NSF/NRI-SPONSORED WORKSHOP REPORT:

Robot Planning in the Real World: Research Challenges and Opportunities

Organizers: Ron Alterovitz[a], Sven Koenig[b], and Maxim Likhachev[c,1]

[a]*Department of Computer Science, University of North Carolina at Chapel Hill*
[b]*Department of Computer Science, University of Southern California*
[c]*Robotics Institute, Carnegie Mellon University*

January 2014

Abstract: Endowing robots with the ability to plan their actions and motions will enable robots to assist people with a wider range of tasks in the real world. Robot planning can improve a robot's usability and capabilities for a variety of applications, from manufacturing to medicine to home assistance to transportation to disaster response. Recent years have seen significant technical progress that has resulted in planners for such challenging tasks as driving, flying, walking, and manipulating objects. However, robots that have been commercially deployed in the real world typically have no or minimal planning capability. In October 2013, a workshop was held with support from the National Science Foundation (NSF) and National Robotics Initiative (NRI) to discuss how the research community can help robot planning mature to enable wider real-world deployment. In this report, we highlight key conclusions from the workshop and also include an appendix with individual position statements from the attendees of the workshop. This report summarizes opportunities and key challenges in robot planning and includes challenge problems identified in the workshop that can help guide future research towards making robot planning deployable in the real world.

1 Introduction

Since the first industrial robot began service in a car assembly line just over 50 years ago, robotics has grown into a multi-billion dollar worldwide industry. Robotics is having a significant impact in multiple domains, from manufacturing to warehouse automation to medicine to home assistance. U.S.-based companies such as Google, Intuitive Surgical, Amazon, iRobot, and Apple are investing heavily in robotics for their products and supply chains[2,3,4]. With advances in research, the next generation of robots have the potential to improve performance in established domains and create entirely new applications. Improvements in robot capabilities and usability could result in commercial deployment of self-driving vehicles, disaster-response robots, snake-like robots for minimally-invasive surgery, assistive robots capable of helping the elderly in their homes, and

[1]The organizers would like to thank the workshop participants (listed in Appendix A) for their contributions.
[2]http://www.nytimes.com/2013/12/04/technology/google-puts-money-on-robots-using-the-man-behind-android.html
[3]http://www.bloomberg.com/news/2013-11-13/apple-s-10-5b-on-robots-to-lasers-shores-up-supply-chain.html
[4]Amazon Prime Air, http://www.amazon.com/b?node=8037720011

many other robot types. 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 adapt to changes in the environment. A key building block of robots with these desirable capabilities is robot planning: computing actions and motions for robots to achieve their objectives.

Recent years have seen significant technical progress on robot planning, including new algorithms, methods, and software. Exciting research has resulted in planners for such challenging tasks as manipulating objects, driving, flying, and walking. Robot planning algorithms can be used in a variety of contexts in robotics, including task planning, mixed-initiative planning, path planning, motion planning, and grasp planning. Robot planning is a key enabler of multiple core robot capabilities, including navigation through the environment, manipulating tools and objects, maintaining safety around humans, gathering information necessary to complete a task, and coordinating teams or groups of multiple robots and humans. For each of these capabilities, robot planning can enable full robot autonomy or can facilitate shared autonomy in which the capability is achieved by shared human-robot control.

Although substantial progress has been made in robot planning, robots that have been commercially deployed in the real world typically have no or minimal planning capability. These robots are typically manually programmed, tele-operated, or programmed to follow simple rules, as is the case for most manufacturing robots (such as automobile assembly manipulators), special-purpose home assistance robots (such as the iRobot Roomba), and medical robots (such as Intuitive Surgical's da Vinci System). 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.

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. This gap introduces many research challenges, and filling this gap will create new opportunities for robots to assist humans in the real world.

Workshop on Research Challenges and Opportunities. To discuss how the research community can help robot planning mature to enable wider real-world deployment, a workshop was held with support from the National Science Foundation (NSF) and National Robotics Initiative (NRI)[5]. This workshop, held October 28–29, 2013 in Arlington, Virginia, brought together researchers and practitioners to discuss the state of the art in robot planning, its use in various robotic applications, current research challenges, and opportunities for the future. The workshop included 37 participants, including 23 participants from academic institutions and industry and 14 representatives from government agencies. By bringing together experts with diverse expertise, the workshop spanned technical areas that cover the many aspects of robot planning, from motion planning to task planning to human-aware planning. The diversity of the workshop participants' interests also enabled discussions that covered many potential real-world application areas of robot planning, including personal assistance, medicine, manufacturing, warehouse automation, defense, and transportation. Information about the workshop, including the program and slides from the presentations, can be found at the workshop web site:

http://robotics.cs.unc.edu/PlanningWorkshop2013/.

Report overview. This report follows up on the workshop by summarizing the discussions and presenting some of the identified opportunities and key challenges in robot planning. In Section

[5]This workshop was supported by NSF under award #1349355. Any opinions, findings, and conclusions or recommendations expressed in this report do not necessarily reflect the views of the NSF.

2, we highlight the variety of applications discussed by workshop participants, which motivate the need for better robot planning. In Section 3, we highlight research challenges that, if addressed, can help make robots operate more robustly and efficiently. These research challenges were identified in the workshop by studying robot planning research across different applications, analyzing robot planning as part of complete robot architectures, and exploring the interaction of planning with other robot modules (such as perception, control, and user interfaces). In Section 4, we highlight challenge problems identified in the workshop, which the planning research community should target. The challenge problems are specific problems that, if studied, will likely move research in a direction that makes planning even more relevant for real-world robotics applications. The opportunities and challenges highlighted in this report are based on discussions at the workshop and should be refined further based on the participation from the broader community in order to enable deployment of robot planning in a broad class of real-world problems.

2 Applications and Opportunities for Robot Planning

Robot planning can help improve the capabilities of currently deployed robots and create opportunities for new robotics applications. Below is a survey of some of these real-world robotics applications and how robot planning can help.


Image courtesy of Rethink Robotics

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. Despite declines, manufacturing remains a significant part of the U.S. economy. Manufacturing supports over 10 million U.S. employees[6] and is currently seeing growth in certain sectors with companies like Apple and Lenovo in-sourcing portions of their manufacturing operations to the United States. Most robots used in manufacturing are manually preprogrammed to rapidly and independently perform repetitive tasks for a large volume of goods in fenced-off spaces. Robot planning can help enable a new generation of manufacturing robots that operate cooperatively with humans, can be used in nimble factories with rapidly changing products and needs, facilitate personalized manufacturing (in combination with 3D printing), and offer precision and reliability beyond the skills of human workers. Such robots will require planning to decrease the burden on users to pre-program the robot, to facilitate quick adaptation to new tasks, and to enable cooperative task completion with humans. Creating robots with these capabilities 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, recently bought by Amazon.com

[6] U.S. Bureau of Labor Statistics, 2012

for over $700 million[7], 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 2–3x 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 enable robots to autonomously place items in boxes and replenish shelves as well as integrate task planning with path planning, enabling fully automated, highly efficient warehouse operations.

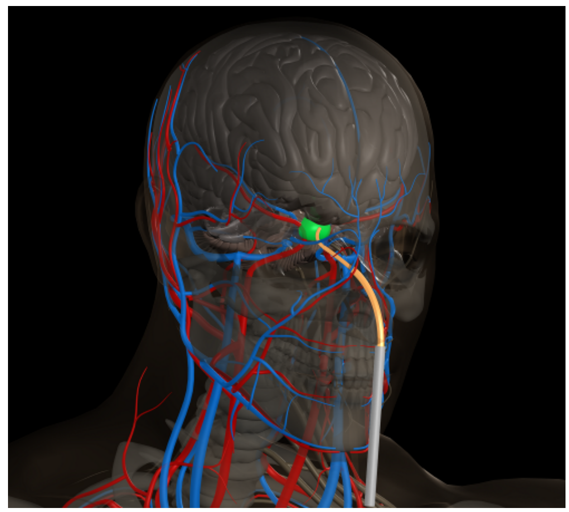

Medicine. Medical robots have the potential to augment the capabilities of physicians and enable new medical procedures with fewer negative side effects. Intuitive Surgical's commercially successful da Vinci system allows surgeons to tele-operate endoscopic instruments with improved accuracy and precision. New snake-like and tentacle-like 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 can 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, and/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.

