Auton Robot (2016) 40:1119–1122
DOI 10.1007/s10514-016-9582-0

Long duration autonomy for maritime systems: challenges and opportunities

Marc Steinberg[1] · **Jason Stack**[1] · **Terri Paluszkiewicz**[1]

Received: 26 May 2016 / Accepted: 15 June 2016 / Published online: 2 July 2016

1 Introduction

Achieving long duration capabilities has long been an important goal of research in autonomous systems. Persistence can enable such systems to be on station to respond to events and situations that were unknown at the start of a mission. It can enable data collection in spatially and temporally variable environments over time scales appropriate for understanding a wide range of phenomena. Long duration systems have improved ability to cover large distances, travel in ways that minimize energy consumption, and go places that may not be practical for manned or shorter duration systems. Long duration is also an important part of the design trade space for minimizing the costs involved in achieving certain kinds of mission capabilities. This is particularly true when thinking of long duration in terms of entire systems and not just individual platforms. For example, a system with an infrastructure that supports autonomously keeping sufficient numbers of individual platforms on station continuously would fall within this category even if each platform or sensor had only limited persistence individually. Additionally, what constitutes long duration with regards to these challenges may vary significantly from system to system.

Much work in the past has focused on extending system limitations through methods such as efficient management of resources, energy harvesting, detection, identification, and mitigation of component failures or other adverse situations that become more likely as mission time increases. More recently, the development of advanced forms of autonomy has led to a focus on new challenges and opportunities with the autonomy algorithms themselves in areas like perception-based control and decision-making, knowledge representation, simultaneous localization and mapping, planning, learning, and collaboration. Such approaches may have fundamental limitations that do not scale with the length of mission time or the size of the physical space that the system explores. They may become brittle and fail as the environment changes due to weather effects that occur over days, seasonal cycles that take place over months, or climate and terrain variations on the timescale of years or decades. For methods that learn and change, they may fail to converge or converge to an undesirable solution if run over a long enough period of time in unstructured and dynamic environments.

2 Attributes and technical challenges

To consider long duration autonomy in a more general sense, a definition is proposed that involves the application of autonomy to problems for which the mission time is much greater than the validity of a priori assumptions and information. This can be thought of along three dimensions. The first dimension is the perishability of a priori knowledge. This knowledge concerns the state of the environment; the context of the mission; the aspects of the world that are relevant to the mission goals; and the states, beliefs, and intentions of other human and machine agents for which these dependencies exist. This dimension is important because the estimation, assimilation, and use of this knowledge is non-trivial and can consume resources such as sensing, computing, energy, and mission time to acquire. Additionally, autonomy methods such as

---

✉ Marc Steinberg
marc.steinberg@navy.mil

Jason Stack
jason.stack@navy.mil

Terri Paluszkiewicz
terri.paluszkiewicz@navy.mil

[1] Office of Naval Research, 875 N. Randolph, Arlington, VA, USA

learning or planning may be very sensitive to the uncertainty in this knowledge. The second dimension is the extent to which the mission times are on scales that approach or exceed system or subcomponent reliability and particularly make it likely that problems will be encountered that were not foreseen in typical failure management and effects analyses. This dimension is important as it creates a need for more metacognitive reasoning, understanding of context, and the ability to solve novel problems and regularly reduce uncertainty as a function of mission time as part of core system capabilities. This need stands in contrast to more traditional fault monitoring, diagnostics, and redundancy management as it requires a much more holistic view of the autonomy and its purpose. Further, this applies not only to hardware but also to software algorithms that may encounter problems if allowed to run for too long in unstructured conditions. It also includes not only outright failures but also slower types of degradation due to prolonged environmental exposure such as biofouling. The third and final dimension is the existence of and/or need for mission goals, constraints, and priorities that must be adjudicated as a function of mission phase and situational context, that may be in competition or confliction as the situation changes over time, and may be only partially known or articulated at mission start. The important aspect of this dimension is that moving among or reordering goals, constraints, or priorities consumes non-trivial resources, involves uncertainty and decision-making, and is distinct from simply transitioning between operating modes. For example, simpler missions may require multiple operating modes such as ingress, egress, survey, launch, and recovery that can be achieved through waypoint following and executed in a preprogrammed, serial fashion. In contrast, a long duration mission may discover in situ that the risks or opportunities are different than expected, emerging constrains may make achieving all goals infeasible, or it may be unclear if any particular goal has been achieved. In such instances, the system must recognize any uncertainties, potentially take action to improve its understanding of the world, and consider changing its goals.

