

Desiderata for Planning Systems in General-Purpose Service Robots

Nick Walker^{*1}, Yuqian Jiang^{*2}, Maya Cakmak¹, Peter Stone²

¹Paul G. Allen School of Computer Science & Engineering, University of Washington, Seattle, USA

²Department of Computer Science, University of Texas at Austin, Austin, USA

{nswalker, mcakmak}@cs.washington.edu

jiangyuqian@utexas.edu, pstone@cs.utexas.edu

Abstract

General-purpose service robots are expected to undertake a broad range of tasks at the request of users. Knowledge representation and planning systems are essential to flexible autonomous robots, but the field lacks a unified perspective on which features are essential for general-purpose service robots. Progress towards planning and reasoning for general-purpose service robots is hindered by differing assumptions about users, the environment, and the overall robot system. In this position paper, we propose desiderata for planning and reasoning systems to promote general-purpose service robots. Each proposed item draws on our experience with research on service robots in the office and home and on the demands of these environments. Our desiderata emphasize support for natural human-interfaces as well as for robust fallback methods when interactions with humans and the environment fail. We highlight relevant work towards these goals.

Introduction

Creating a robot that is able to carry out a broad range of tasks in a household or office environment is a long-standing grand challenge for A.I. and robotics. Envisioned in science fiction for example as “Rosie the Robot” in the *Jetsons*, and the focus of competitions such as RoboCup@Home, we refer to such a futuristic embodied personal assistant as a *General Purpose Service Robot* (GPSR). A GPSR, like the one shown in Figure 1, integrates skills for navigating, observing and manipulating the environment, and interacting with users to service requests.

The problem of coordinating low-level skills to accomplish complex goal-driven behavior is often addressed using AI planning techniques. AI planning systems reason about a robot’s feasible actions and generate a trajectory in an abstract action space based on models of the environment. These systems define structures that are used in all aspects of the robot’s behavior including: the actions it can take and their effects on the world, the information it gathers, and attributes of the environment it perceives. For example, when a GPSR is asked to “Grab a fruit from the kitchen,” it might rely on its current knowledge about the layout of the environment and what objects are in the kitchen to produce a preliminary plan. As the GPSR executes the plan, its perception modules will store information in its representation

^{*}Equal contribution.



Figure 1: A Toyota Human Support Robot, representative of typical GPSRs, handing an apple to a user in a mock home environment.

of knowledge. If an action fails, the robot may invoke a reasoning process to infer the source of the failure, or to suggest alternative goals it could pursue. As this example illustrates, the knowledge representation and planning system play a critical role in orchestrating the complex behavior of a general-purpose robot. A well designed system could be beneficial to a range of tasks in various environments across different platforms.

Building a general-purpose service robot is a difficult task. In practice, researchers focus their efforts on individual aspects of the general problem with the hope that progress on a part will contribute to progress on the whole. As a result, the community pursuing service robots has sought solutions to many different problems, at the expense of progress towards integrated general-purpose service robots. Unlike other subsystems which are difficult to disentangle from hardware or particular environments, the knowledge representation and planning system of a service robot is a central and portable component that can be designed to address a wide set of requirements. It can undergird progress towards more general robots.

Three main factors give rise the constraints that GPSRs must operate in. Differing assumptions about these factors can motivate different designs of GPSR planning systems:

- Users. Researchers make varying assumptions about how users will interact with robots. These motivate disparate

approaches. What will users actually expect from systems? Can we derive from these a common set of functions that our systems should support?

- The environment. Expense and technical challenge makes it difficult to field robots in multiple environments. Often, researchers design systems to be flexible but rely on assumptions about the domains they are targeting. Which of these assumptions will still hold in other domains? Can we consolidate characteristics of various environments to better understand what features of planning systems support generalization across different domains?
- The robot system. Platforms are highly diverse in their design and capabilities, and it's rare that researchers have access to more than a handful. As a result, they may assume a level of computation or a level of connectivity that doesn't hold with similar platforms. Can we extract the core set of hardware constraints that service robot planning systems should assume?

Progress towards the creation of planning systems for general-purpose service robots depends on our ability to agree on a basic set of goals. Therefore, an important task for the AI planning community should be understanding the requirements and constraints of general-purpose service robots. Towards this end, we propose desiderata for knowledge representation and reasoning systems in general-purpose service robots. Our purpose is to make explicit a position on how users, the environment, and the robot system drive the design of the knowledge representation and planning component. We hope that these points prompt further discussion and lay a path for work towards the creation of practical service robots in a wide range of environments.