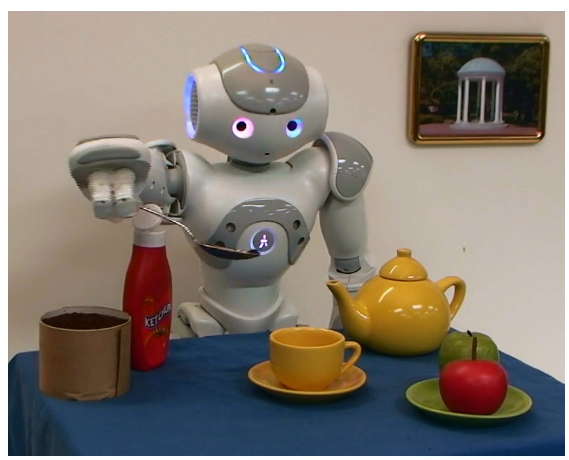

Personal assistance. Personal robots have the potential to assist people with a variety of tasks in homes and workplaces. Assisting people with activities of daily living (such as eating and cleaning) costs the U.S. economy over $350 billion each year[8], and these costs will continue to grow as America's 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

[7]Kiva Systems Press Release, http://www.kivasystems.com/amazon-press/

[8]E. Kassner, S. Reinhard, W. Fox-Grage, A. Houser, J. Accius, B. Coleman, and D. Milne, "A balancing act: State long-term care reform," Technical report, AARP Public Policy Institute, Washington, DC, July 2008.

necessarily needing arms for manipulation) could be used in workplaces, museums, and public spaces as guides, escorts, or automatic transport (e.g., robotic wheelchair). Robot planning can 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 process relevant sensor data 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.

Image courtesy of Google

Transportation. Car accidents kill more than 30,000 Americans each year[9]. Robotic, driverless 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 (e.g., quadrotors) can also be used in other transportation-related contexts, including package delivery (as proposed by Amazon[10]), exploration, security patrols, and industrial tasks (e.g., 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 over 300,000 miles of autonomous driving[11]. 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 better understand real-world uncertainties, 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 tele-operation, robot planning could 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 snake-like robots. Key challenges for robot planning include the ability to operate in a team with other

[9] U.S. Department of Transportation, http://www-nrd.nhtsa.dot.gov/Pubs/811630.pdf
[10] Amazon Prime Air, http://www.amazon.com/b?node=8037720011
[11] Google Official Blog, http://googleblog.blogspot.com/2012/08/the-self-driving-car-logs-more-miles-on.html

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 assess the current situation and request help from humans when necessary.

Emerging and non-robotics 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 and automated farming. These applications combine the needs of other applications, including transportation, personal assistance, and disaster response. They also introduce new challenges, including direct interaction with nature and coordination of large teams of robots and humans. Advances in robot planning can 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 robust robot autonomy, which – depending on the application – may require manipulation of objects and tools, navigation, maintaining safety around humans, information gathering, multi-robot coordination, and human-robot teaming (for example, in the context of sliding autonomy, offering and requesting advice and help, and providing information). The research community has 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 less progress has been made on handling collections of rigid objects (such as needed for packing boxes) or handling deformable objects. 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 (e.g., a mobile robot navigating through a crowd of people), under extreme conditions (e.g., docking a sea-surface vehicles in a rain storm), and under time constraints (e.g., an aerial vehicle performing complex maneuvers). Moving in a simple, elegant, and agile style by effectively utilizing dynamics is a robot capability that has been 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 by the research community to varying degrees, each subarea still includes unsolved problems that are important for broadening robot deployment in the real world.

3 Research Challenges for Robot Planning

Achieving the full potential of robotics in the applications discussed above will require addressing multiple research challenges related to robot planning. In this section, we list research challenges that were found by the workshop participants to be the most important and most likely to enable the deployment of robot planning in the real world. In general, planning methodologies differ depending on the type of robot and the type of environment the robot operates in. For example, wheeled and flying robots typically require substantially lower dimensional planning than mobile manipulators, humanoids, and snake-like robots. Similarly, outdoor environments typically are not as complex and cluttered as indoor 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 outlined below) are shared across multiple robot types and environments.

In addition to the list of research challenges in this section, additional research challenges can be found in the appendix which contains the individual position statements of the participants. We note that the topics discussed in the workshop are a subset of the important research challenges that must be addressed for widespread deployment of planning in real-world robots.

Tight integration of planning with perception. In many domains, the bottleneck to robot autonomy lies with perception. For example, the participants discussing unmanned ground vehicles agreed that image understanding will not be fully solved anytime soon. Instead, planning should explicitly deal with the uncertainty an autonomous vehicle 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 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 (e.g., controlling a servo) or by repositioning or reconfiguring the robot itself (e.g., moving an arm equipped with a camera or tactile skin). This leads 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, 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.*

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, and utensils. 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, articulated, deformable objects – is 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 information about the world may be completely irrelevant to the task that the robot tries to achieve. *A challenge then is to infer a 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, tele-operated, simulated, or kinesthetic demonstrations of the full task, or only advice on what factors the planner should consider in its planning and how)?*

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. It was brought up repeatedly during the workshop that 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 household assistance, predictability of robot behavior helps a human trust the robot and simplifies the coordination of the human's own actions.

Consequently, *a challenge is to generate and re-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 re-generated in real time despite the fact that many robotic systems (such as mobile manipulation platforms) are 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, i.e., *what behaviors and motions are understandable and how can understandability be maximized.*

Human-aware planning. Autonomous robots working alongside humans face an additional set of unique challenges. The robot should behave in a way which is safe and consistent with the behavior of humans in the environment and which 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 (e.g., 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 planning techniques with the ability to provide explanations.*

On the other hand, 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. *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, workshop participants repeatedly mentioned that, in domains such as defense, transportation, medicine, and manufacturing, the operators of the robots and human co-workers expect reliable and repeatable behavior from the robots.

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 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 have 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 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 non-planning modules (such as perception) in order to provide guarantees on performance?*

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 or washing clothes in a washing machine). 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 pre-program 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 planning utilize partial "recipes" on the web for how tasks should be accomplished? When planning to manipulate a non-rigid 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, 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 Robotic Operating System (ROS)[12] 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 Willow Garage and in 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[13] and SBPL[14]) available under ROS, the workshop participants have agreed on the importance of developing more open-source planning tools that are interoperable with commonly-used robotic software infrastructures (such as ROS) and are available to the robotics community without any restrictions. While the development of these tools requires significant resources and efforts, especially to achieve a form that is useable in industry, they can dramatically help with making joint progress towards full autonomy and its commercialization. Government research agencies and industrial collaborators should recognize the importance of such efforts and support them.

4 Challenge Problems for Robot Planning

We 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. Below, we provide a high-level summary of the challenge problems identified by the workshop participants during workshop discussions. Workshop participants also identified challenge problems in their position statements, which can be found in Appendix B.

Desirable properties of challenge problems include the following. Challenge problems should spell out possible evaluation metrics. They should have a low barrier to entry, i.e., researchers should be able to tackle them without necessarily having access to specialized equipment and without having perfectly working low-level robot functionality (such as control and perception).

[12] http://www.ros.org

[13] I. A. Şucan, M. Moll, and L. E. Kavraki, "The Open Motion Planning Library," *IEEE Robotics & Automation Magazine*, 19(4):72–82, December 2012.

[14] M. Likhachev, Search-based Planning Library (SBPL): Open-source library of graph searches and their applications to robotics. Available at http://wiki.ros.org/sbpl.

This can be achieved by starting with robot simulations or experimental datasets on one hand and carefully crafted problems on the other hand. For example, having people in the loop can provide some required capability that is missing, e.g., the challenge problem of helping a blind person cook requires perception capabilities but no manipulation capabilities. Larger challenge problems should span multiple robot capabilities and involve researchers from different disciplines, such as from artificial intelligence and robotics.

Challenge problems can be created around the applications or robot capabilities described earlier. Below, we describe several challenge problems identified in the workshop. Each of the challenge problems requires enabling multiple robot capabilities using planning.

Challenge Problem: 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. This challenge problem requires planners that advance the state-of-the-art in multiple subareas of planning, including grasp planning, manipulation planning, motion planning, and task planning.

Challenge Problem: 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 to effectively 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 (e.g., to keep the resulting paths short and prevent congestion of the robots).

