Robot Planning in the Real World: Research Challenges and Opportunities

Ron Alterovitz, Sven Koenig, Maxim Likhachev

■ *Recent years have seen significant technical progress on robot planning, enabling robots to compute actions and motions to accomplish challenging tasks involving driving, flying, walking, or manipulating objects. However, robots that have been commercially deployed in the real world typically have no or minimal planning capability. These robots are often manually programmed, teleoperated, or programmed to follow simple rules. Although these robots are highly successful in their respective niches, a lack of planning capabilities limits the range of tasks for which currently deployed robots can be used. In this article, we highlight key conclusions from a workshop sponsored by the National Science Foundation in October 2013 that summarize opportunities and key challenges in robot planning and include challenge problems identified in the workshop that can help guide future research toward making robot planning more deployable in the real world.*

More than 50 years have passed since the first industrial robot began service in a car assembly line and since development began on Shakey, the first robot capable of running the full cycle of autonomy, from sensing to planning to execution. Since that time, robotics has grown into a multibillion dollar worldwide industry. In addition to industrial robotics companies such as ABB, KUKA, Yaskawa, and FANUC, in recent years a variety of companies such as Google, Intuitive Surgical, Amazon, SoftBank, iRobot, Apple, and Uber are increasingly investing in robotics in a variety of application areas, from warehouse management to medicine to home assistance to transportation. With advances in research, the next generation of robots has the potential to improve performance in established domains and create entirely new applications.

Achieving the full potential of robotics will require improving the ability of robots to reason about how to accomplish a task, process sensor data in real time, utilize available resources effectively, cooperate with humans, and

  ISSN 0738-4602

adapt to changes in the environment. An important building block of robots with these desirable capabilities is robot planning. A plan is "a detailed proposal for doing or achieving something" (Stevenson and Lindberg 2010), and planning in robotics typically corresponds to computing actions and motions for a robot to achieve a specified objective. In 1992, Drew McDermott noted that planning requires "reasoning about possible courses of execution" (McDermott 1992). Robot planning is often necessary for navigating through an environment, manipulating tools and objects, maintaining safety around humans, and gathering information necessary to complete a task (Lavalle 2006). Typically, a human user provides the robot with a high-level description of the task objectives and the robot planner computes actions and/or low-level motions for the robot to autonomously or semiautonomously accomplish the task. Robots have to plan on multiple interacting levels, from low-level control to high-level task planning. The robot's planner is tightly integrated with a robot's other components; the robot planner utilizes information from the robot's sensing and perception systems, might be guided by the input of an operator, and outputs instructions to the control and actuation systems. Thus, progress in robot planning must go hand in hand with progress in computer vision, human-computer interfaces, and other areas. Robot planning is a critical component of enabling full robot autonomy, or it can facilitate shared autonomy in which a human and the robot share control. Robust and effective robot planning capabilities are needed to realize fully the potential of robots in a wide variety of applications, including self-driving vehicles, disaster response, minimally invasive surgery, and assistance for people in their homes and workplaces.

Although recent years have seen substantial technical progress on robot planning, robots that have been commercially deployed on a large scale in the real world typically have no or minimal planning capability. These commercial robots are typically manually programmed for a specific task, teleoperated, or programmed to follow simple rules, as is the case for most manufacturing robots (such as automobile assembly manipulators), medical robots (such as Intuitive Surgical's da Vinci System), and special-purpose home assistance robots (such as the iRobot Roomba). These commercial robots are highly successful in their respective niches. Although progress is being made (especially for self-driving vehicles), the lack of robust and effective planning capabilities is limiting the range of tasks for which robots can be used for long-term autonomous operation in real-world settings with unstructured environments.

There is currently a substantial gap between the potential of robot planning to enable exciting robotics applications and the reality of the limited deployment of robot planning in the real world. Researchers from robotics and artificial intelligence study closely related planning topics (see, for example, the textbook *Planning Algorithms* by robotics researcher LaValle [2006] and the textbook *Automated Planning: Theory and Practice* by artificial intelligence researchers Ghallab, Nau, and Traverso [2004]) and their cooperation can thus help to close this gap. In this article we summarize key conclusions from a workshop[1] sponsored by the United States National Science Foundation held in October 2013. The workshop included robotics and artificial intelligence researchers and practitioners from academic institutions, government agencies, and industry (see the Acknowledgment for the list of participants who all contributed to the ideas presented in this article). The expertise of the participants roughly reflected the current distribution of research directions on robot planning, ranging from lower-level motion planning to higher-level task planning to planning under uncertainty and to interaction between planning, perception, and control. We discuss application areas, new opportunities, key research challenges, and specific challenge problems involving robot planning that can help guide future research toward making robot planning more deployable in the real world.