Desiderata

A general-purpose service robot is a capable and intelligent helper. In the hospital, such a robot may be asked to fetch supplies. At home, a GPSR might set the table or store groceries. In the office, a GPSR might scout meeting rooms or locate coworkers. The key aspect of a GPSR is that its tasking is flexible. If a robot has some set of skills, we would designate it a general-purpose robot if it can sequence these skills to achieve more complex behavior, even when the particular goal has not been encountered before and the environment or context changes. Progress in individual capabilities like autonomous navigation and localization, object manipulation or visual recognition are all necessary to implement such a system, but in order to achieve generality, these components must operate at the behest of a larger planning system.

Researchers have long pursued general-purpose service robots. Since 2006, the international RoboCup@Home competition has challenged teams to field domestic service robots in a mock apartment. The competition tasks have included entertaining guests, cleaning, as well as an explicit *General Purpose Service Robot* challenge, where the robot carries out tasks that an operator requests in spoken language (Holz and Iocchi 2013). This task is designed to confront the robot with key challenges from the GPSR problem,

such as interacting with the user to clarify a goal, planning without full knowledge of objects in the environment, and handling execution time failures. Success requires comprehensive integration of a robot planning and reasoning system with a real robot system that can operate under fixed time constraints. This task is perhaps the largest benchmarking effort for general-purpose service robots, but this vision is shared widely in the robotics research community (Alterovitz, Koenig, and Likhachev 2016; Agostini, Torras, and Wörgötter 2017; Torras 2016).

In the following, we describe some of the essential aspects of GPSRs and how they lead to our desiderata.

Users

Whether in a domestic, a medical, or a scholastic setting, GPSRs will interact primarily with non-expert users. These users may have detailed knowledge of the environment and their roles, but will have minimal familiarity with robots. This attribute drives several desiderata:

- *Human-interfaces*: GPSR planning systems must be able to accept commands and information from a wide variety of users. Non-experts will not be able to provide input in planning systems' native representations, but will instead need to use special interfaces. Such interfaces should facilitate all manner of user demands, including goals, constraints on planning and scheduling, and suggestions. Research towards the design of these interfaces continues, spanning natural language interaction (Cantrell et al. 2012; Thomason et al. 2019) to graphical user interfaces (Huang and Cakmak 2017). Planning systems must be designed to support one or more of these human-interfaces.
- *Active error resolution*: Often, commands given by non-expert users cannot be formulated in the robot's representation because they are incomplete or ambiguous. Instead of discarding the command and requiring the user to start over, it's preferable to actively resolve these problems through dialog or other interactions with the user.
- *Responsiveness*: Users will expect the robot to service commands promptly. Planning processes must not impede the robot's ability to quickly carry out tasks. If there are further interactions during the command taking phase, it is critical that queries can be answered within reasonable time to keep the users engaged. Even during execution, the system must always be ready to respond to interruptions from users asking questions or making new requests.
- *Online adaptation*: User needs are diverse and change over time. As a consequence, a GPSR will be asked to carry out tasks for which it does not have adequate skills or planning knowledge. A GPSR's planning system must facilitate the acquisition of both new actions as well as the knowledge necessary to plan with these actions, either through end-user programming (Huang and Cakmak 2017), interactive learning (Laird et al. 2017), or other methods.
- *Explainability*: Explanations of the robot's plan can help build trust with the users. The planning system should be

able to answer questions about the generated plans, and in case execution fails, explain what in the environment is different from the robot's model at planning time (Fox, Long, and Magazzeni 2017).

The Environment

We expect that GPSRs will find use in a wide range of environments, but a constant across these is the presence of people. Beyond handling robot-to-operator interactions, the robot must be able to cope with the demands of human-populated environments. This leads to additional desiderata:

