

A systematic approach to the development of long-term autonomous robotic systems for agriculture

Tim Bohne¹, Gurunatraj Parthasarathy¹ and Benjamin Kisliuk¹

Abstract: For robots to compete with conventional machinery in agriculture, improving their long-term autonomy seems necessary. In this article, we provide concepts and intermediate results of our work aimed at a long-term autonomous robotic system performing a plant monitoring task. Based on notions from literature, we introduce a structured approach to defining long-term autonomy and report on the system we develop that meets this definition in practical experiments. Finally, we present intermediate results from simulation and discuss further avenues of research.

Acknowledgements: The DFKI Niedersachsen (DFKI NI) is sponsored by the Ministry of Science and Culture of Lower Saxony and the Volkswagenstiftung. The paper describes work carried out in the context of the funded projects PORTAL (BMEL, FKZ: 28DK111B20) and ZLA (NiMWK, Volkswagenstiftung, ZDIN 11-76251-14-3/19).

Keywords: long-term autonomy, agricultural robotics

1 Introduction

The agricultural domain poses distinctive challenges for robotic systems to act independently of human supervision over prolonged periods of time due to the typical farm environment being neither under full control nor fully known to the system at any point in time. This might be the reason why most available products and prototypes deployed over long durations fulfill only very basic tasks and often require a controlled environment. However, as long-term autonomy (LTA) promises to be a cornerstone for robots competing with conventional machinery in more complex tasks in the future, it remains a challenge to develop flexible robotic systems for LTA deployment. To enable systematic improvement and evaluation of LTA capabilities, a sound definition of the concepts of *long-term* and *autonomy* in this context is essential. We reuse the AROX platform from [Ho20], seen in Figure 1a. Figure 1b shows the test site created in 2022 to facilitate long-term autonomy experiments by providing agricultural test beds in an enclosed area. The goal of the system is to accomplish the plant monitoring application described in [Ti22], i.e., to continuously acquire canola plant data.

¹ Deutsches Forschungszentrum für Künstliche Intelligenz (DFKI) GmbH, Plan-Based Robot Control, Berghoffstraße 11, 49090 Osnabrueck, Germany, <firstname.lastname>@dfki.de



Fig. 1: (a) AROX platform and base station [Ho20] (b) Test plot for canola monitoring

2 Relevant work and definitions

Although work regarding robotic LTA reaches back at least a decade [BKS13], the special case of LTA in agriculture as an application domain with very specific requirements [BV16] is gaining traction with many publications in recent years. Kunze et al. describe technological building blocks for addressing LTA in robotic systems, with work concerning agriculture focusing on navigation and mapping in changing environments [Ku18]. The following sections rely on the work described in [Pa22; Bo22].

2.1 What is autonomy?

In practice, robots are not expected to operate in complete isolation, but to act and self-preserve within the scope of their task without external supervision. Moreover, some functions of a robot may operate autonomously while others do not, rendering the integrated robot “partially autonomous”. Figure 2 shows concepts of different degrees of autonomy, derived from [He15]. For this work, the AROX should be able to operate fully autonomously for what we call an *LTA episode*, i.e., a period of time that is plausible in the context of plant monitoring.




Fig. 2: Spectrum of autonomy

2.2 How long is long-term?

Likewise, LTA is usually not expected to persist indefinitely, and other authors have defined varying time spans for their respective applications [Ha17; Ku18; Br18]. While the hardware and process can be considered static for the scope of our work, the environment introduces dynamics that occur on various time scales such as hours (day-night), days (predictable weather changes), weeks (plant growth), seasons (snow, leaf discoloration), and years (tree growth / removal) posing requirements for the implemented system. Moreover, robotic hardware constraints introduce a “temporal action radius” by enforcing interruptions for refueling, data storage clearing, or maintenance. A typical agricultural process might consist of multiple refueling cycles for single events such as harvesting, where a robot must return to a power supply station or unload harvested crops several times. For continuous applications, e.g., weeding or monitoring, such multi-cycle missions may be repeated and spread over a certain growth period spanning weeks. If other factors lead to multiple upper bounds, the lowest one defines the limit for an LTA episode. In our case, we found that the hardware maintenance interval sets a plausible upper bound of one week for LTA episodes.

2.3 Simplified LTA plant monitoring scenario

To approach the complex problem of LTA in a structured way, the core idea of our research is to start from a simplified practical scenario that allows to conduct experiments with a likewise simplified system that nevertheless satisfies the previously given definition. In earlier work, we described a monitoring application for plant breeding [Ti20], which is outlined in Figure 3. With a time frame of one week per LTA episode and a few hours per multi-cycle mission derived from practical experience, we can initially assume unexpected changes in energy consumption as the only dynamic to consider, which occurs frequently in practice.




Fig. 3: Multi-cycle mission for plant monitoring [Ti20]

3 Implementation of the system

In literature, there are overviews of robotic capabilities related to LTA that have been the subject of extensive research [Ku18; BV16], focusing on motion control, navigation, mapping, and localization for agricultural scenarios, for which a plethora of established software modules are available in the ROS ecosystem. For top-level control and plan execution, a simple state machine was implemented that can handle external interruptions from modular health monitoring solutions such as the battery watchdog (cf. sec. 3.2). Since robustness and verifiability of robotic systems is mentioned in many publications as a focus of future research, we anticipate this already in the design of the control architecture. Initially, we consider only energy consumption as an uncontrolled dynamic, so we concentrate on autonomous energy supply (cf. sec. 3.1) and a battery monitoring system (cf. sec. 3.2) as a substitute for other monitoring systems to be included in the future.