Applications and Opportunities for Robot Planning

Improvements in robot planning could help improve the capabilities of currently deployed robots and create opportunities for new robotics applications. We begin by presenting a survey of real-world robotics applications and how advances in robot planning could help.

Manufacturing

The needs of the manufacturing industry led to the birth of modern robotics; the first industrial robot entered service in a General Motors assembly line in New Jersey in 1961. Most robots used in manufacturing are manually preprogrammed to perform repetitive tasks rapidly and independently for a large volume of goods in fenced-off spaces. Improvements in robot planning could contribute to a new generation of manufacturing robots that operate cooperatively with humans, for example by autonomously computing safe motions that meet human expectations or by effectively splitting the burden of completing a task through shared autonomy. Robot planning could also be used in nimble factories with rapidly changing products and needs by facilitating quick adaptation to new tasks and reducing the effort of manually reprogramming robots if workspaces or products are modified. Creating robots with the planning capabilities needed for these new scenarios will require research on manipulation planning, efficient user interfaces for conveying how tasks should be performed, human-robot cooperation, enabling situational awareness, compensating for environmental and operational uncertainty, and assuring performance.

Warehouse Automation

Most warehouses today are labor intensive, and robotics has the potential to make warehouses operate more efficiently. Kiva Systems, which was acquired by Amazon.com for more than $700 million,[2] uses fleets of small robots to move inventory shelves (pods) around in warehouses. The robots carry inventory shelves to people on the perimeter of the warehouse who manually complete tasks (such as placing items in boxes and/or replenishing the shelves). Kiva's robots make the workforce 2 to 3 times more efficient by eliminating walking on the warehouse floor, but more sophisticated planning technology could reduce the number of robots needed and thus result in further cost savings. Robot planning is already used for coordinating the motions of the many small robots. Advances in robot planning could help enable robots autonomously to place items in boxes and replenish shelves through improvements in manipulation planning and better integration with a robot's sensing and perception systems.

Medicine

Medical robots have the potential to augment the capabilities of physicians and enable new medical procedures with fewer negative side effects. Intuitive Surgical's commercially successful da Vinci system allows surgeons to teleoperate endoscopic instruments with improved accuracy and precision. New snakelike and tentaclelike medical robots could maneuver along curved, winding paths to reach anatomical targets in highly constrained spaces, enabling minimally invasive access to previously unreachable sites. Robot planning could help medical robots reach their full potential by facilitating intuitive operation of complex robots, for example, by passively suggesting a path for the robotic mechanism to follow, by actively guiding the surgeon's motion to respect motion constraints, or by guaranteeing safety by automatically avoiding anatomical obstacles and sensitive structures. Robot planning for medical applications is challenging because of complex constraints on robot motion, large robot configuration spaces, the common need to pass through highly constrained spaces, the need to reason about uncertainty and deformable environments, and the need for fast, high-quality plans with safety guarantees.

Personal Assistance

Personal robots have the potential to assist people with a variety of tasks in homes and workplaces. Assisting people with activities of daily living (such as eating and cleaning) costs the economy more than $350 billion each year in the United States alone (Kassner et al. 2008), and these costs will continue to grow not just in the United States but in many nations as the aging and disabled population increases. Personal robots with manipulation capabilities could assist people with activities of daily living, thus enabling the elderly and people with disabilities to remain independent in their homes longer without needing to move to expensive institutions. Personal robots with the ability to navigate in human environments (without necessarily needing arms for manipulation) could be used in workplaces, museums, and public spaces as guides, escorts, or automatic transport (for example, a robotic wheelchair). Advances in robot planning could help enable personal robots to efficiently, autonomously navigate and manipulate objects in people's homes and other environments designed for humans. Navigating and manipulating objects in environments designed for humans raises numerous challenges for robot planning. The robot planner must utilize information from the sensing and perception systems in real time, must be fast and reactive, and must consider the presence of humans and animals in the environment. The robot planner must also handle unstructured and dynamic environments and consider uncertainty. Furthermore, the robot should generate consistent, intuitive plans such that humans in the environment can safely anticipate the robot's motions. The robot planner must also take as input a vague description of a task and create a plan that satisfies the user's intent in a manner that is flexible and robust.