- *Interpretability*: While people are skilled at perceiving each other's gross intentions, they are not generally capable of discerning the intentions of robots. Robots carrying out complex behaviors autonomously for long durations will create uncertainty in passersby. In addition to being friendly to direct users, a GPSR planning and reasoning system must include transparency mechanisms that support producing interpretable plans to build trust and minimize disturbance to humans in the environment.
- *Ethical*: GPSRs must respect the values of the people around them. A recent European Commission report identified respect for human autonomy, prevention of harm, fairness and explicability as the core ethical principles that trustworthy AI systems must uphold (2019). Balancing adherence to ethical concerns in general will be a challenging task, so GPSR planning and reasoning systems must support straightforward encoding and tuning of these principles.
- *Safety*: GPSRs will be expected to guarantee that their actions bring no harm to the humans in their environment. Ensuring safety is challenging for any autonomous robot, but is particularly difficult given the dynamic nature of a GPSR's tasking. How can planning systems help ensure that the robot only undertakes safe tasks?
- *Resource efficiency*: GPSRs will operate in domains that have inherent resource constraints, such as the number of elevators, or the available supply of human patience. Planning systems that interact with these limited resources must model and optimize for the costs of the robot's actions.
- *Handle uncertainty*: Since GPSRs operate in highly dynamic environments, the robot's model at planning time often turns out to be inaccurate during execution. The planning system should therefore be able to robustly handle these uncertainties at planning time or through replanning at execution time.
- *Adaptability*: Some deviations in a GPSR's environment will represent a major or lasting change. Beyond simply tolerating these at execution time, the planning system should incorporate this information and not generate plans that are likely to fail given the current state of the environment. Rather, the planning system should generate better initial plans that adapt to new situations, based on recent experience.
- *Open-world*: The real world for a GPSR is "open"; the robot's knowledge base cannot capture the full state of the

environment. Certain information may not be available to a GPSR at planning time, such as the current location of an object that the user requested. Therefore, the planning system should not make the "closed-world" assumption and should be capable of planning for knowledge acquisition.

Robot System

A robot is a complex combination of hardware and software components, and each is unique in its set of tradeoffs. Planning systems operate atop representations that are notionally abstract to the particulars of a platform, however, in practice, the overall robot system still imposes important constraints on the design of the planning system. Work towards practical planning systems for GPSRs must strive for:

- *Robustness*: An unexpected failure in a GPSR's planning system will most immediately erode user trust and faith. Further, inopportune issues have the potential to endanger the robot or its surroundings, and the long duration of deployments makes such occurrences almost unavoidably probable if not adequately addressed. Planning and reasoning systems for GPSRs must achieve high standards of robustness, making them a prime target for formal verification efforts.
- *Portability*: Robots have unique strengths and weaknesses, but the core set of necessary capabilities for GPSRs are fixed; the robot will have to navigate through and manipulate its environment, as well as interact with people. To enable portability, planning systems should model at least this minimum set of capabilities, regardless of the details of particular platforms.
- *Compute efficiency*: Mobile robots are constrained by battery technology. In practice, the majority of a GPSR's power budget is spent on computation, which leads robot designers to adopt limited, power efficient compute solutions. As a result, GPSR planning systems must be able to operate within a limited compute budget.
- *Connectivity tolerance*: Network reliability in robotics laboratory environments is an upper bound on the conditions that real GPSRs will encounter when deployed. Though standards and hardware for consumer grade networking continue to improve, robots' preferred bandwidth and signal characteristics will not be well served by typical equipment for some time. As a result, practical GPSR planning systems cannot assume reliable network connectivity.

Recent Directions

There is still a significant gap between existing systems and the fulfillment of the desiderata. Taking as an example work from Khandelwal et. al. (2017) and Thomason et. al. (2019), it is possible to achieve high levels of responsiveness and support dialog interactions in a deterministic classical planning regime. The efficiency enabled by this paradigm however comes at the expense of rich expression of the probabilistic aspects of the environment. The authors

could have pursued probabilistic systems that have straightforward means for expressing the openness of the world or the uncertainty of the robot's knowledge, but with current methods, they would have necessarily found the responsiveness of the resulting system limited.

Besides difficulties in integrating existing technologies, fulfillment of the desiderata is challenging because some items have received less attention from the planning community and have few mature approaches. In the remainder of this section, we highlight directions where there is ongoing work and opportunities for progress towards our desiderata. For a comprehensive survey on advances in robot planning technologies we refer the reader to the recent survey by Ingrand and Ghallab (2017).

There is growing interest in human-aware and explainable planning (Sreedharan, Kambhampati, and Others 2017). Recent progress includes new problem formulations, metrics, and initial approaches in generating explicable plans (Zhang et al. 2017) and plan explanations (Chakraborti et al. 2017). Recent works have advanced techniques for learning and refining planning operators (Agostini, Torras, and Wörgötter 2017), a critical component in interactive learning. Efforts to refine planning components in cognitive architectures are creating some of the first integrated interactive learners (Mininger and Laird 2016).