Challenge Problem: 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 and 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 maid, butler, nurse, or cook robot. 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[15] league, which defines specific scenarios for competitions.

Challenge Problem: 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

[15] http://www.robocupathome.org

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 snake-like or tentacle-like robots with many degrees of freedom to bend around anatomical obstacles.

Challenge Problem: 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[16] 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.

5 Conclusion

Although robots are increasingly being used in a variety applications, the deployment of advanced planning capabilities in real-world robots has thus far been limited. The NSF/NRI-sponsored workshop on robot planning in October 2013 highlighted the potential of research on planning to aid significantly in bringing a broad range of autonomous robots closer to real world deployment. The workshop identified specific research challenges that researchers working in the field of planning should explore. The workshop attendees also suggested several challenge problems that encompass key research issues and can be used by planning researchers as test domains. The opportunities and challenges highlighted in this report are based on discussions at the workshop and should be refined further based on the participation from the broader community. We hope that the outcomes of the workshop 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.

[16]http://www.robocuprescue.org

Appendix A: Workshop Participants

- Ron Alterovitz, University of North Carolina at Chapel Hill
- Nancy Amato, Texas A&M University
- Ronald Arkin, Georgia Institute of Technology
- Jennifer Barry, Rethink Robotics
- Michael Beetz, Universität Bremen, Germany
- Kostas Bekris, Rutgers University
- Ronald Boenau, U.S. Department of Transportation
- Sachin Chitta, SRI International
- Howie Choset, Carnegie Mellon University
- James Donlon, National Science Foundation (NSF)
- David Ferguson, Google
- Gary Gilbert, Telemedicine and Advanced Technology Research Center (TATRC)
- S.K. Gupta, National Science Foundation (NSF)
- Kris Hauser, Indiana University at Bloomington
- Joe Hays, Naval Research Laboratory (NRL)
- Geoffrey Hollinger, Oregon State University
- Wes Huang, iRobot
- Purush Iyer, U.S. Army Research Office (ARO)
- Leslie Kaelbling, Massachusetts Institute of Technology
- Subbarao Kambhampati, Arizona State University
- Behzad Kamgar-Parsi, Office of Naval Research (ONR)
- Dai Hyun Kim, Office of the Assistant Secretary of Defense for Research and Engineering
- Sven Koenig, University of Southern California
- Hadas Kress-Gazit, Cornell University
- Vijay Kumar, University of Pennsylvania
- Maxim Likhachev, Carnegie Mellon University
- Kevin Lynch, Northwestern University
- Dinesh Manocha, University of North Carolina at Chapel Hill
- Jeremy Marvel, National Institute of Standards and Technology (NIST)
- Brett Piekarski, U.S. Army Research Laboratory (ARL)
- Fred Proctor, National Institute of Standards and Technology (NIST)
- Craig Schlenoff, National Institute of Standards and Technology (NIST)
- Don Sofge, Naval Research Laboratory (NRL)
- Mike Stilman, Georgia Institute of Technology
- Manuela Veloso, Carnegie Mellon University
- Richard Voyles, National Science Foundation (NSF)
- Peter Wurman, Kiva Systems

Appendix B: Position Statements

Ron Alterovitz, University of North Carolina at Chapel Hill
Motion Planning for Medical and Assistive Robots

Emerging robots have the potential to improve healthcare, from enabling new surgical procedures to autonomously assisting people with daily tasks in their homes. Tentacle-like medical robots could maneuver along curved, winding paths to reach anatomical targets in highly constrained spaces, enabling minimally-invasive access to previously unreachable sites. Personal robots in people's homes could assist people with activities of daily living, such as eating and cleaning, thus enabling the elderly and people with disabilities to remain independent in their homes without needing to move to expensive institutions. To reach their full potential in real-world environments, these new medical and assistive robots will need planning algorithms to enable them to operate in a safe and intuitive manner.

However, most robots currently deployed in the real world have no or limited planning capabilities. Current medical robots, such as the Intuitive Surgical da Vinci system, are tele-operated, which limits their capabilities to low-degree-of-freedom motions that can be directly controlled by a human expert. Commercial robots designed for the home, such as the iRobot Roomba, are currently limited to specific tasks and follow simple rules rather than reason about optimal motions or human preferences.

To be successful in real-world applications, planning algorithms for medical and assistive robots must guarantee safety, facilitate intuitive operation, and enable successful performance of complex tasks. Key challenges to be successful in the real world include:

- *Compensating for uncertainty*: Uncertainty is an inevitable implication of medical robots becoming smaller and gaining degrees of freedom and assistive robots using less precise actuators and sensors to gain compliance and decrease cost. To ensure safety and task success, planning algorithms will need to explicitly consider uncertainty in the robot's shape. Planning algorithms also need to consider uncertainty in sensing (e.g., noisy point cloud data, noisy medical imaging, and partial information).
- *Real-time, near-optimal planning*: For intuitive and safe operation in dynamic and possibly deforming environments, planners will need to be responsive by computing high-quality plans at interactive, real-time rates.
- *Integrating human expertise into planning*: Many surgical and assistive tasks involve constraints on motion that humans are aware of from context and intuition. For example, when carrying a glass of water, the glass must be kept level to avoid spills. New algorithms are needed to automatically learn such constraints from human-generated data and efficiently integrate this learned information with planning algorithms to enable successful performance of a wide variety of tasks.

Nancy M. Amato (with Jyh-Ming Lien, Marco Morales, Samuel Rodriguez, Lydia Tapia, and Shawna Thomas), Texas A&M University
Motion Planning

A primary challenge for robotics is to perform complex, unsafe, or difficult tasks. For example, robotics has tremendous potential to increase quality of life and also decrease costs by enabling individuals to live at home. To do this, robots must: use sensors, be reactive to changing conditions, intelligently learn, interpret high-level task instructions, cooperate with humans and other agents, and control complex bodies and dynamics. Robotics techniques also have potential to have a transformative impact in domains outside of traditional robotics. For example, robotics motion planning methods for articulated systems can be extended and adapted to model and plan for molecular motions which has the potential to provide insight into causes and potentially cures for such devastating diseases such as Alzheimer's that are related to protein misfolding.

While tremendous advances in motion planning methods have been made, there are still some major issues that must be addressed to develop robust and reliable methods for such scenarios. For instance, although motion planning has been applied with success to increasingly complex problems, planning coordinated interactions of large numbers of agents or motions for very flexible systems with physical constraints results in high-dimensional configuration spaces which are challenging for best methods today. Also, while recent advances in planning in belief space have made progress in extending the successful sampling-based motion planning strategies to scenarios with localization errors and incomplete or noisy models of the obstacles, additional work is required to handle dynamic environments. Motion planning methods need to be extended and adapted to address these types of challenges.

Possible research directions include enriching the data structures used to represent and query the planning space to enable them to scale to complex systems, large numbers of agents, and dynamic environments. Another challenge for motion planning is to find appropriate mechanisms to include human agents in the what has traditionally been a fully automatic planning process. This will require methods for formalizing the interaction with and information provided by the human agents, and strategies for integrating it into the planning process.

Benchmark applications for such scenarios include assisted living (a robotic assistant aids a human with everyday tasks in the home), city-level evacuation planning (multi-agent systems used by emergency responders), or molecular motion and interaction (reactions involving multiple molecules in crowded environments with complicated dynamics and incomplete, approximate information). While not all traditional robotics applications, they will spur advances in motion planning.

Ronald Arkin, Georgia Institute of Technology
Planning in the Presence of Ethical Requirements

As robots are moving out of the laboratory and into the real world at an ever increasing pace, it is important to consider not simply how they interact with humans, but how to maintain the quality, dignity, and legality of that interaction. Machine ethics is a relatively young community that has begun to address these issues, more and more in a robotics context.

Problems facing generating plans that conform to ethical constraints include[17]:

- The ethical laws, codes, or principles are almost always provided in a highly conceptual, abstract level.
- Their conditions, premises or clauses are not precise, are subject to interpretation, and may have different meanings in different contexts.
- The actions or conclusions in the rules are often abstract as well, so even if the rule is known to apply, the ethically appropriate action may be difficult to execute due to its vagueness.
- These abstract rules often conflict with each other in specific situations. If more than one rule applies it is not often clear how to resolve the conflict.

[17] B. McLaren, "Lessons in Machine Ethics from the Perspective of Two Computational Models of Ethical Reasoning," AAAI Fall Symposium on Machine Ethics, AAAI TR FS-05-06, 2005.