Transportation

Car accidents kill more than 30,000 people each year in the United States.[3] Robotic, self-driving vehicles have the potential to reduce death and injury due to car accidents and to increase the efficiency of our road network, particularly in highly congested urban areas. Autonomous ground, water, and air vehicles (for example, quadrotors) could also be used in other transportation-related contexts, including package delivery (as proposed by Amazon),[4] exploration, security patrols, and industrial tasks (for example, object transport with a forklift). Exciting progress has been made in recent years. The DARPA Grand and Urban Challenges were completed successfully by multiple entrants and the Google driverless cars have already completed more than 300,000 miles of autonomous driving.[5] But key challenges remain before self-driving cars and other autonomous vehicles will be widely adopted. Robot planning is a necessary component of an autonomous vehicle, and the integration of perception and planning needs to be improved. Autonomous vehicles need to understand real-world uncertainties better, and planners must be robust to uncertainties arising from limitations in robot perception capabilities.

Disaster Response

Robots have the potential to assist in a variety of emergency response situations, including search and rescue operations, firefighting, bomb diffusion, and surveillance. Although robots can already be used in many of these situations under teleoperation, robot planning could help enable wider deployment in the real world by requiring less human effort to make robots accomplish their tasks and by making the robots easier to use in high-stress, dynamically evolving disaster response situations. A variety of robot architectures are relevant to disaster response, including ground vehicles, aerial vehicles, underwater vehicles, humanoid robots, and snakelike robots. Key challenges for robot planning include the ability to operate in a team with other robots and/or humans, handling large degrees of freedom with significant constraints on motion (for example, for humanoids and snake robots), situational awareness, integrating perception with planning, real-time performance, and the need to operate at a tempo beyond the capabilities of most current robotics systems.

Surveillance and Monitoring

Robots have the potential to assist with tasks such as surveillance, inspection of structures (both on the ground and underwater), and environmental monitoring. Aerial vehicles such as drones and quadrotors are ideally suited for above-ground surveillance and monitoring tasks. Key challenges for robot planning are similar to the challenges for applications such as transportation and disaster response. The planning algorithms will need to enable a robot to operate in a team with other robots and/or humans, to integrate perception with planning, to compute and execute motions in real time, and to request help from humans when necessary.

Emerging and Nonrobotics Applications

Robot planning is likely to be used in many additional robotics applications, some of which have not yet been thought of. Emerging robotics applications that could benefit from enhanced robot planning capabilities include construction of structures, automated farming, physical therapy, and rehabilitation. These applications combine the needs of other applications, including transportation, personal assistance, and disaster response. They also introduce new challenges, including manipulative interaction with nature and coordination of large teams of robots and humans. Advances in robot planning could also be integrated with education; robots with integrated planning can help inspire interest in computer science and engineering in children. Planning algorithms are also used for applications beyond robotics. Robot planning algorithms have made their way into diverse applications, such as modeling protein folding, animating agents for games and virtual environments, and simulating large crowds for optimizing security and emergency evacuation procedures. Thus, addressing the research challenges in real-world robotics applications will likely benefit other domains as well.