Achieving long duration in this sense raises many challenges. Currently fielded maritime autonomy has been most successful on short duration missions and has been limited to modest numbers of mobile platforms. Skilled human operators are typically relied upon for management of the systems goals, interpreting engineering data that indicates a malfunction or degradation, and handling any unexpected contingencies or environmental factors. Automation is typically tailored for particular uses and can be brittle outside of those circumstances. These systems traditionally have limited understanding of their own status and limited ability to appraise how a particular environmental context or event may relate to their ability to achieve their goals. At the stage in the science, systems may have no awareness about being in an uncertain state. Even when problems can be detected, the system may lack the ability to reason on and recover from such problems. In grappling with these problems, some approaches may ultimately run into computational limits or fail after a certain amount of time. Other approaches may rely on brute force and power intensive use of payloads to solve problems in ways that limit the total operational time. Overall, the current state of the art in integrating theory, algorithms, software, and hardware do not necessarily allow these systems to carry out longer duration missions in the maritime environment.

To make progress in this area, systems are needed that are non-myopic when compared to traditional types of automated planning and decision-making methods, and this must manifest in multiple ways. First, systems must be able to understand how an action today might impact their ability to do something far in the future even when there is substantial uncertainty in the future state. Second, there is a need for new methods that can ensure robustness even when the environment, context, and world state change repeatedly and at different time scales. In the most extreme cases, this may involve operating in harsh environments with substantial risks from a poor decision or in an "open world" where there is an arbitrarily large number of objects and entities that might relate to the mission that can each have many attributes, functions, affordances, and relationships. Third, there is a need for methods to deal with circumstances in which not only vehicle states but also task progress and achievement are not known with certainty or are only quantifiable in a probabilistic sense. Fourth, systems that operate over very long time periods will create new challenges in the types of interactions they have with people. It will be important to rethink the roles and responsibilities of humans and autonomy for these types of long duration missions. Additionally, there is a need for methods that can ensure adequate communication and comprehension between humans and autonomy. This is particularly challenging for in-mission interactions that are infrequent, latent, or bandwidth limited as well as pre- or post-mission interactions that are high bandwidth but limited in duration. A final significant challenge is performing verification, validation, test, and evaluation for achieving long duration in unstructured, uncertain, and dynamic environments. Generally, these problems become much harder to solve in instances where one cannot assume a constrained and largely known environment, narrow and well-defined mission tasks, and a well-trained supervisor with the ability to mitigate many of the problems the system might encounter.

3 Metrics and research testbed needs

These kinds of systems may drive a need to reconsider the types of metrics that are most important. Traditional metrics such as Operational Availability may still be applicable.

However, the unique nature of long duration missions will likely drive the need for new and more appropriate metrics to quantify these additional capabilities and attributes. While the community is only in the early stages of defining what these metrics might be, the following categories represent some of the areas that should be incorporate in metric development:

1. Time, frequency, or effort between undesired human interventions to solve or prevent problems—In this context, "undesired" means that it would be possible for all required information and physical capacities to be locally resident so that the situation theoretically could have been solved using onboard autonomy only. This would not apply to situations in which human interaction was desirable, such as to redirect the system because the human's goals have changed.
2. Time, frequency, or effort for learning/adaptation—This concerns how much time, effort, or the number of instances are required for the system to learn or adapt to new situations, contexts, or environments. Metrics in this space include both upper bounds that should be minimized (i.e., enable the autonomy to learn and adapt quickly) as well as lower bounds that could be used to avert premature human intervention (e.g., if the human control loop is out of phase with the system dynamics due to communication latency or inaccurate information, then premature human intervention could actually degrade adaptation or its convergence rate).
3. Safety and robustness—These metrics concern how close the system comes to undesirable or failure conditions, how frequently unsafe or unacceptable actions are selected, and the frequency of errors or constraint violations.
4. Collaboration and scalability—Collaboration concerns the costs associated with working as part of a team (e.g., cost of communicating and sharing information with other agents, cost of changing teaming structure, cost of task handoff, etc.). Scalability concerns how system capabilities, attributes, and approaches scale with mission time or scope.
5. Information sharing—Percentage of time the system passes the correct information to support decisions and the percentage of time information is received in time by either the human users or other elements of the autonomous system to make critical decisions.
6. System cost—These metrics concern the cost to achieve a given mission with a long-duration autonomous system and may be normalized against a traditional way of achieving that same mission. For example, a sensing mission could be measured by cost per bit or cost per data product.