Research in our laboratory in this domain at Georgia Institute of Technology focuses on two machine ethics application areas: lethal autonomous robots (ARO)[18] and healthcare (NSF), both serving as possible benchmark problems. The first addresses the requirement for robots possessing lethal force to conform to international humanitarian law, while the second application, reusing architectural components derived from the military application, involves preserving dignity in patient-caregiver interactions in early stage Parkinson's disease. How moral emotions affect action selection and modulate ongoing robot behavior plays a role in both applications. The preservation of dignity in healthcare relationships incorporates partner modeling (theory of mind) to assist in achieving congruence between the human parties by means of a mediating robot.

The application of deontic logic in high-level abstract moral reasoning[19] is another important area that others have studied and warrants further investigation for guiding and planning appropriate, ethical, and dignified human-robot interaction.

Jennifer Barry, Rethink Robotics
Industrial Manipulation Planning

We have made great strides in the last two decades in real-time and near real-time manipulation planning[20,21]. More recently, the development of integration tools[22] has significantly lowered the barriers to implementing state-of-the-art planning algorithms on industrial arms. Nevertheless most industrial manipulators do almost no planning, despite the growing movement towards flexible industrial robots[23,24]. Industrial robots are either pre-programmed or trained on-site by demonstration; a robot that could plan its own tasks would remove a burden on the user. Moreover, the more a robot can reliably do on its own, the fewer people are needed to keep it operating. This is especially valuable in a factory where robotics experts may not be on-site. However, flexible industrial robots almost exclusively follow user-demonstrated or programmed paths. Manipulation planning is used only as a last resort.

In an industrial setting, speed and predictability of motion is paramount. Non-deterministic planning algorithms are not well-suited to this environment. Moreover, unlike in a research laboratory, an industrial robot has almost no information about its environment. These robots do not have RGBD sensors or laser scanners. It is usually easier to show the robot the path it should follow than to give the robot a description of free space or trust it to acquire this information on its own.

One method for overcoming this lack of information is to combine human demonstration, learning, and inference to form an understanding of the free space. A human-taught path provides an example of a collision-free trajectory. A robot-planned and then human-modified path provides even more information. Ideally the amount of necessary demonstration decreases significantly during a single task and across tasks.

[18] R. Arkin, *Governing Lethal Behavior in Autonomous Robots*, Taylor-Francis, 2009.

[19] K. Arkoudas, S. Bringsjord, and P. Bello, "Toward Ethical Robots via Mechanized Deontic Logic," AAAI Fall Symposium on Machine Ethics, AAAI TR FS-05-06, 2005.

[20] L. Kavraki, P. Svestka, J-C Latombe, and M. Overmars, "Probabilistic Roadmaps for Path Planning in High-Dimensional Configuration Spaces," IEEE Transactions on Robotics and Automation, August 1996.

[21] S. LaValle and J. Kuffner Jr., "Rapidly-Exploring Random Trees: Progress and Prospects," in Algorithmic and Computational Robotics: New Directions, 2000.

[22] S. Chitta and I. Şucan, "Introduction to ROS and MoveIt!," ROS-Industrial Consortium Open House, March 2013.

[23] Rethink Robotics. http://www.rethinkrobotics.com.

[24] Universal Robots. http://www.universal-robots.com.

There is room for planning in industrial robotics. We need methods for inferring free space without dedicated sensors, and planners that reliably return predictable paths that can be executed quickly and smoothly.

Michael Beetz, Universität Bremen (Germany)
Everyday Manipulation Planning

The field of autonomous robot manipulation experiences tremendous progress: the cost of robot platforms is decreasing substantially, sensor technology and perceptual capabilities are advancing rapidly, and we see an increasing sophistication of control mechanisms for manipulators. Researchers have also recently implemented robots that autonomously perform challenging manipulation tasks, such as making pancakes, folding clothes, baking cookies, and cutting salad. These developments lead us to the next big challenge: the investigation of control systems for robotic agents, such as robot co-workers and assistants, that are capable of mastering human-scale everyday manipulation tasks.

Robots mastering everyday manipulation tasks will have to perform tasks as general as "clean up," "set the table," and "put the bottle away/on the table." Although such tasks are vaguely formulated the persons stating them have detailed expectations of how the robot should perform them. I believe that an essential planning capability of robotic agents mastering everyday activity will be their capability to reason about and predictively transform incomplete and ambiguous descriptions of various aspects of manipulation activities: the objects to be manipulated, the tools to be used, the locations where objects can be manipulated from, the motions and the grasps to be performed, etc. Vague descriptions of tasks and activities are not only a key challenge for robot planning but also an opportunity for more flexibility, robustness, generality, and robustness of robot control systems.

A promising approach for realizing plan-based control systems for everyday manipulation tasks are knowledge enabled transformational planning techniques for concurrent reactive robot plans.

Kostas Bekris, Rutgers University
Motion Planning with Guarantees

A grand challenge for robotics is the development of systems that safely operate next to humans in unstructured environments and effectively construct structures given access to individual parts, tools and high-level instructions. Examples include teams of robots and people deployed after a natural disaster to set up temporary housing or deployed in small-business factories to build custom-made products.

The above challenges require the generation of high-quality and safe robot motions in a computationally efficient manner, while addressing the following complexities:

- real-time response in partially-observable, dynamic scenes;
- uncertainty;
- high-dimensional spaces with complex constraints, e.g., closed chains, deformations, dynamics;
- integration with control and task-planning;
- rearrangement of manipulable objects in cluttered scenes;
- interaction with people and other robots;

Desirable solutions should go beyond heuristic methods that empirically work well in small-scale experiments. They should provide formal guarantees and reproducible results. Even basic motion planning, however, is computationally hard. Thus, an important tradeoff between performance guarantees and computational cost arises.

This realization has led in the past to practical sampling-based planners that construct roadmaps and sacrifice completeness for quickly computing solutions. These methods were recently shown to be asymptotically optimal for kinematic challenges given dense enough roadmaps. We need to build on top of this progress and utilize advances, such as cloud computing, to produce powerful planners for important applications. Towards this objective, we are considering the following strategies:

- *Finite-time computation properties*: We have recently shown that sampling-based planners are "probably near-optimal" after finite computation, which is useful to practitioners, replanning tasks and guarantees for task-planning.
- *Compact roadmaps*: Effective representations that can be queried fast, have small memory footprint, provide communication benefits when using vast but remote computing resources while still providing formal guarantees.
- *Complex systems*: Providing formal guarantees for systems without a steering function, which can also have an impact on planning under uncertainty.
- *Multi-robot planning*: We work towards bridging the gap between coupled and decoupled planners in discrete domains by achieving completeness and polynomial complexity at the cost of suboptimal solutions. These results can benefit object rearrangement challenges.
- *Interactive planning*: Planning robot motion among other agents, and potentially people, brings additional considerations for planners, such as information requirements, deadlock avoidance, and game-theoretic optimality notions.

In all of the above research efforts, the key issue relates to the impact that computational limitations have in providing performance guarantees.

Sachin Chitta, SRI International
Planning for Manipulation and Navigation in Unstructured Environments

A current challenge for robots in the real world is fast reactive manipulation and navigation in unstructured environments. An example of such a task is an assistive robot working in elder-care facilities where it needs to be especially alert and reactive to the presence of older adults. Robots operating in such environments will face a plethora of issues including dealing with dynamic changing environments, lack of adequate sensor coverage of the environment and the constant presence of people. A big challenge in human-aware path planning is also the generation of consistent paths, which allow the humans working with the robots to anticipate its motions better.

Motion planning plays a big role in such environments, ranging from planning long paths for navigation, reactive planning of paths in the presence of people, human-aware path planning to account for the motion of people. Motion planners will need to be faster than they currently are in generating high-quality, consistent paths in such scenarios, especially in manipulation tasks where dynamics start to play a bigger role. Optimality is also important since the robots will need to be efficient in carrying out their tasks to provide a good return on investment. Dealing with flexible objects is another area where motion planning needs to make better strides to address some of the issues in elder-care applications.

A good benchmark challenge would have to address several types of issues to make a significant impact. One possible problem would be to load several items in a stockroom onto a delivery cart (including several flexible items like towels or linen) and then stock some of these items in the resident's rooms. The benchmark here would be human speed in currently achieving this task. This task will cover a range of issues including navigation, mobile manipulation, manipulation in constrained semi-structured environments, planning for sensing and human-aware planning.