Enhancing and Expanding Robot Capabilities

Advances in robot planning will have significant impact on a variety of robot capabilities which span multiple existing and emerging robotics applications. Planning is critical for the long-term autonomous operation in unstructured environments, which — depending on the application — may require manipulation of objects and tools, navigation, maintaining safety around humans, information gathering, multirobot coordination, and human-robot teaming (for example, in the context of sliding autonomy, offering and requesting advice and help, and providing information). The research communities in robotics and artificial intelligence have investigated planners that enable these core capabilities, but there remains a large gap with respect to applicability in real-world conditions. For example, significant progress has been made on handling single rigid objects, but current methods are often not robust for handling collections of rigid objects (such as needed for packing boxes) or handling deformable objects in real-world scenarios. Similarly, significant progress has been made on robot navigation for ground, aerial, and marine vehicles, but less progress has been made on navigation in dynamic and unstructured environments with time pressure. Examples of challenging scenarios that require these capabilities include navigating in the presence of people (for example, a mobile robot navigating through a crowd of people), under extreme conditions (for example, docking a sea-surface vehicle in a rain storm), and under time constraints (for example, an aerial vehicle performing complex maneuvers). Moving in a simple, elegant, and agile style by effectively utilizing dynamics is a robot capability that has not been studied sufficiently. Additionally, robots should be able to explain their behavior, the plans need to be understandable by humans, and the planner should be able to quantify how likely it is to succeed and communicate this information to the human user if necessary. Although the subareas of robot planning have been studied to varying degrees, each subarea still includes unsolved problems that are important for broadening robot deployment in the real world.

Research Challenges for Robot Planning

We next list some of the important research challenges in robot planning that need to be addressed to achieve the full potential of robotics in the real-world applications discussed above. Addressing these research challenges requires progress on robot planning algorithms and proper planning representations rather than just more processor speed and memory. In general, planning methodologies differ depending on the type of robot and the type of environment the robot operates in. For example, mobile manipulators, humanoids, and snakelike robots often require planning in higher-dimensional configuration spaces relative to wheeled and flying robots. Similarly, indoor environments are often more complex and cluttered relative to outdoor environments, although indoor environments are sometimes more structured. These differences lead to different approaches to robot planning. However, many robot planning research challenges (including many of those outlined below) are shared across multiple robot types and environments.

Tight Integration of Planning with Perception

In many domains, the bottleneck to robot autonomy lies with perception. For example, many people share the view that automatic, detailed image understanding will not be fully solved anytime soon for scenes encountered by unmanned ground robots. Instead, robot planning for unmanned ground robots should explicitly deal with the uncertainty the robot has in its perception of the world. Similarly, lightweight micro-aerial vehicles and surgical robots have poor or limited sensing capabilities, and planning must compensate for this. A challenge is how to plan with uncertainty in perception in a way that scales, especially when it is hard to quantify the uncertainty. Are there robot planning representations that are amenable to real-time requirements on planning yet capture critical elements of uncertainty?

Modern robots are sometimes equipped with an array of sensors, many of which are controllable either directly (for example, controlling a servo) or by repositioning or reconfiguring the robot itself (for example, moving an arm equipped with a camera or tactile skin). This may lead to massive amounts of incoming sensory data. Some of the data may even be contradictory due to noise in sensing. To help with uncertainty in perception, robot planning should reason about when and how the robot can control its sensors in order to obtain information that disambiguates the uncertainty that jeopardizes the robustness of the robot completing its task.

How Should Planning Represent Infinite Dimensionality

The world has infinite dimensionality. How should planning represent it? As a robot moves in the real world, the planner faces the challenges of what it should model in its environment as well as when and how. For example, a typical kitchen may contain hundreds of relevant objects, such as pots, dishes, utensils, and food items. A personal assistance robot operating in the kitchen should not have to model all of these objects for planning a specific task. Furthermore, even if the robot could model everything computationally, the question is how these objects should be modeled in the first place. Geometric information about the world is relatively easy to obtain and represent, but the physical behavior of objects — for example, their articulation or deformability — and the affordances of objects are much harder to represent and estimate. In medical robotics applications as well as cooking applications for home assistance, understanding the deformation of objects such as tissues or foods is often critical to task success. The brittleness of autonomy in the real world often comes from failures to account for certain factors or from errors in the model. On the other hand, much of the information about the world may be completely irrelevant to the task that the robot tries to achieve. A challenge then is to infer a robot planning representation that is reasonable and useful for a given task. This inference process may also be combined with robot actions that explore the world and lead to better model estimates. The planner can aid in this exploration given its knowledge of the task and the potential solutions it considers.

For a robot to come up with a compact representation for planning without any prior experiences or human input is challenging, if not impossible. Exploring the role of human demonstrations for planning could help with this potential avenue of research. Can a planner utilize human demonstrations in building a compact planning representation for the task at hand? Can the planner figure out when to ask for demonstrations and then learn from them the "right" planning representation? In what form should these representations be given (for example, teleoperated, simulated, or kinesthetic demonstrations of the full task, or advice on what factors the planner should consider)?

