] K. Konolige and J. Bowman, “Towards lifelong visual maps,” in *IROS*, 2009, pp. 1156–1163.
- [43] W. S. Churchill and P. Newman, “Experience-based navigation for long-term localisation,” *IJRR*, 2013.
- [44] P. Mühlfellner *et al.*, “Summary maps for lifelong visual localization,” *J. Field. Rob.*, 2015.
- [45] C. Valgren and A. J. Lilienthal, “SIFT, SURF & seasons: Appearance-based long-term localization in outdoor environments,” *RAS*, 2010.
- [46] M. Milford and G. Wyeth, “SeqSLAM: Visual route-based navigation for sunny summer days and stormy winter nights,” in *ICRA*, 2012.
- [47] F. Dayoub *et al.*, “Long-term experiments with an adaptive spherical view representation for navigation in changing envs,” *RAS*, 2011.
- [48] P. Neubert, N. Stünderhauf, and P. Protzel, “Appearance change prediction for long-term navigation across seasons,” in *ECMR*, 2013.
- [49] H. Porav, W. Maddern, and P. Newman, “Adversarial training for adverse conditions: Robust metric localisation using appearance transfer,” in *ICRA*, 2018.
- [50] Y. Latif, R. Garg, M. Milford, and I. Reid, “Addressing challenging place recognition tasks using generative adversarial networks,” in *ICRA*, 2018.
- [51] S. Lowry and M. J. Milford, “Supervised and unsupervised linear learning techniques for visual place recognition in changing environments,” *IEEE T-RO*, 2016.
- [52] D. Withers and P. Newman, “Modelling scene change for large-scale long term laser localisation,” in *ICRA*. IEEE, 2017.
- [53] G. D. Tipaldi, D. Meyer-Delius, and W. Burgard, “Lifelong localization in changing environments,” *IJRR*, 2013.
- [54] T. Kucner *et al.*, “Conditional transition maps: Learning motion patterns in dynamic environments,” in *IROS*, 2013.
- [55] T. Krajník, J. P. Fentanes, J. M. Santos, and T. Duckett, “FreMen: Frequency map enhancement for long-term mobile robot autonomy in changing environments,” *IEEE T-RO*, 2017.
- [56] S. Griffith and C. Pradaliar, “Survey registration for long-term natural environment monitoring,” *J. Field. Rob.*, 2016.
- [57] J. Dong *et al.*, “4d crop monitoring: Spatio-temporal reconstruction for agriculture,” in *ICRA*, 2017.
- [58] D. F. Wolf and G. S. Sukhatme, “Mobile robot simultaneous localization and mapping in dynamic environments,” *Auton. Rob.*, 2005.
- [59] C.-C. Wang *et al.*, “Simultaneous localization, mapping and moving object tracking,” *IJRR*, no. 9, 2007.
- [60] R. Ambrus *et al.*, “Meta-rooms: Building and maintaining long term spatial models in a dynamic world,” in *IROS*. IEEE, 2014.
- [61] J. Biswas and M. Veloso, “Episodic non-markov localization: Reasoning about short-term and long-term features,” in *ICRA*, 2014.
- [62] N. Bore *et al.*, “Multiple Object Detection, Tracking and Long-Term Dynamics Learning in Large 3D Maps,” *ArXiv e-prints*, 2018.
- [63] S. Kasaei *et al.*, “Towards lifelong assistive robotics: A tight coupling between object perception and manipulation,” *Neurocomp.*, 2018.
- [64] C. Eriksen, A. Nicolai, and W. Smart, “Learning object classifiers with limited human supervision on a physical robot,” in *IRC*, 2018.
- [65] Z. Kalal, K. Mikolajczyk, and J. Matas, “Tracking-learning-detection,” *IEEE TPAMI*, 2012.
- [66] Z. Yan, T. Duckett, and N. Bellotto, “Online learning for human classification in 3d lidar-based tracking,” in *IROS*, 2017.
- [67] S. M. Mellado *et al.*, “Modelling and predicting rhythmic flow patterns in dynamic environments,” in *TAROS*, 2018.
- [68] L. Beyer, S. Breuers, V. Kurin, and B. Leibe, “Towards a Principled Integration of Multi-Camera Re-Identification and Tracking through Optimal Bayes Filters,” *CVPR Workshops*, pp. 1444–1453, 2017.
- [69] C. Coppola, T. Krajník, T. Duckett, and N. Bellotto, “Learning temporal context for activity recognition,” in *ECAI*, 2016.
- [70] J. Mason and B. Marthi, “An object-based semantic world model for long-term change detection and semantic querying,” 2012.
- [71] R. Toris and S. Chernova, “Temporal persistence modeling for object search,” in *ICRA*, 2017, pp. 3215–3222.
- [72] T. Krajník *et al.*, “Where’s waldo at time t ? using spatio-temporal models for mobile robot search,” in *ICRA*, 2015, pp. 2140–2146.