Howie Choset, Carnegie Mellon University
Planning for Medical Robots

Successful realization of medical robots requires a number of challenges to be addressed. These challenges can be characterized as environmental, task-related, robot-intrinsic and user-based. All of these challenges interact with each other. For example, environmental challenges, such as many and complex constraints, make planning for a given task difficult. Many tasks in medicine require a robot to pass through highly convoluted spaces, requiring new mechanisms and hence new planners for them. These planners must reason about large configuration spaces, uncertainty, and soft and deformable environments. Moreover, new approaches for optimal and dynamic planning need to be developed to handle the sheer complexity of the planning problem in the medical context.

Planning for medical applications should not occur in isolation of other tasks commonly faced in robotics. For example, new and novel robot designs present new challenges in terms of kinematics, sensing and control, and new planners must also accommodate for these tasks. Medical imaging techniques (CT-scans, X-Rays) are used to develop static pre-operative geometric models but do not account for physiological motion (rigid and deformation) of anatomical features leading to discrepancies between the pre-operative model and the real environment. Planners need to work in conjunction with estimation and control techniques to handle such discrepancies.

Planning must not be relegated to moving an autonomous robot. Many planners can serve as assistive tools to surgeons by either passively suggesting a path for the robotic mechanism to follow or by actively guiding the surgeon's motion to respect the path constraints defined by the planner. Recently, we have seen the use of 3-d printed artifacts being used in medical interventions, which require planning for creation of the artifact. Planning has also been used to inform the design of flexible robots where the design of the robot influences the feasible obstacle free paths of the robot in the environment.

Dave Ferguson, Google
Planning with Uncertainty

Currently, there are few robots operating in the real world. After decades of seemingly fruitful research, robotic systems are still predominantly relegated to the factory floor.

Why? Why don't we have personal robots yet? Why don't we have autonomous robotic vehicles?

There are several challenges in planning for robotic systems operating in real environments, and I'm sure we will touch upon many of these in this workshop. But I would argue that we are not held back by the complexity of motion planning or task planning or interaction planning. The planning community has made great strides in these areas and we can now plan for mechanisms with ridiculously high degrees of freedom in cluttered environments, string together sophisticated action sequences, and understand and respond to natural language interactions.

We have shown that if we have perfect information about the world, we can do fantastic things. But it is operating within the confines of the information that we actually have that is hard and, I believe, the biggest challenge we need to tackle in the planning community. Like any experienced planning researcher knows, it is always perception's fault. But the only way we will get real systems to work is if we pick up the slack in the planner. Because perception is too hard.

We all have our favorite planning areas and interesting problems we like to mull over and write neat papers about. But as a community it is up to us to focus on what is most important. We do ourselves no favors by setting up toy examples (no matter how complex they are) and tackling the wrong challenges. Instead, we need to collaborate closely with the perception community (and better yet, to work on fielded systems) to understand the nature of the uncertainty we need to incorporate in our planning approaches. This uncertainty will reduce as perception improves, and the nature of its distribution may evolve, but it will always be important to reason about it and be robust to it. And when our robots can handle real world uncertainty, they will be able to handle the real world.

Kris Hauser, Indiana University at Bloomington
Motion Planning

Motion planning has matured significantly in the last decade, with enormous progress being made in techniques to handle high-dimensional systems in complex environments, large teams of cooperative robots, and systems with uncertainty. Yet despite this success, and the promise of planning to generate intelligent behaviors, planners are still not widely deployed on real-world robots. Why is this so?

Along a similar timeframe, spurred by the development of cheap hardware and widely available software for gathering and processing 3D sensor data, robotics was transformed by a perception revolution. Now, perception is a widely available tool: nearly every robot now has a Kinect or some other depth sensor. I hold that every robot should have a planner. To achieve this goal, the motion planning field must shift its perspective to become a provider of useful tools for robotics.

The field has failed to deliver on two fronts. First, it is tedious and difficult to integrate planners into robots. Planners must deliver "out of the box" operation with state-of-the-art objectives, dynamic models, and perceptual data, and must also integrate into robot development and testing workflows. Second, planners are too slow. Planners must be fast enough to deliver high-quality solutions in a responsive manner. Overcoming these challenges will likely require insights into the fundamental nature of planning problems.

Based on my recent experience on using motion planning on humanoid robot locomotion problems in the DARPA Robotics Challenge, and on human-operated semiautonomous robots, I suggest a few promising research directions:

- *Real-time optimal replanning.* Approach the goal of computing high-quality solutions with time constraints.
- *Structured motion planning.* Planners should accept a larger variety of dynamic models, objective functions, and constraints, including hybrid models and semantic knowledge. Planners will need to examine problem structure to devise good solution strategies.
- *Black-box motion planning.* Abstract the choice of planning algorithm from the API. Provide knobs for controlling solution time/quality/repeatability tradeoffs.
- *Large models from perception.* Plan quickly in the face of large, possibly dynamic point cloud models (e.g., tens or hundreds of millions of points). Incorporate uncertain initial states and predictive models from state estimators (e.g., particle filters).
- *Use real dynamics.* Cope with the nuanced behavior of low-level controllers and actuators, e.g., motor controllers, passive damping, friction, back-EMF, backlash, drivetrains, etc.

- *Debugging.* Setting up large planning problems is error-prone, as a single typo can cause an incorrect planner failure. Provide a method for explaining / debugging planner reasoning.

Geoffrey Hollinger, Oregon State University
Planning for Autonomous Information Gathering

Planning and coordination of autonomous vehicles for information gathering in unstructured outdoor environments is a key future research challenge. Emerging application domains include environmental monitoring, aerial surveillance, emergency response, and agricultural sensing. Teams of autonomous vehicles are uniquely suited for performing these large-scale monitoring and surveillance tasks because of their potential for collecting, processing, and reacting to incoming sensor data. In this sense, multi-vehicle planning is effectively a big data problem where vast quantities of sensor data are streaming continuously during operation.

Existing planning algorithms are insufficient for coordinating teams of vehicles in these large-scale monitoring domains. Key advancements must be made in (1) online adaptation and learning with changing objectives, (2) scalability to large environments and increasing team sizes, and (3) real-time reactive decision making during in-the-loop planning. Effective surveillance and monitoring platforms will need to sense and react to incoming information in real time and also communicate information back to human operators. Planning and decision making in these scenarios will allow for supervised autonomy during which autonomous vehicles act to assist scientists, first responders, and intelligence officers.

Potential avenues for solving these research challenges can be found in combining statistical machine learning methods with motion planning techniques. Recent advances in data processing and machine learning have allowed for adaptation in real time with impressive accuracy. Similarly, motion planning algorithms are utilizing sampling-based techniques to scale up and generate complex coordination of multi-robot teams with performance guarantees. By combining scalable machine learning algorithms with sampling-based motion planning techniques, we can potentially improve the adaptability and scalability of robot planning methods.

One key benchmarking problem in autonomous surveillance and monitoring is the effective reconstruction of a time-varying scalar field over a large space. For example, a team of autonomous underwater vehicles may be deployed to determine the dissolved oxygen content in a bay during an ecological event. Given some underlying model of the scalar field and a team of assets with limited speed, maneuverability, and sensing radius, the planning problem is to choose the optimal trajectories for the entire team that minimize the reconstruction error over the field. Prior solutions have employed techniques like gradient descent, submodular optimization, sampling-based planning, and model predictive control to optimize trajectories in these scenarios. However, the development of a general solution that is broadly applicable and can effectively adapt to changing objectives remains an open problem.

Wes Huang, iRobot
Planning in Human Environments

There are many planning challenges as robots increasingly operate in human environments, from gracefully navigating around those environments to assisting people with various tasks in those environments.

Robots need to navigate in shared human environments without causing any disruption. Although the problem of getting from point A to point B is well understood, doing so gracefully in a busy dynamic environment is still a challenge – people are not "just" obstacles. Aside from traffic-like conventions of right-of-way and turn-taking, nonverbal social cues and sometimes explicit verbal communication are used. No robot currently has a human adult level understanding of these issues. While there are perception problems to be solved here, there are plenty of planning problems. For example, how should a robot use its gaze to naturally or more effectively navigate? When should a robot ask people to make way for it to pass versus take another route?