Another important research direction is to explore the benefits of experiences. Can planning learn from prior planning and execution episodes what the "right" planning representation is for a given task? Past failures in execution may suggest the necessity for additional factors in planning, whereas the analysis of successful plans that do not exercise certain degrees of freedom in the world may allow the planner to construct a more compact representation for the given task.

Consistency, Predictability, and Understandability of Robot Behavior

The behavior of robots needs to be consistent, predictable, and understandable. This is especially the case in manufacturing where an operator needs to be able to anticipate what action the robot is going to perform next and how it will perform the action so that the operator can intervene when necessary. Predictability also simplifies the operator's task of coordinating multiple robots. The same holds in defense applications where a soldier needs to predict and understand the behavior of the robot in order to trust it and plan his/her own actions. In the domain of home assistance, predictability of robot behavior helps a human trust the robot and simplifies the coordination with the human's own actions.

Consequently, a challenge is to generate plans that are consistent (for example, similar for similar scenarios) and easy to understand by a human. These plans need to be generated and updated in real time despite the fact that many robotic systems (such as mobile manipulation platforms) have many degrees of freedom (DOFs). Computing feasible and optimal plans is already computationally challenging for high-DOF systems. Being able to compute, in real time, consistent plans for high-DOF systems is therefore a considerable research challenge. In addition to pure computational challenges, there is a question of the understandability of robot behavior and motion, that is, what behaviors and motions are understandable and how understandability can be maximized.

Human-Aware Planning

Autonomous robots working alongside humans face an additional set of unique challenges. The robot should behave in a way that is safe and consistent with the behavior of humans in the environment and that helps humans accomplish their tasks without becoming a nuisance. To achieve this, the robot needs to infer human intentions and goals and incorporate them into its plans so that the robot helps the humans without causing delay, danger, or confusion. However, human intentions are typically impossible to predict perfectly. Instead, based on the context and prior observations, intentions are typically inferred probabilistically. A challenge for planning is, therefore, to compute safe plans that account for the uncertainty in human intentions as well as utilize actions that disambiguate this uncertainty (for example, asking clarifying questions) when necessary and possible. Robots also need to be able to explain their behavior to humans on request. A challenge is therefore to create robot planning techniques with the ability to provide explanations.

The presence of humans presents not just challenges but opportunities for robots to improve their reliability. Robots can ask humans for help in accomplishing tasks that are hard to complete autonomously and when perception fails or is not sufficiently accurate. A challenge for planning is to reason about the chances of successfully accomplishing a task without human help and the possibility, cost, and utility of asking humans for help. Furthermore, humans can also be asked to provide demonstrations. Robot planning should reason about when demonstrations should be provided and how demonstrations can be used to infer what behaviors and motions are expected from the robot, what planning representation is best suited for planning, and what constraints need to be obeyed during task execution. This can be even more challenging if human inputs are only partial demonstrations or advice as opposed to full demonstrations of how a task can be accomplished.

Robotic Systems with Guarantees on Performance

Robots for many applications are becoming more and more complex, with higher degrees of freedom and/or massive arrays of sensors. As a consequence, the software architectures of robots are also becoming more and more complex, incorporating numerous distinct software modules. Given such complexity, it becomes difficult to assure that the behavior of the robot is going to be correct under different conditions, and the lack of such assurance jeopardizes the employment of autonomous robots in many domains. For example, human coworkers and robot operators expect reliable and repeatable behavior from the robots in domains such as defense, transportation, medicine, and manufacturing.

Consequently, the software modules of robots need to be designed in a way that the reliability, the repeatability, and the performance of the overall system can be analyzed. Since robot planning is responsible for decision making, this places a significant burden on the planner. That is, in addition to the requirement that the planner itself has guarantees on its performance and generates consistent solutions, we need to reason about its interaction with other components. More specifically, it brings up several challenges for the design of robot planning architectures. How should different levels of planning (such as task-level planning, motion-level planning, and low-level controls) be combined in a principled way? What properties does each of these modules need to satisfy in order to maintain guarantees on the performance of the overall system? How should planning interact with nonplanning modules (such as perception) in order to provide guarantees on performance?

