Mobile Manipulation Hackathon

Moving Into Real-World Applications

By Máximo A. Roa, Mehmet Dogar, Jordi Pages, Carlos Vivas, Antonio Morales, Nikolaus Correll, Michael Görner, Jan Rosell, Sergi Foix, Raphael Memmesheimer, and Francesco Ferro

The Mobile Manipulation Hackathon was held in late 2018 during the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) to showcase the latest applications of wheeled robotic manipulators. The challenge had an open format, where teams developed an application using simulation tools and integrated it into a robotic platform. This article presents the competition and analyzes the results, with information gathered during the event and from a survey circulated among the finalist teams. We provide an overview of the mobile manipulation field, identify key areas required for further development to facilitate the implementation of mobile manipulators in real applications, and discuss ideas about how to structure future hackathon-style competitions to enhance their impact on the scientific and industrial communities.

Mobile Manipulation

Autonomous mobile manipulation combines two fundamental robotic skills: mobility in an environment and the manipulation of objects. The ability to do both simultaneously opens numerous applications in diverse areas, including manufacturing, logistics, home automation, and health care. Such applications typically require complex (structured and unstructured) manipulation. They also demand navigation in large spaces, possibly in cooperation with human beings and other robotic systems.

Mobile manipulation is a complex field. Mobility introduces additional pose uncertainty to the manipulation problem while limiting the available perception systems and introducing constraints to the navigation issue, which now must include one or more arms mounted on a robot. Mobile manipulation is also a systems challenge, requiring designers to draw on multiple fields: perception, navigation, tasks, path and grasp planning, control, error recovery, human–robot interaction, and robotic hardware development. Each field is an area of research in its own right, but the particular challenge in mobile manipulation is to obtain an integrated system that can combine a large variety of hardware and software components to increase the range of tasks that a robot can perform, while decreasing the dependency on prior information and broadening the awareness the robot has of its current situation.

Digital Object Identifier 10.1109/MRA.2021.3061951
Date of current version: 17 March 2021

1070-9932/21©2021IEEE

Since the complexity of mobile manipulation lies at the interface of the fields mentioned in the previous paragraph and because any significant experimentation will require not only mastery of a variety of techniques but system integration and hardware acquisition, it is difficult to establish mobile manipulation as a field of its own. Similarly, it is not clear what the commercial applications of mobile manipulation really are. While performing truly human-like tasks is only possible when combining mobility and manipulation, the high cost and limited performance suggest commercial solutions that are mobile only (such as floor cleaning and transportation), that are manipulation only (such as conventional robotic assembly lines), or that constrain a system in such a way that manipulation remains trivial, for example, picking up and moving entire shelves in warehouses. However, other applications, such as telepresence and remote assistance systems, are moving toward demanding some way to interact with objects and people, for instance, in assistance scenarios for senior citizens. While industrial use cases might be able to complete multiple tasks using fixed-base manipulators, a single, flexible mobile platform could autonomously take over multiple responsibilities in different locations, thus possibly improving returns on investment, which is especially important for small and medium enterprises that cannot afford multiple static robotic platforms.

To address these challenges and build a community around mobile manipulation, the IEEE Robotics and Automation Society (RAS) Technical Committee (TC) on Mobile Manipulation, with its members and collaborators, organized a hackathon that ensured participants had an even playing field by providing a complete mobile manipulation system that offered a basic level of operation. This enabled the community to showcase 1) its work in relevant subfields, such as grasping, manipulation, perception and motion planning, and 2) application domains that might truly benefit from a mobile manipulation solution.

Related Hackathons and Competitions

The hackathon phenomenon has been described, in the context of digital innovation, as an appropriate vehicle to bring people from different disciplines together as well as to actually engage a community with a particular topic [1]. Consequently, a body of work exists on how to design a hackathon to optimize the desired outcome in terms of networking [2], learning [3], and broadening participation in computing [4]. In its purest form, the hackathon format unites groups of unrelated people to share knowledge and work toward a solution, learn from one another, and potentially form long-term connections. Given the current state of the art in hardware and software, we deemed it unlikely that significant insights would emerge from an ad hoc event in which teams were formed at a conference venue, with no previous contact and no opportunity to learn about available tools. Instead, the hackathon was organized as a multistage competition from which finalist teams were selected based on an initial entry derived mostly from simulation results.