Planning under uncertainty is attacked with many different methods. There have been recent efforts to combine logical and probabilistic planning and reasoning methods in integrated architectures for robots (Sridharan et al. 2018). A state-of-the-art integrated robot task planning system by Hanheide et al. (2017) leverages a probabilistic planner to tackle the issues of reasoning in an open and uncertain world, as well as interactive goal acquisition and failure explanation. Recent work approaches adaptive planning by feeding information gathered during plan execution, such as action costs and user preferences, back to the planning system (Jiang et al. 2018; Wilde, Kulić, and Smith 2018). Other methods for combining learning and planning have demonstrated that deliberative controllers can enhance the exploration capabilities of deep reinforcement learning agents in complex visual simulators (Gordon, Fox, and Farhadi 2019).

The importance of speed and responsiveness are recognized in the planning community, and benchmarked in the International Planning Competition (López, Celorrio, and Olaya 2015; Vallati et al. 2015). Some researchers are beginning to consider how to design systems that work across related domains and platforms (Lima, Ventura, and Awaad 2018; Hart et al. 2018). Numerous other deployed systems with planning capabilities highlight enthusiasm for creating robust systems that can operate autonomously for long durations in real environments (Veloso et al. 2015; Hawes et al. 2017; Khandelwal et al. 2017).

Conclusion

The objectives laid out in this paper are challenging, but they also present a unique opportunity for collaboration in the planning and robotics communities. Research that advances any of the desiderata is valuable, but progress towards general-purpose service robots is enhanced when we

actively consider how our contributions interface with an integrated system. We hope that by laying out the most important objectives for a general-purpose service robot, the community will not only pursue each direction as an intriguing problem on its own, but also develop new approaches and integrated systems that aim to address the full set of desiderata.

References

- Agostini, A.; Torras, C.; and Wörgötter, F. 2017. Efficient interactive decision-making framework for robotic applications. *Artificial Intelligence* 247:187–212.
- Alterovitz, R.; Koenig, S.; and Likhachev, M. 2016. Robot planning in the real world: research challenges and opportunities. *Ai Magazine* 37(2):76–84.
- Cantrell, R.; Talamadupula, K.; Schermerhorn, P.; Benton, J.; Kambhampati, S.; and Scheutz, M. 2012. Tell me when and why to do it!: Run-time planner model updates via natural language instruction. In *Proceedings of the Seventh Annual ACM/IEEE International Conference on Human-Robot Interaction, HRI '12*, 471–478. New York, NY, USA: ACM.
- Chakraborti, T.; Sreedharan, S.; Zhang, Y.; and Kambhampati, S. 2017. Plan Explanations As Model Reconciliation: Moving Beyond Explanation As Soliloquy. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence, IJCAI'17*, 156–163. AAAI Press.
- Fox, M.; Long, D.; and Magazzeni, D. 2017. Explainable Planning. In *IJCAI-17 Workshop on Explainable AI (XAI)*, 24.
- Gordon, D.; Fox, D.; and Farhadi, A. 2019. What should I do now? marrying reinforcement learning and symbolic planning. *CoRR* abs/1901.01492.
- Hanheide, M.; Göbelbecker, M.; Horn, G. S.; Pronobis, A.; Sjö, K.; Aydemir, A.; Jensfelt, P.; Gretton, C.; Dearden, R.; Janicek, M.; Zender, H.; Kruijff, G. J.; Hawes, N.; and Wyatt, J. L. 2017. Robot task planning and explanation in open and uncertain worlds. *Artificial Intelligence*.
- Hart, J.; Yedidsion, H.; Jiang, Y.; Walker, N.; Shah, R.; Thomason, J.; Padmakumar, A.; Fernandez, R.; Sinapov, J.; Mooney, R.; and Stone, P. 2018. Interaction and Autonomy in RoboCup@Home and Building-Wide Intelligence. In *AAAI Fall Symposium on Artificial Intelligence and Human-Robot Interaction*.
- Hawes, N.; Burbridge, C.; Jovan, F.; Kunze, L.; Lacerda, B.; Mudrova, L.; Young, J.; Wyatt, J.; Hebesberger, D.; Kortner, T.; Ambrus, R.; Bore, N.; Folkesson, J.; Jensfelt, P.; Beyer, L.; Hermans, A.; Leibe, B.; Aldoma, A.; Faulhammer, T.; Zillich, M.; Vincze, M.; Chinellato, E.; Al-Omari, M.; Duckworth, P.; Gatsoulis, Y.; Hogg, D. C.; Cohn, A. G.; Dondrup, C.; Fentanes, J. P.; Krajník, T.; Santos, J. M.; Duckett, T.; and Hanheide, M. 2017. The STRANDS Project: Long-Term Autonomy in Everyday Environments. *IEEE Robotics Automation Magazine* 24(3):146–156.
- Holz, D., and Iocchi, L. 2013. Benchmarking Intelligent Service Robots through Scientific Competitions : The RoboCup @ Home Approach. Technical report.