Assisting seniors at home or in long-term care facilities will be an increasingly relevant source of planning problems. Most seniors prefer to stay in their own homes instead of moving to a long-term care facility; however, people's abilities decline with age. Formal and informal care and support can help seniors continue to live independently, but there can be significant cost for formal care and a significant burden for families and friends. Robots will eventually help seniors perform many personal and household tasks, but there are many planning challenges. How does a robot help a person get dressed? How does a robot prepare, serve, and clean up after a meal? How can a robot offer assistance before it is requested? To identify some general themes, these are tasks in which: a robot needs to cooperate with a human; the environment or objects involved in a task may be ill-defined or not easily modeled; or the robot needs to observe a situation and decide where it can or should play a role.

Purush Iyer, U.S. Army Research Office (ARO)
Planning for Dealing with Unawareness

Traditional planning is considered a separate module, in autonomous systems, that employs powerful scheduling algorithms to meet a set of constraints. In a future where autonomous systems are expected to rely less on detailed input from the operator and where the need to operate as part of a human-machine team will be greatest, automated planning systems could play a major role in advancing learning and reflection to improve the skill set of a robot. In particular, the automated planning component could help in staging experiments to learn and discover, as part of its current task set, the boundaries of the autonomous system's physical capabilities (say, ability to jump) and also user preferences (is its operator left-handed, right-handed or ambidextrous). The latter (for instance) would be needed in collaborative activities as, for instance, in a situation where a robot and a human need to clear debris. Thus, the desiderata are:

- Planning systems should be layered throughout the architecture of autonomous systems; in particular, they cannot be independent modules that have a fixed functionality or API. Furthermore, planning systems should be able to infer the needs of the various components of a system rather than be explicitly given a task list/ goal to schedule.
- Planning systems should improve with time, by learning, from past experience.
- Planning systems should be actively involved in the learning mechanism of the autonomous system. In particular, planning needs to account for reducing unawareness, or increasing awareness, of the system with respect to its abilities that have not been programmed explicitly.

Leslie Pack Kaelbling, Massachusetts Institute of Technology
Planning in Unstructured Environments

One critical set of challenges faced by robots come about in highly unstructured environments, such as unmodified private homes or disaster areas. The problems in these domains are numerous: the dimensionality is high, the horizon for planning is long, and the uncertainty is significant.

One important challenge is representational: how can we represent the robot's state of knowledge about its environment? It is especially difficult when the input data is multi-modal, including, for example, 3d and 2d vision, touch information, natural-language or other knowledge-base input, tactile or auditory information, etc. High-dimensional spaces must ultimately be represented with factored representations. I would suggest that a collection of heterogeneous representations, which are combined at query time, may be the most effective solution to this problem. It is common wisdom that some form of hierarchical solution is appropriate for addressing long planning horizon, but there are as yet very few workable concrete suggestions about how to do this. I believe that the only way to effectively handle uncertainty is to model it explicitly and consider it during planning. Decision theory provides a solution to optimal sequential decision-making uncertainty, but it is wildly computationally intractable. We should aggressively seek efficient alternatives, but do so in a way that allows us to understand trade-offs between computation and utility.

Household robots are a great challenge domain, because it is relatively easy to set up test-beds in regular lab space. It is important to note that the fundamental insights gained and methods developed to solve problems in the household robotics domain will generalize to a wide variety of mobile manipulation domains.

Subbarao Kambhampati, Arizona State University
Planning for Human-Robot Teaming

Most current applications of robots involve viewing them as sensors and effectors that can be remotely operated by humans. Much of the role for planning in robotics has thus been limited to path and manipulation planning (and in the case of multiple robots, task assignment). An increasing number of applications, including search and rescue, however demand that robots be full-fledged members of human-robot teams. Such teaming scenarios present a significantly richer set of challenges and opportunities for planning that go beyond path planning and task assignment. In particular, robot team members operating in such human-robot teams need task planning and replanning to engage in goal-directed reasoning, and plan and intention recognition techniques to anticipate the actions of their human team members, and even dialog planning to effectively communicate with the human team members. The teaming context poses significant novel challenges to each of these planning roles, thus precluding the direct adoption of techniques developed for their traditional settings. Our main position with respect to this workshop is to advocate investigation of the challenges faced by each of these roles of planning in human-robot teaming.

Our own ongoing work, supported in part by ONR and ARO, has been addressing several of the challenges of planning for human-robot teaming, with particular focus on task planning/replanning in the context of teaming. The planning challenges here stem both from the long-term nature of teaming tasks, distribution of world model across team-members, the open-world nature of the environment, as well as the need for supporting effective communication between the human and the robot. These in turn demand that the planner guiding the robot must deal with incompletely specified models, uncertain objectives, open and dynamically changing worlds, as well as the ability to take continual human instructions (including those that change and/or modify goals and action models), and return meaningful status reports.

We are working towards handling incomplete models through generation of robust plans; handling uncertain objectives and partial preferences through diverse plans and open world conditional goals; and handling continual state, goal and model-updates with the help of commitment and opportunity sensitive replanning. We are also collaborating with human-factors researchers to better understand planning challenges arising from the communication with the humans in the team. An interesting problem here is the automatic generation of excuses as a way of explaining planning failures to the humans in the loop, which in turn helps elicit additional domain knowledge from the humans. Finally, we are adapting plan/intent recognition techniques to help the robot better coordinate with the human team members.

Behzad Kamgar-Parsi, Office of Naval Research (ONR)
Planning for Area Surveillance

A particularly important and challenging problem is area surveillance. Currently area surveillance is performed by one or a few agents (e.g. mobile and stationary cameras) without coordination. Progress is being made in wide area surveillance with many agents, where planning and task assignments are done in a centralized manner. This requires robust communications with a lot of band-width, a powerful central processing station, and typically involves delays. Future trend in wide area surveillance is to use many autonomous heterogeneous agents that plan coordinated surveillance in a decentralized manner. Usual challenges for planning algorithms include the following: (a) information is often uncertain and incomplete, and sometimes contradictory, (b) in changing environments new information is collected that may necessitate changing goals and priorities, which may require rapid re-planning, and (c) the space in which agents make decisions (be it state space, information space, belief space) grows rapidly hence algorithms become computationally intractable. Therefore we want to develop approaches to planning that shrink the space while preserving its essential features, can deal with uncertainties and changing information and priorities, and can assess their own performance and provide a measure of optimality.

Sven Koenig, University of Southern California
Combining Task and Motion Planning for Multi-Robot Systems

Robots are becoming better at solving low-level planning tasks (such as navigation and manipulation planning). At the same time, robots also have to solve high-level planning tasks (such as task planning). It is an open problem how to best combine low-level and high-level planning since they often work on different data structures (with no clear correspondence) and according to different principles (for example, due to the continuous nature of low-level planning and the symbolic nature of high-level planning). Low-level and high-level planning are also predominantly studied by different research communities, namely robotics and artificial intelligence.

We speculate that progress on heuristic search (in particular, the study of suitable variants of the heuristic search method A*) has good chances to result in a systematic interface between low-level and high-level planning, for the following reasons: Heuristic search-based planning is now predominant for high-level planning, and there has been progress on using heuristic search for low-level planning (for example, in the context of lattice-based planning). Hierarchical search could then be used to combine low-level and high-level planning. It is very encouraging that there has been recent progress on variants of A* that are helpful in this regard, such as

- hierarchical versions of A* (that search on different abstraction levels),
- any-angle versions of A* (that do not restrict the resulting paths to lie on given graphs),
- any-time versions of A* (that produce paths quickly and then improve them over time), and
- incremental versions of A* (that use experience with previous similar search problems to re-plan faster).

However, a lot more research is required on both the individual techniques in isolation and their integration (including understanding their behavior and taming combinatorial state-space explosions) with an eye on achieving real-time planning with performance guarantees. Facilitating additional interaction between the robotics and artificial intelligence research communities would likely help to accelerate progress.

Finally, planning becomes more complicated as teams of robots need to solve tasks cooperatively, for example, because single robots are unable to solve them in isolation or because one wants to achieve robustness or parallelism. More research is required on both low-level and high-level planning for multi-robot systems in isolation and their integration.

A good benchmark problem for combining task and motion planning for multi-robot systems is construction, where a team of robots cooperatively needs to assemble a structure.