Robotic competitions have aims that are very similar to those of hackathons, but they operate on a different timescale (months of preparation versus a single day, for example) and emphasize robust solutions more than prototypes. Competitions have a long history in robotics and artificial intelligence, with their entries often determining the state of the art for years to come, such as in localization [5] and autonomous driving [6]. They can also lead to unexpected insights into a systems challenge's real problems. For example, the Amazon Picking Challenge [7] has shown that warehouse picking is less a grasping and manipulation challenge (the majority of teams used suction) than a perception problem. Similarly, the Industrial Assembly Challenge [8] has demonstrated that perception and planning are secondary when dealing with sufficiently restricted and well-defined problems.

Despite much progress in these research domains, open-loop control and mechanical templates and fixtures usually excel in such scenarios. These insights can then be used to refine a competition format to push a community in a desired direction. Many successful contests focused on robotic manipulation have been organized in recent years. IROS Robotic Grasping and Manipulation Competitions were organized for 2016 [9], 2017, 2019, and 2020 (https://rpal.cse.usf.edu/competition_iros2020/). They included a fixed set of tasks in areas such as service (spooning peas and preparing iced tea), manufacturing (assembly/disassembly), and logistics (bin picking). The undertakings did not require mobility. The Real Robot Challenge (https://real-robot-challenge.com/) was organized by the Max Planck Institute for Intelligent Systems (MPI-IS) in 2020. It was based on the remote execution of submitted software on a robotic hand. There was a fixed set of tasks, such as grasping and pushing, which did not require mobility. The IEEE International Conference on Soft Robotics also holds a competition (http://www.robosoft2019.org/robosoft_competition.html) with a manipulation challenge that emphasizes soft manipulators. Similarly, the tests do not require mobility.

There have also been recent competitions that targeted mobile manipulation. The FetchIt! Mobile Manipulation Challenge was held at the 2019 IEEE International Conference on Robotics and Automation [10]. The task was to assemble a kit formed by six objects obtained from stations around a designated arena, combining navigation and manipulation skills. Similarly, the RoboCup@Home competition (https://athome.robocup.org/), using the Toyota Human Support Robot (HSR) [11] as the official platform, includes a set of tidying-up and service tasks in living room and kitchen setups, requiring mobile manipulation. RoboCup@Home also encourages teams to make "open challenge" demonstrations (i.e., free demonstrations determined by the teams, instead of a fixed set of tasks), although these are not the main focus, as they are performed during off hours and do not necessarily include awards [12]. The Smart Cities Robotics Challenge [13], which is organized as part of the European Robotics

League and builds on the success of the European Robotics Challenge (EuRoC) [14], also includes a fixed set of mobile manipulation tasks, such as delivering coffee shop orders and shopping pick-and-pack procedures.

The unique feature of our hackathon, compared to the preceding competitions, is that it positions mobile manipulation together with open demonstrations at center stage. As explained, multiple mobile manipulation competitions have focused on fixed sets of tasks. This has the advantage of creating benchmarks that enable progress to be objectively measured. Therefore, they are crucial to the community. However, we believe that an open format also has a place. It enables 1) teams to demonstrate their core research innovations more directly and 2) the community/audience to be informed about the state of the art for a rich variety of tasks. With the Mobile Manipulation Hackathon, our goal has been to push teams to perform their own research demonstrations and to identify tasks that the community is working on.

The Field of Mobile Manipulation

Merging mobility and manipulation, mobile manipulation systems need to overcome some of the most difficult challenges in robotics, including the following:

- *Generality*: Mobile manipulation systems must perform a variety of tasks, acquire new skills, and apply those abilities in novel situations. They must be able to continuously adapt and improve their performance.
- *High-dimensional state space*: Versatile robotic systems must be equipped with many actuators and sensors, resulting in high-dimensional state spaces for planning and control.
- *Uncertainty*: The ability to locomote, the required generality in task execution, and the use of multiple sensors and actuators make it impractical to engineer an entire environment for a task. As a result, mobile manipulation systems have to explicitly address problems that arise due to the uncertainty of sensing and actuation.
- *System complexity*: Mobile manipulation systems require the integration of a large number of hardware components for sensing, manipulation, and locomotion as well as the orchestration of algorithmic capabilities in perception, manipulation, control, planning, and so on.

The mobility of these systems can take multiple forms depending on the environment: air/space (drones, planes, helicopters, and satellites), water (ships and submarines), and land (wheeled and legged robots). In air and space, mobile manipulation systems often take the shape of aerial vehicles carrying some sort of manipulator [15], [16], e.g., grippers [17] and multilink arms [18], [19] attached to a rotorcraft; they may also be built as manipulators endowed with some flying mechanism, e.g., rotors [20]. A significant challenge for these systems is to maintain flight stability during object manipulation, which limits the range of operations that can be performed. This coupling between the control of mobility and manipulation also exists in water, where robots need to maintain a stable pose while experiencing additional forces due to object manipulation [21], [22]. Land is the most common environment for mobile manipulation. Humans live on land, and therefore a larger variety of mobile manipulation tasks can be found there. Furthermore, the control of mobility and manipulation can be decoupled more easily on land, compared to in-air and underwater manipulation. A land robot can attain a statically stable configuration and, for small enough forces, avoid the need to balance during manipulation.