Robot Planning That Utilizes the Availability of Massive Amounts of Data.

Much of the brittleness of current autonomous robots comes from the fact that they lack a deeper understanding of the world. It is much easier to plan motions for simple tasks (such as pick-and-place tasks) than to generate plans to accomplish more complex tasks (such as cooking). Geometric information about the world can be relatively easily perceived by a robot, but the semantics of perceived objects are much harder to derive. A robot often has a good understanding of how its own body moves, but knowledge about how other objects, especially articulated or deformable objects, can be manipulated is difficult and sometimes impossible to encode beforehand and is often impractical to try to estimate online. Finally, many tasks require a prior knowledge of "recipes" for how they can be achieved. These recipes provide an abstract and potentially partial specification for how to achieve a task. It is impractical to preprogram the recipes for all tasks the robot may encounter during its lifetime.

On the other hand, modern robots typically have access to the Internet and consequently massive amounts of data available on the web. Can this data be utilized to empower robots with a deeper understanding of the world and to improve their robustness? For example, can robot planning utilize partial "recipes" on the web for how tasks should be accomplished? When planning to manipulate a nonrigid object, can a planner collect data from the web about how this object can be manipulated and utilize the data to build an effective planning representation and guide the search for a plan?

Furthermore, given the network connectivity of robots and the importance of having vast amounts of knowledge for planning complex tasks, the knowledge and experiences gathered by one robot can and should be shared among other robots when possible. Sharing information among robots has the potential to accelerate their understanding of the world. With this in mind, the question is how to build a common shared database of knowledge and experiences for robots and what information should go into the database given the vast differences in modern robotic systems.

Open Source Planning Libraries

The development of the Robot Operating System (ROS)[6] has had an enormous effect on sharing research results between academic groups and transitioning robot technologies into the commercial world. ROS is now being used by numerous companies and nearly every university that does research in robotics. Part of this success can be attributed to the fact that many ROS components were built in joint efforts between researchers at companies, foundations, and academia. Equally important is the fact that ROS and its components are under an open source license that allows for the unrestricted use of the software.

While there are several planning libraries (such as OMPL [Sucan, Moll, and Kavraki 2012], SBPL,[7] and ROSPlan [Cashmore et al 2015]) available under ROS, it is important to develop more open source planning tools that are interoperable with commonly used robotic software infrastructures (such as ROS) and are available to the research community without any restrictions. While the development of these tools requires significant resources and efforts, especially to achieve a form that is usable in industry, they can dramatically help with making joint progress toward full robot autonomy and its commercialization. Government research agencies and industrial collaborators should recognize the importance of such efforts and support them.

Challenge Problems for Robot Planning

We next present a set of challenge problems, which are problems that, if studied, will likely move research in a direction that makes planning even more relevant for real-world robotics applications. Challenge problems can be created around the applications or robot capabilities described earlier. Each of the described challenge problems requires enabling multiple robot capabilities using planning.

Desirable properties of challenge problems include the following. The challenge problems should spell out possible evaluation scenarios and metrics. They could become progressively more difficult, for example, require longer and longer periods of autonomy. The challenge problems should be designed such that they are as robust as possible to overfitting, that is, solutions customized for the challenge problem should be generalizable to real-world scenarios. The challenge problems should be scoped well. They should currently be out of reach and thus result in clear advances of the state of the art yet have a low barrier to entry, that is, enable progress with small teams and a limited amount of resources. For example, researchers should be able to tackle them without necessarily having access to specialized equipment and without having perfectly working nonplanning robot capabilities such as control and perception. This can be achieved by starting with robot simulations or carefully crafted problems that minimize the need for certain capabilities. For example, if a challenge problem requires both perception and actuation capabilities, including a human in the loop could eliminate the requirement for one of these capabilities; for example, helping a blind person cook requires perception capabilities but no actuation capabilities. Larger challenge problems should span multiple robot capabilities or application scenarios and involve researchers from different disciplines, such as from artificial intelligence and robot control theory.

The impact of challenge problems can be increased by supporting common data sets, simulation environments, and hardware platforms and by requiring the participants to make open source software available to the research community.