- Huang, J., and Cakmak, M. 2017. Code3: A System for End-to-End Programming of Mobile Manipulator Robots for Novices and Experts. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction*, 453–462.
- Ingrand, F., and Ghallab, M. 2017. Deliberation for autonomous robots: A survey. *Artificial Intelligence* 247:10–44.
- Jiang, Y.; Yang, F.; Zhang, S.; and Stone, P. 2018. Integrating Task-Motion Planning with Reinforcement Learning for Robust Decision Making in Mobile Robots. *arXiv preprint arXiv:1811.08955*.
- Khandelwal, P.; Zhang, S.; Sinapov, J.; Leonetti, M.; Thomason, J.; Yang, F.; Gori, I.; Svetlik, M.; Khante, P.; Lifschitz, V.; Aggarwal, J. K.; Mooney, R.; and Stone, P. 2017. BWIBots: A platform for bridging the gap between AI and human–robot interaction research. *The International Journal of Robotics Research (IJRR)* 36(5-7):635–659.
- Laird, J. E.; Gluck, K.; Anderson, J.; Forbus, K. D.; Jenkins, O. C.; Lebiere, C.; Salvucci, D.; Scheutz, M.; Thomaz, A.; Trafton, G.; Wray, R. E.; Mohan, S.; and Kirk, J. R. 2017. Interactive Task Learning. *IEEE Intelligent Systems* 32(4):6–21.
- Lima, O.; Ventura, R.; and Awaad, I. 2018. Integrating classical planning and real robots in industrial and service robotics domains. In *The 6th ICAPS Workshop on Planning and Robotics (PlanRob-18)*.
- López, C. L.; Celorrio, S. J.; and Olaya, Á. G. 2015. The deterministic part of the seventh International Planning Competition. *Artificial Intelligence* 223:82–119.
- Mininger, A., and Laird, J. 2016. Interactively Learning Strategies for Handling References to Unseen or Unknown Objects. *Advances in Cognitive Systems* 5:1–16.
- Sreedharan, S.; Kambhampati, S.; and Others. 2017. Balancing explicability and explanation in human-aware planning. In *2017 AAAI Fall Symposium Series*.
- Sridharan, M.; Gelfond, M.; Zhang, S.; and Wyatt, J. 2018. REBA: A Refinement-Based Architecture for Knowledge Representation and Reasoning in Robotics.
- The European Commission. 2019. *Ethics guidelines for trustworthy AI*. The European Commission.
- Thomason, J.; Padmakumar, A.; Sinapov, J.; Walker, N.; Jiang, Y.; Yedidsion, H.; Hart, J.; Stone, P.; and Mooney, R. J. 2019. Improving Grounded Natural Language Understanding through Human-Robot Dialog. *arXiv preprint arXiv:1903.00122*.
- Torras, C. 2016. Service Robots for Citizens of the Future. *European Review* 24(1):17–30.
- Vallati, M.; Chrupa, L.; Grześ, M.; McCluskey, T. L.; Roberts, M.; and Sanner, S. 2015. The 2014 international planning competition: Progress and trends. *Ai Magazine* 36(3):90–98.
- Veloso, M.; Biswas, J.; Coltin, B.; and Rosenthal, S. 2015. CoBots: Robust Symbiotic Autonomous Mobile Service Robots. In *Proceedings of the 24th International Conference on Artificial Intelligence, IJCAI'15*, 4423–4429. AAAI Press.
- Wilde, N.; Kulić, D.; and Smith, S. L. 2018. Learning user preferences in robot motion planning through interaction. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 619–626. IEEE.
- Zhang, Y.; Sreedharan, S.; Kulkarni, A.; Chakraborti, T.; Zhuo, H. H.; and Kambhampati, S. 2017. Plan explicability and predictability for robot task planning. In *Proceedings - IEEE International Conference on Robotics and Automation*, 1313–1320.