Two common forms of mobility on land are legs and wheels. Legged locomotion and bimanual manipulation are typically combined in humanoid robots, e.g., [23]. Even though planning and control for legged locomotion can be more complex than for wheeled locomotion, legs can be advantageous depending on ground characteristics. Particularly for search-and-rescue operations, where debris, obstacles, and steps are expected, legged mobile manipulation is preferred. Such systems dominated, for example, the DARPA Robotics Challenge [24].

The most common and versatile mobile manipulation systems, however, are wheeled. They strike the right balance between ease of mobility and manipulation and access to most human environments. The development of wheeled mobile manipulators has unfolded during the past 35 years. The first prototype of a mobile manipulator was the Mobile Robot (MORO), in 1984 [25]. Initial attempts to mount robotic arms on mobile platforms happened during the 1990s, with robots such as the Hostile Environment Robotic Machine Intelligence Experiment Series (HERMIES) (Hostile Environment Robotic Machine Intelligence Experiment Series) [26] and KAMRO (Karlsruhe Autonomous Mobile Robot) [26] and the Karlsruhe Autonomous Mobile Robot (KAMRO) [27]. Efforts toward coordinating base and arm motions also received seminal contributions during these years [28], [29]. Since then, there have been many developments and highlights in wheeled manipulation systems. Hvilshøf et al. [30] surveyed up to 30 different prototypes developed up to 2011. The main application domains of mobile manipulation systems ranged from domestic service [31], [32] through space [33] to industry, with commercial solutions from, e.g., KUKA (https://www.kuka.com/en-gb/products/mobility/mobile-robots) and Neobotix (https://www.neobotix-roboter.de/produkte/mobile-manipulatoren).

Around 2010, a wave of more advanced, bimanual, multipurpose wheeled manipulators arrived (Figure 1), with systems such as the Personal Robot 2 (PR 2) [32] developed at Willow Garage, the Care-O-Bot 3 [34] developed at the Fraunhofer Institute for Intelligent Analysis and Information Systems, HERB [35] developed at Carnegie Mellon University, Rollin' Justin [36] developed at the German Aerospace Center (DLR), and the ARMAR series developed at the Karlsruhe Institute of Technology [37]. This wave represented a milestone since it coincided with the introduction to the community of the Robot Operating System (ROS) [38] which, through its modular structure and components such as the ROS Navigation Stack (http://wiki.ros.org/navigation) and MoveIt! (https://moveit.ros.org/), made it easier to build complex software systems. 2010