- [73] J. Santos *et al.*, “Lifelong information-driven exploration to complete and refine 4-d spatio-temporal maps,” *IEEE RAL*, 2016.
- [74] M. Beetz *et al.*, “Open-EASE – a knowledge processing service for robots and robotics/ai researchers,” in *ICRA*, 2015.
- [75] F. Balint-Benczedi *et al.*, “Storing and retrieving perceptual episodic memories for long-term manipulation tasks,” in *ICRA*, 2017.
- [76] F. Pratama *et al.*, “Long-term knowledge acquisition using contextual information in a memory-inspired robot architecture,” *JETAJ*, 2017.
- [77] J. Winder, “A framework for anomaly reasoning: Interpretation through concept formation for knowledge transfer and lifelong learning,” in *IJCAI*, 2016.
- [78] M. Veloso, J. Biswas, *et al.*, “Cobots: Collaborative robots servicing multi-floor buildings,” in *IROS*, 2012.
- [79] E. Marder-Eppstein *et al.*, “The Office Marathon: Robust navigation in an indoor office environment,” in *ICRA*, 2010.
- [80] B. Lacerda *et al.*, “Optimal and dynamic planning for markov decision processes with co-safe LTL specifications,” in *IROS*, 2014.
- [81] J. Fentanes *et al.*, “Now or later? predicting and maximising success of navigation actions from long-term experience,” in *ICRA*, 2015.
- [82] S. Rosenthal, M. Veloso, and A. K. Dey, “Is someone in this office available to help me?” *JINT*, 2012.
- [83] J. Miller and J. P. How, “Predictive positioning and quality of service ridesharing for campus mobility on demand systems,” in *ICRA*, 2017.
- [84] L. Iocchi *et al.*, “A practical framework for robust decision-theoretic planning and execution for service robots,” in *AAAI*, 2016.
- [85] L. Riek, “Wizard of Oz Studies in HRI: A Systematic Review and New Reporting Guidelines,” *Journal of HRI*, pp. 119–136, 2012.
- [86] I. Nourbakhsh, C. Kunz, and T. Willeke, “The mobot museum robot installations: a five year experiment,” in *IROS*. IEEE, 2003.
- [87] M. M. de Graaf, S. Ben Allouch, and J. A. van Dijk, “Long-term evaluation of a social robot in real homes,” *Interaction Studies*, 2016.
- [88] M. E. Foster *et al.*, “The MuMMER Project: Engaging Human-Robot Interaction in Real-World Public Spaces,” in *ICSR*, 2016.
- [89] R. Pinillos *et al.*, “Long-term assessment of a service robot in a hotel environment,” *RAS*, 2016.
- [90] T. Kanda *et al.*, “A Two-Month Field Trial in an Elementary School for Long-Term Human-Robot Interaction,” *IEEE T-RO*, 2007.
- [91] A. Jacq *et al.*, “Building successful long child-robot interactions in a learning context,” in *HRI*, 2016.
- [92] P. Baxter *et al.*, “Robot education peers in a situated primary school study: Personalisation promotes child learning,” *PLoS ONE*, 2017.
- [93] I. Leite, C. Martinho, and A. Paiva, “Social Robots for Long-Term Interaction: A Survey,” *ICSR*, 2013.
- [94] S. Rossi, F. Ferland, and A. Tapus, “User profiling and behavioral adaptation for HRI: A survey,” *Pattern Recognition Letters*, 2017.
- [95] J. Young *et al.*, “Towards Lifelong Object Learning by Integrating Situated Robot Perception and Semantic Web Mining,” in *ECAI*, 2016.
- [96] P. Duckworth *et al.*, “Unsupervised learning of qualitative motion behaviours by a mobile robot,” in *AAMAS*, 2016.
- [97] M. Korein, B. Coltin, and M. Veloso, “Constrained scheduling of robot exploration tasks,” in *AAMAS*, 2014.
- [98] L. Sun *et al.*, “3DOF pedestrian trajectory prediction learned from long-term auton. mobile robot deployment data,” in *ICRA*, 2018.
- [99] D. R. Olsen and S. B. Wood, “Fan-out: Measuring human control of multiple robots,” in *CHI*, 2004.
- [100] L. Riazuelo *et al.*, “Roboearth semantic mapping: A cloud enabled knowledge-based approach,” *IEEE TASE*, 2015.
- [101] B. Kehoe, S. Patil, P. Abbeel, and K. Goldberg, “A survey of research on cloud robotics and automation,” *IEEE TASE*, 2015.
- [102] C. Schlegel *et al.*, “Model-driven software systems engineering in robotics: Covering the complete life-cycle of a robot,” *IT*, 2015.
- [103] S. A. Seshia and D. Sadigh, “Towards verified artificial intelligence,” *CoRR*, vol. abs/1606.08514, 2016.