Hadas Kress-Gazit, Cornell University
Combining Task and Motion Planning with Behavioral Guarantees

Most people cannot program and everyone expects machines to perform their tasks predictably, successfully and safely. These facts lead to two major challenges in robotics that must be overcome in order to transition robots from niche application to wide impact in many domains; people must be able to instruct robots in an abstract and intuitive way, and robots must be "well behaved" at all times and in all situations. More specifically, this means that robots must a) be able to interact at a high-level, at a specification (the "what") and not at an implementation (the "how") level and b) must provide guarantees for safe and reasonable behavior most of the time, and graceful degradation when failure is inevitable.

To address these two challenges, planning algorithms must span different abstraction layers; they must seamlessly integrate low-level control designed for specific robot platforms with high-level reasoning targeted at creating complex behaviors as specified by the task. There must be formal mechanisms allowing the planning problem to be refined and abstracted while preserving the correctness of the solution at the different layers.

Since robots are physical systems with noisy sensor and imperfect actuators operating in complex environments around humans, providing absolute guarantees that a robot will never fail is impossible. To prove correctness properties, allow for graceful degradation and to inform users of possible failure modes, planning algorithms must reason about their own robustness; they must examine the assumptions that they make about the problem, they must reason about failure situations and they must provide meaningful feedback when tasks cannot be done or cannot be completed.

Vijay Kumar, University of Pennsylvania
Motion Planning

Most real world applications of robots require an operational tempo that is currently beyond the state of the art in robotics. The two exceptions to this are in semi-structured or completely structured environments such as warehouses (the Kiva systems example) or carefully mapped roads (the Google car). The central challenge for us is to be able to design planners, and more broadly perception-action loops, that can exhibit the level of performance seen in the Google car and the Kiva robots, but in less structured settings while being able to reason about uncertainty. There is a second set of challenges surrounding manipulation and other tasks in which robots must exhibit controlled contact interactions with a non-cooperative, physical environment. The DARPA Robotics Challenge will hopefully address this set of problems.

There are three directions of research that need to come together in order address the main challenges. First, in order to reason about uncertainty, efficient probabilistic approaches that successfully solve the dual integration problem - integrating over the belief space and integrating over the set of measurements that are possible at future steps - need to be developed. Second, the branch of control theory dealing with model-predictive control or receding-horizon control needs to be developed further to deal with completeness and convergence guarantees in settings with uncertainty and non-trivial dynamics. Finally, as we start building complex systems (flying and running robots) and systems that physically interact with the real world, it is necessary to come to grips with nested perception-action loops. Traditional approaches where a planner provides a feedforward or reference trajectory for the entire system do not scale well for complex systems.

Possible benchmarks include

- *3-D Kiva*: Imagine a three-dimensional indoor environment where you want to send 100 flying robots to respond to requests by humans by flying to goals safely and efficiently. What architecture (decentralized versus centralized or hybrid) must we consider for solving this problem? Does the cloud change the fundamental multi-agent coordination problem space? What if the environment were not fully known? How do robots adapt to unstructured environments and how do they share information to build a model that others can leverage? Does the cloud lend itself to new real-time, planning and control algorithms?
- *Humanoid*: DRC but real-time performance at a tempo that is significant and meaningful for search and rescue.

Maxim Likhachev, Carnegie Mellon University
Integrated Planning in Complex Worlds

Real-world in its full glory contains too many factors and unknowns to account for all of them upfront while developing an autonomous robot. The first challenge for planning therefore is to operate on a "deeper" level than just a static representation of the world. Planning representations need to "adapt" to the tasks the robot executes, the environment it works in, the experiences it gathers and demonstrations it receives. For example, human demonstrations for how to open doors, for how to sort items arriving on a conveyor or for how to assemble furniture shouldn't be used as simple re-play motions but instead should be used by the robot to infer how relevant objects behave and how to construct planning representations that are effective for given tasks.

The second challenge is in developing planning frameworks that integrate tighter with other modules in the system such as perception, control and task planning. The planners need to reason about the weaknesses and strengths of these modules and generate plans that avoid their weak spots and capitalize on their strengths. Furthermore, these weaknesses and strengths should be revised based on the robot experiences. For example, a micro-aerial vehicle has limited payload leading to noisy sensing and actuation. Generated plans should therefore exercise as much as possible the controllers that are robust such as visual servoing towards easily detectable landmarks, and avoid generating segments of the paths that require the use of fragile controllers. Similarly, a motion planner for a manipulator needs to reason about the limitations of grasping and the constraints of higher-level task planning. To address these challenges without the explosion of the state-space, we need to research planning representations that vary in their dimensionalities, model limitations and strengths of modules and revise all of it based on experience.

Third, for robots to operate alongside humans, we need to make them more consistent and predictable. Consider, factory workers working side-by-side with industrial robots, surgeons getting help from robots bringing medical tools and soldiers working with robots manipulating IEDs. In all of these scenarios, humans need to be able to predict what action the robot is going to execute next and how. This allows humans to intervene when necessary as well as plan their own actions. The predictability of motions is also beneficial to task-level planning and remote execution monitors which need to predict the behavior of the robot and re-plan or intervene as necessary.

Kevin Lynch, Northwestern University
Real-Time Motion Planning for Uncertain Hybrid Mechanical Systems

As robots become increasingly dynamic (examples include the early hoppers and runners from Raibert's group, Boston Dynamics' Atlas, RHex, Cheetah, and the ParkourBot), extreme locomotion becomes possible. Methods for real-time motion planning, estimation, and control are required to take advantage of these robots' capabilities.

As our challenge problem we consider motion planning for robot parkour. Given a dynamic model of the robot and a well characterized environment (including possible footholds and handholds), the problem is to plan a sequence of runs, jumps, vaults, swings, etc., to traverse a complex environment by taking advantage of the various constraints present.

This motion planning problem has four important characteristics:

- *Hybrid*: The equations of motion of the system change depending on the contact constraints currently active (e.g., which hands or feet are in contact, and whether these contacts are sliding or sticking). Transitions between different regimes may be characterized by impacts.
- *Mechanical*: The equations of motion in each regime are not arbitrary, but are characterized by an inertia tensor, Coriolis terms due to the Christoffel symbols of the inertia, potential forces, etc.
- *Uncertain*: Uncertainty is introduced due to errors in the robot and environment model as well as a lack of controllability to recover from all perturbations within a particular regime.
- *Real time*: The usual decoupling of "planning a trajectory" and "controlling to follow the plan" does not suffice due to lack of controllability within each regime separately. Therefore real-time multi-step replanning is required.

A successful motion planning framework will account for the several possible dynamic regimes and their transitions; make use of the structure of the equations of motion for increased efficiency; explicitly model the propagation of uncertainty and its mitigation, shaping, or increase by various dynamic primitives; and fine tune the nominal plan during execution.

Such a framework is likely to include a non-real-time component that addresses the combinatorial problem of searching for a sequence of maneuvers and generating a nominal motion plan with a high likelihood of success (e.g., modeled on traceurs' careful assessment of an environment) as well as gradient-based real-time tuning of the nominal plan to account for the evolving state uncertainty (belief distribution) maintained by a belief filter during execution.

The basic framework is likely to apply to hybrid robot manipulation as well.

Dinesh Manocha, University of North Carolina at Chapel Hill
Motion Planning

Algorithmic motion planning has been actively studied in robotics and related areas for more than three decades. There is a rich collection of motion planning algorithms based on computational geometry and algebraic methods, local or potential field techniques, randomized sampling, handling kinodynamic or non-holonomic constraints, optimization methods, etc. Most of these algorithms have been successfully used for CAD/CAM, bioinformatics, computer gaming and other applications. At the same time, advances in manufacturing technologies, sensing, and actuator devices have led to the development of powerful robots, including humanoid robots and general-purpose and programmable mobile manipulators. However, there has been limited use or application of motion planning algorithms on these "physical robots" to perform various autonomous or navigation tasks. This is due to issues related to dynamic constraints, modeling uncertainty, perception, as well as real-time computation of motion strategies on the robotic platform.

Recent developments in programmable mobile manipulators, along with open-source operating systems and environments (e.g. ROS), seem to open up many new possibilities to bridge this gap. Furthermore, the availability of better sensors (depth sensors, for example) and high computing power (multi-core CPUs and many-core GPUs) makes real-time planning algorithms for high-DOF physical robots feasible. However, there is relatively less work in terms of developing fast planning capabilities that can take into account model and sensor uncertainty.