3.1 Autonomous energy supply

The ability to recharge is crucial for multi-cycle missions. Arvin et al. divide it into three subproblems: observing the charge level, localization and navigation to the charging station, and docking [ASR09]. The latter is implemented by a *Hough transform* based shape detection of the base station in a continuously updated laser scan (cf. fig. 4), upon which we compute a goal pose inside that can be adopted using the navigation module. As a prerequisite, the robot must be in the vicinity of the base station and roughly oriented towards it.




Fig. 4: (a) Robot facing base station (b) Base station perception based on lidar data (c) Successful detection of base station shape

3.2 Battery watchdog system

In practice, even with good planning, there is a risk of the robot not reaching the base station due to battery depletion because of unexpectedly high energy consumption. In this

work, this is ensured by a real-time battery watchdog system that interrupts current actions and drives to the base station for recharging when the battery level drops below a safety threshold, which is the sum of a pre-set value and a dynamically calculated prediction of how much energy the return to base will consume. This prediction is continuously updated using the navigation module, which means that improving the navigation planner will also improve battery management. If the remaining energy is determined to be insufficient to return to base safely, human assistance is called for by the robot. Naturally, the parameters for these calculations were optimized for a tradeoff between safety and not having to return to recharge too frequently.

4 Experiments and evaluation

At the time of writing this article, first successes have been realized in running the described components on the physical AROX platform, but no integrated runs have yet been conducted. Therefore, we report on experiments performed using the *Gazebo* simulator for ROS systems. When an LTA episode in a physical experiment were to span a week, this would roughly correspond to 2 to 4 missions repeated on non-consecutive days, with the time in between waiting in the base station. In simulation, this number of missions can be reasonably completed in about 5 hours of runtime by skipping idle time. A single multi-cycle mission is represented as a plan incorporating all the steps of the simplified scenario, including several planned returns to base for recharging and all the actions described. Moreover, it covers a variety of situations by navigating to different locations.

To evaluate the experiments, two metrics were used, namely *mission failures* and *autonomy percentage* based on [Ha17]. The autonomy percentage is the ratio of time spent performing autonomous tasks to the total runtime; this is expected to be higher than in real applications, though, as we omit the main waiting times. We define a mission as failed when a timeout $t = 900\text{ s}$ is exceeded without action completion. Also, any occurrence of a failure signal in the state machine is considered a failed mission, e.g., in case of a completely discharged battery. The experiment was performed three times with no failures, at an average of 5.09 hours, 3.67 completed missions, and 106.67 tasks. It required an average of 9.67 charge cycles, and the average total distance traveled was 1102.68 m. The autonomy percentage was 97.15 % on average. Although three runs do not guarantee flawless usage of the system, they demonstrate that the expected basic functionality is achieved in a setup meeting the definition of long-term autonomy.

5 Discussion and future avenues

The experiments show that the introduced definitions can be used for the development and evaluation of an LTA system. We argue that the choice of a particular time frame for an LTA episode depends on the combination of robot, application, and domain. In turn, the

challenges that must be solved for successful integration of such a system depend on the application and episode length. To narrow down the initial tests, we only considered energy supply and consumption. However, it can be argued that longer LTA episodes are of interest for actual productive use, leading to additional dynamics that need to be addressed in the future, especially along seasonal cycles. These include weather extremes, harsh ground conditions, lighting changes, obstacles (static / dynamic) blocking the robot's path, but also less tangible aspects such as perceptual aliasing, which involves recognizing certain objects or places under altered circumstances, which is far from trivial. In addition, system dynamics such as connectivity issues that do not necessarily affect the environment but the robot itself, must be managed adequately to achieve true long-term autonomy. We plan to integrate the described system into the physical AROX and extend it to enhance its LTA capabilities. This includes a more general framework for monitoring various relevant aspects of system health and the ability to deal with broader environmental dynamics. Additionally, the navigation and mapping modules will be improved to enable intelligent adaptation of the used planners.

Bibliography

- [ASR09] Arvin, F. et al.: Swarm robots long term autonomy using moveable charger. International Conference on Future Computer and Communication. IEEE, 2009.
- [BKS13] Barfoot, T.; Kelly J.; Sibley, G.: Special issue on long-term autonomy. IJRR, 2013.
- [Bo22] Bohne, T.: Execution Monitoring for Long-Term Autonomous Plant Observation with a Mobile Robot, Osnabrück University, 2022.
- [Br18] Brommer, C. et al.: Long-Duration Autonomy for Small Rotorcraft UAS Including Recharging. IEEE International Conference on Intelligent Robots and Systems, 2018.
- [BV16] Bechar, A.; Vigneault, C.: Agricultural robots for field operations: Concepts and components. Bio. Eng., 2016.
- [Ha17] Hawes, N. et al.: The STRANDS Project. IEEE Robotics & Automation Magazine, 2017.
- [He15] Hertzberg, J.: Technische Gestaltungsoptionen für autonom agierende Komponenten und Systeme. In: Das Recht vor den Herausforderungen der modernen Technik, 2015.
- [Ho20] Hoellmann, M. et al.: Towards Context-Aware Navigation for Long-Term Autonomy in Agricultural Environments. IEEE/RSJ International Conference on Intelligent Robots and Systems, 2020.
- [Ku18] Kunze, L. et al.: Artificial Intelligence for Long-Term Robot Autonomy: A Survey. IEEE Robotics and Automation Letters, 2018.
- [Pa22] Parthasarathy, G.: Towards Long-Term Autonomous Survivability in Agriculture: A real-time safety watchdog to ensure the survivability of an outdoor plant monitoring robot., KTH Royal Institute of Technology, 2022.
- [Ti22] Tieben, C. et al.: Erste Schritte zu einem virtuellen Zuchtgarten. 42. GIL-Jahrestagung. Bonn: Gesellschaft für Informatik e.V. (S. 295-300), 2022.