Addressing these kinds of metrics at a research level will drive a need for new long duration platform and multi-platform simulators and hardware testbeds. This need exists both for higher fidelity integrated simulation testbeds that can affordably address specialized maritime problems and "lite" testbeds that can be an easy entry point to engage the broader robotics, control, and computational intelligence communities. With air and ground systems, there are now standard laboratory configurations used very broadly by researchers with different areas of expertise. There is not yet an equivalent for maritime systems. One consequence is that researchers from other fields often do not fully understand the severity and uniqueness of the constraints, coupling, and limitations in communications, sensing, power management, and control in this environment. More general methods may rely on assumptions that are not valid in the maritime domain. For example, in acoustic communications, bandwidth is a function of range and may interfere with sensing in a way different from electromagnetic propagation. One possibility would be laboratory configurations of small, easily used maritime systems in pools with some ability to control environmental conditions to simulate ocean phenomena with appropriate scaling. Another approach would be the use of surrogate systems such as air platforms that have attributes or dynamics more analogous to those of interest on the intended maritime system. In both cases, there would also be issues in simulating sensing and communications limitations. It is also rapidly becoming feasible to talk about field testing of affordable maritime systems including floats with limited mobility and the new generation of more affordable ocean gliders. Nevertheless, there are still challenges with keeping the systems affordable and easy to use when also integrating mission payload and communication components. A final option is to create remote access testbeds that researchers without specialized expertise can use to test out their software. This could be done through an interface that allows for experimentation while ensuring systems will remain safe.

4 Future directions

Building on these types of testbeds, there are a number of areas where advancement could be targeted to strengthen future long duration autonomy systems. First, there has been increasing work in recent years in decentralized and distributed autonomy methods that scale to larger numbers of systems. Some of these methods, like consensus and market-based techniques methods probably still require too much communications for typical maritime problems. However, others based on combinations of biological inspiration, game theory, information theory, and control theory may be suitable for adaptation. This is true not just for control purposes but also for capabilities such as communication-limited multi-platform simultaneous localization and mapping and decentralized perception. Second, there is a need for learning, planning, and exploration strategies that are feasible to use over large areas and time periods in dynamic environments. Such approaches must have the meta-cognition and introspection to deal with the unexpected, take advantage of exploration, and recover from mistakes and bad hypotheses. Third, there is a need for persistent environmental understanding, machine learning, and adaptation at useful timescales. Fourth, there is a need for improvements in understanding of the allocation of roles and responsibilities between humans and machines for these types of long duration missions. There also is a need for human interaction methods such as dialogue and explanation systems that support very long time periods of operations in constrained communications and sensing environments. Fifth, there is a need for improvements in high reliability verification and validation including methods that support composability of components and moving from design time verification to real-time verification. Real-time verification is particularly important for long-term missions because it may be easier to do verification when more is known about the platform's local constraints and situation rather than trying to cover all possible future circumstances. Similarly, there is a need for methods for testing and designing complex systems that increase the likelihood of finding worst-case situations or that at least bound the worst cases that may be encountered when using advanced autonomy methods.

As advances are made in these areas, it becomes possible to consider wholly new types of systems to explore and understand the ocean domain. Truly long duration systems may change the types of scientific questions that can be asked and the types of problems it is practical to solve. However, it will take advances across many different fields of engineering, computer science, oceanography, and human factors to make these new systems possible.