We need to develop motion planning algorithms for physical robots that cannot be separated from the underlying application or context. These include necessary interfaces with low sensory-motor control and symbolic reasoning than to explore geometric issues that do not correspond to major bottlenecks in the physical world. In many ways, we have reached a level of maturity in terms of classic geometric motion-planning problems, like the ones related to the curse of dimensionality or planning in "narrow passages". However, this maturity can be considered at a conceptual level and we need to extend that along other dimensions. For example, some of the widely used motion planning techniques are based on probabilistic approaches that perform reasoning not only in robot configuration spaces but can also be extended in "augmented" spaces, like state spaces that are used to model control issues or belief spaces used to account for decisional issues. Moreover, we need to continue development of good software tools, such as FCL, OMPL, MoveIT, OpenRAVE, to integrate these algorithms into new applications.

Jeremy Marvel, National Institute of Standards and Technology (NIST)
Planning for Robot Collaboration in Manufacturing

From my perspective, one of the largest challenges facing current and next-generation robotics in manufacturing is that of intrinsic collaboration. Tasking and re-tasking robots for manufacturing applications is already quite difficult due to limitations stemming from proprietary hardware and interfaces. Getting these robots to also collaborate intelligently requires considerably more effort. Currently, robot-robot collaboration is achieved by explicitly coordinating motions based on a set of event-based constraints. In contrast, human-robot collaboration is limited to instances of simple colocation, rule-based reactive actions, and safety-related mechanisms. These, however, are not true collaborations in the sense of robots and humans coopering with a shared understanding of the manufacturing process to achieve some desired end goal. Getting robots to collaborate with humans or other robots requires a significant advancement in autonomy and performance assurance. This necessitates improvements in situational awareness, real-time coordination of motions based on sensor feedback, dynamic planning and re-tasking to compensate for environmental and operational uncertainty, and task decompositions and representations to enable more accurate mappings of observations to responses. Combined with correlated test methods and metrics, advances in these fields will enable greater confidence in the capability of collaborative robots to complete their assigned tasks correctly while meeting their assembly performance objectives.

Don Sofge, Naval Research Laboratory (NRL)

Safe Navigation Planning in Crowded Environments

A key planning challenge faced by mobile robots in the real world is the need to operate safely amongst people, and other robots, while avoiding collisions. People possess an innate skill for moving through crowds and passing one another (both on foot and while driving) while generally avoiding collisions. This skill relies on understanding how we move through space, observing others, predicting others' paths, and automatically adjusting our own motion to avoid collisions. One possible approach to a solution would be for robots to have a different (i.e. better) representation of space with direct feedback to motion control parameters (e.g., swerve left, or veer right). In order to be successful the robot should not only sense the presence of others (both human and robot), but also anticipate where they will move next. A benchmark problem that would spur research in this area might involve having a mobile robot navigate through a crowded mall setting, or perhaps a crowded conference or an outdoor concert, while avoiding collisions with others.

Mike Stilman, Georgia Institute of Technology
Planning for MacGyver Robots

Future robots should have the capability to use their entire bodies and environments in order to solve complex tasks. Previously, we have developed planning algorithms for Navigation Among Movable Obstacles (NAMO) where robots could move objects out of the way in order to achieve navigation or manipulation tasks.

The MacGyver Concept: Our group is moving forward from the initial ability to interact with environment objects. We believe that robots should not only remove objects, but also use them to help in achieving tasks by building simple machines. For instance, consider robots that will use a board to build a bridge or use a broom as a lever to create additional force from distance. Such static and dynamic concepts should be incorporated into autonomous planning by humanoid robots and mobile manipulators. We believe that utilizing the entire environment and the robot's entire body is critical to future advances in robot autonomy.

Planning in Constraint Space: In order to make it possible for robots to utilize their environments we propose to plan in the space of constraints. Our preliminary work, presented at ICRA 2013, shows that a robot can create structures from environment objects without explicitly deciding their locations. At each step of the plan the robot solves an optimization problem to verify the validity of the plan step. Once a plan is determined, the constraint satisfaction problem is fully solved and the plans are finalized.

Planning with Uncertainty: It is challenging for a robot to work with environment objects that have unknown or uncertain properties. We are building on our work from the NAMO domain to incorporate uncertainty through decision theory as well as hierarchical planning in the now. We aim to maximize the robot's capacity to execute actions and quickly identify novel circumstances and plans that can overcome them.

Manuela Veloso, Carnegie Mellon University
Planning for Symbiotic Robot Autonomy

Actual robots in the real world are mostly used in targeted industrial, in space, and military applications. They execute complete planned motions, motion plan for given high-level goals, or are remote operated. Besides the Roomba robots, and instances of service robots in hospital environments, there are not many autonomous robots in the real world. There are the Google and other self-driving cars, which are becoming a reality.

The real world offers challenges to robot autonomy, at all levels, namely perceptual, cognitive, and actuation. Focusing on the real world where humans habitate, robots will further concretely face navigation, timing, quality, and interaction challenges. Planning involves having models of the world, and models of actions. It is challenging, or most probably impossible, to find appropriate representations for general-purpose models, independent of the task and goals. Furthermore, the real world is dynamic, includes lots of uncertainty, unpredictability, and other external, at best poorly modeled, actuators, such as other robots and people.

I believe that the challenges offered by the real world towards our real autonomous robots may be unsurmountable in the close future, unless the robot limitations themselves are explicitly included in the planning, so that the robot can ask for help from available external resources, such as humans, other robots, and the web for information and crowdsourcing.

Planning can then include the representation of world features (e.g., location of objects, understanding of language), and actions (e.g., going up a floor, picking up any object) that the robot may not be able to assess or execute, but for which the robot will be able to ask for help.

Planning can also help attack autonomy problems at different levels of abstraction and different levels of uncertainty, such that robots can evaluate alternatives courses of action, and select ones that address complex joint objectives.

The science of "planning for robots in the real world" is an integrated approach for planning, execution, replanning, and continuous multi-model learning, driven by experience.

The RoboCup@Home tasks are successful examples in home environments. Mobile service robots, like the CoBot robots at Carnegie Mellon University, are also good examples of tasks that include several levels of planning can be benchmarked. A benchmark restaurant servicing task would be able to include multiple types of robots with different capabilities, such as non-movable Baxter-like robots with arms and non-armed movable CoBot-like robots. We could also have mobile robots collect situated, localized environment data, e.g., temperature, wifi, noise, pollution levels.

Peter Wurman, Kiva Systems
Multi-Robot Path Planning

Kiva Systems use fleets of small robots to move inventory shelves (pods) around in warehouses. The inventory shelves are carried to replenishment or pick workers who stand at stations along the perimeter. By eliminating the walking that these people would do in a traditional warehouse, the workforce becomes 2-3x more efficient.

The largest facilities that use Kiva have over a thousand robots on a single, continuous floor, typically laid out as a grid of cells. The robot density can reach as high as one robot per every 6 cells of floor space. One of Kiva's key challenges is what is generally known as multi-vehicle routing, however there are aspect of Kiva's problems that are not well-captured in typical research treatments. Thus, the Kiva scenario represents a multi-vehicle path planning problem of enormous size that is ideally suited to a challenge problem.

A Kiva floor is generally broken up into three regions. In the storage area, pods are arranged in dense, city-like blocks. The perimeter is usually lined with queuing areas for each station where robots with the next deliveries for that station can line up. Between the station queues and the storage areas is an interchange area.

A typical robot mission starts with the robot having just set down an inventory pod. Having completed one mission, it immediately becomes available for a new mission. It is assigned a new goal of fetching a specific inventory pod and bringing it to a specific station. At the station, the human will complete the pick or replenishment task. The duration of the human's task is imperfectly predictable. Once the human is done interacting with the pod, the robot must return the pod to the storage area.

Three things distinguish the Kiva multi-vehicle path planning problem:

- The number and density of robots.
- The delivery pattern in which a continuous stream of robots descends on a small number of perimeter stations.
- The variable human task at the station that causes robots to release from stations as a stochastic process.

A challenge problem that presented an abstract version of a Kiva floor with its three key distinguishing features would spur research into multi-vehicle coordinated path planning. A breakthrough in this area has the potential to significantly reduce the number of robots needed to run these facilities.