Box and Bin Handling

Creating robots that can handle boxes, bins, and the small rigid and flexible items in them is a challenge problem that has implications for multiple applications, including warehousing, manufacturing, and home assistance. Specific tasks include opening and unpacking boxes, placing items into bins, finding items in bins, and packing items into boxes. The

Amazon Picking Challenge,[8] announced after the workshop on which this article is based, addresses a subset of the challenges above. The overall challenge of box and bin handling requires planners that advance the state of the art in multiple subareas of planning, including grasp planning, manipulation planning, motion planning, and task planning.

Warehousing for Manufacturing on Demand

Manufacturing on demand allows a company (especially a small business) to manufacture customized products in small batches as they are purchased. Warehousing includes the close coordination of multiple robots that navigate in tight spaces to transport objects between different locations in a warehouse. In the current state of the art, planners are typically provided with the start and goal locations for each robot. It is an open problem how effectively to integrate low-level path planning with high-level task planning, which is critical for effective automated manufacturing on demand. In this challenge problem, because products can be customized, it is necessary to determine sequences of goal locations for the robots that not only achieve the task-planning objective but consider the impact of the selected sequences on path planning (for example, to keep the resulting paths short and prevent congestion of the robots).

Fetching and Cleaning in Home Environments

Personal robots in homes, assisted living centers, and nursing homes must operate in human spaces, which are typically cluttered, unstructured, and include humans and pets. A challenge problem in this domain is to fetch items for a person with a disability. Another challenge problem is to clean up a cluttered room. Planning in this domain requires awareness of humans, which raises multiple challenges as discussed in the prior section. For example, the motions of robots need to feel natural to humans so that they are predictable and enable cooperation. This challenge problem can be extended to substantially more complex tasks. Possible extensions include building a robotic maid, butler, nurse, or cook. A home assistant robot, for example, could be required to perform tasks such as delivering daily medication, doing laundry, changing linens, cooking, serving food, and helping with personal hygiene, eating, and dressing. A subset of these tasks is currently covered by the RoboCup@Home[9] league, which defines specific scenarios for competitions.

Surgical Manipulation

Many surgical tasks require manipulating deformable tissues inside the human body. Planning motions for surgical robots that account for deformation could enable surgeons to perform safer and more efficient surgery. A representative challenge problem is retraction and exposure: the objective is for a laparoscopic surgical robot to grasp a flap or section of tissue and lift it to expose tissue underneath. This challenge problem requires robust manipulation of deformable objects as well as an appropriate level of perception of the objects in the scene in real time as they deform. This challenge problem can be made more realistic (and more difficult) by considering constraints on visibility of tissues and high levels of uncertainty in the motion of instruments and tissue. Furthermore, some surgical sites may only be safely accessed by maneuvering along curved trajectories through constrained cavities, which would require planning motions for snakelike or tentaclelike robots with many degrees of freedom to bend around anatomical obstacles.

Search and Rescue

Searching for and rescuing a human or animal in an unstructured environment raises numerous planning challenges for a robot. A representative challenge problem in search and rescue is for a mobile manipulator robot to navigate over rubble, search for an object, grasp the object, and then bring it to a new location. Many aspects of the tasks in search and rescue scenarios are included in the RoboCup Rescue[10] league competitions. This challenge problem can be extended to consider situations with large crowds of humans, which introduces large numbers of degrees of freedom that must be considered during planning. Other extensions include using ground, marine, and aerial vehicles and considering larger and more complex environments and tighter time constraints.

Conclusions

Although robots are increasingly being used in a variety of real-world applications, the deployment of advanced robot planning capabilities in real-world robots has thus far been limited. Progress will require the collaboration of planning researchers from robotics and artificial intelligence with researchers from neighboring disciplines, such as computer vision, haptics, natural language processing, and human-computer interfaces. We hope that this article will help guide researchers, inspire new research directions, and lead to new programs that stimulate research on robot planning with the goal of making robots with advanced planning capabilities ready for real-world deployment.

Acknowledgments

This article is based on the findings of a workshop sponsored by the United States National Science Foundation under award 1349355 that included participants from academia, government agencies, and industry (including Google, iRobot, Kiva Systems, and Rethink Robotics). This article has been shaped by all of the workshop participants, namely Ron Alterovitz, Nancy Amato, Ronald Arkin, Jennifer Barry, Michael Beetz, Kostas Bekris, Ronald Boenau, Sachin Chitta, Howie Choset, James Donlon, David Ferguson, Gary Gilbert, S. K. Gupta, Kris Hauser, Joe Hays, Geoffrey Hollinger, Wes Huang, Purush Iyer, Leslie Kaelbling, Subbarao Kambhampati, Behzad Kamgar-Parsi, Dai Hyun Kim, Sven Koenig, Hadas Kress-Gazit, Vijay Kumar, Maxim Likhachev, Kevin Lynch, Dinesh Manocha, Jeremy Marvel, Brett Piekarski, Fred Proctor, Craig Schlenoff, Don Sofge, Mike Stilman, Manuela Veloso, Richard Voyles, and Peter Wurman. We also thank the reviewers for their suggestions, which have been incorporated into this article as well. The opinions, findings, conclusions, and recommendations expressed in this article do not necessarily reflect the views of the National Science Foundation.

Notes

1. Robot Planning in the Real World: Research Challenges and Opportunities. A 2013 NSF-Sponsored Workshop cochaired by R. Alterovitz, S. Koenig, and M. Likhachev. Details are available at robotics.cs.unc.edu/PlanningWorkshop.

2. See Kiva Systems 2013 Press Release (www. kivasystems.com/amazon-press).

3. U.S. Department of Transportation Traffic Safety Facts 2010. Data is available at www-nrd.nhtsa.dot.gov/Pubs/811630.pdf.

4. Information on Amazon Prime Air is available at www.amazon.com/b?node=8037720011.

5. An official Google blog on the self-driving car that logs more miles on new wheels is available at googleblog.blogspot.com/2012/08/the-self-driving-car-logs-more-miles-on.html.

6. Available at www.ros.org.

7. The Search-Based Planning Library (SBPL), an open-source library of graph searches and their applications to robotics, is available at wiki.ros.org/sbpl.

8. Information about the Amazon Picking Challenge is available at amazonpickingchallenge.org.

9. Information about the RoboCup@Home League 2013 is available at www.robocupathome.org.

10. Information about the 2013 RoboCup Rescue League is available at www. robocuprescue.org.

References

Cashmore, M.; Fox, M.; Long, D.; Magazzeni, D.; Ridder, B.; Carrera, A.; Palomeras, N.; Hurtós, N.; and Carreras, M. 2015. ROSPlan: Planning in the Robot Operating System. In *Proceedings of the Twenty-Fifth International Conference on Automated Planning and Scheduling,* 333–341. Palo Alto, CA: AAAI Press.

Ghallab, M.; Nau, D.; and Traverso, P. 2004. *Automated Planning: Theory and Practice*. San Francisco: Morgan Kaufmann Publishers.

Kassner, E.; Reinhard S.; Fox-Grage, W.; Houser, A.; Accius, J.; Coleman, B.; and Milne, D. 2008. *A Balancing Act: State Long-Term Care Reform*. Technical Report, AARP Public Policy Institute. Washington, DC: American Association of Retired Persons.

LaValle, S. M. 2006. *Planning Algorithms*. Cambridge, UK: Cambridge University Press.

McDermott, D. 1992. Robot Planning. *AI Magazine* 13(2): 55–79.

Stevenson, A., and Lindberg, C. A., ed. 2010. *New Oxford American Dictionary.* New York: Oxford University Press.

Sucan, I. A.; Moll, M.; and Kavraki, L. E. 2012. The Open Motion Planning Library. *IEEE Robotics and Automation Magazine* 19(4): 72–82.

Ron Alterovitz is an associate professor in the Department of Computer Science at the University of North Carolina at Chapel Hill. He develops motion-planning algorithms for medical and assistive robots to enable robots to perform new, less invasive medical procedures and to assist people in their homes.

Sven Koenig is a professor in the Department of Computer Science at the University of Southern California. He develops techniques for decision making that enable agents and teams of agents to act intelligently in their environments and exhibit goal-directed behavior in real time.

Maxim Likhachev is a research assistant professor with the Robotics Institute and National Robotics Engineering Center (NREC) at Carnegie Mellon University. He develops techniques for high-dimensional planning in real time, planning that "learns how to plan" based on experience, automatic generation of compact representation of planning problems, and high-dimensional planning under uncertainty in real time.