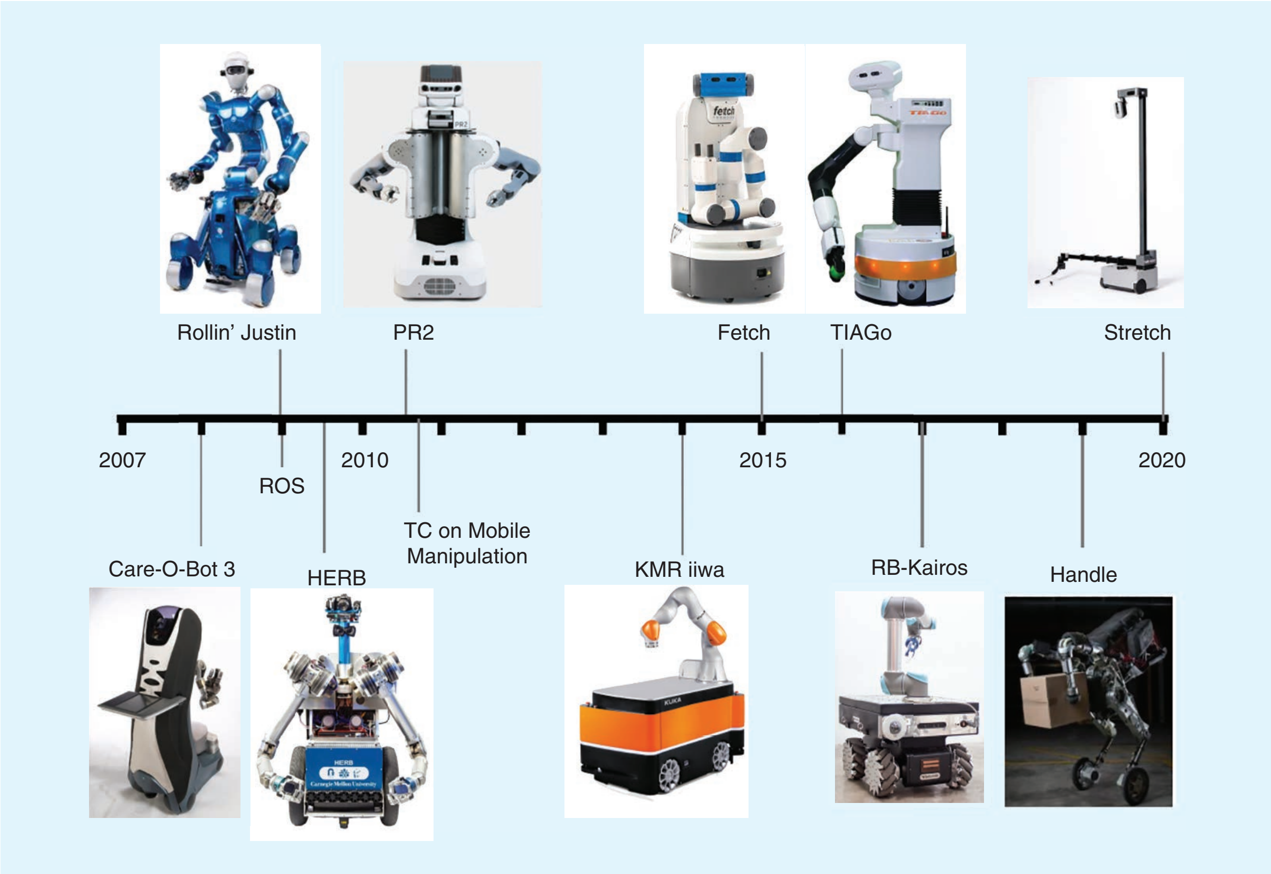

Figure 1. Wheeled robotic manipulator development during the past decade.

was also the year when the RAS TC on Mobile Manipulation was established.

Although this series of wheeled manipulation systems created a lot of excitement and interest in mobile manipulation and its applications, it also revealed challenges. The cost of building such systems was especially prohibitive for large-scale use and adoption, hampering the development of a larger research community. Early adopters of mobile manipulators included the military and law enforcement, both of which used robots for dangerous missions such as defusing bombs and the remote inspection of installations. In the past few years, a greater number of simpler yet fully integrated and commercially oriented wheeled manipulation systems have been observed. These developments include Take It and Go (TIAGo) (http://pal-robotics.com/robots/tiago/), which is unimanual, and TIAGo++ (bimanual) by PAL Robotics; the Fetch Mobile Manipulator (https://fetchrobotics.com/robotics-platforms/fetch-mobile-manipulator/) by Fetch Robotics (available for researchers); Swift (https://www.iamrobotics.com/our-solution/) from IAM Robotics; RB-1, RB-Kairos, RB-Eken, and RB-Vulcano systems from Robotnik (https://robotnik.eu/products/mobile-manipulators/); industrially oriented KUKA Mobile Robotics (KMR) (https://www.kuka.com/en-gb/products/mobility/mobile-robots/kmr-iiwa); and the assistance-oriented Toyota HSR (https://www.toyota-global.com/innovation/partner_robot/).

The field is still evolving, and interesting concepts have recently been presented, such as Handle (https://www.bostondynamics.com/handle) from Boston Dynamics and Stretch (https://hello-robot.com/product) from Hello Robot. Figure 1 presents a timeline of the development of these wheeled robotic manipulators. The systems target applications such as supplying and transporting parts in manufacturing and logistics operations and object delivery and human interaction in health and personal care. Yet the mobile manipulation market remains a niche, and estimates of the market for these types of systems are difficult to obtain. For instance, the latest report from the International Federation of Robotics does not include mobile manipulation systems as a separate domain but combines the information with overall statistics according to application areas (industrial, logistics, medical, field robotics, defense, and so on) [39]. However, it is recognized that the combination of mobile platforms with collaborative robots opens the door to new use cases and could substantially increase demand for robotic systems.

With advances in the development of mobile manipulators and the number of potential applications comes the need for standardization, especially in areas related to safety during human–robot collaboration. There have been important efforts in this direction, even though there is still uncertainty about regulations covering the use of mobile manipulators. Depending on the area of application, different regulations apply. For example, in industrial settings, many integrators apply International Organization for Standardization (ISO) 10218-1 regarding the safety of industrial robots, ISO Technical Specification 15066 related to collaborative robots when manipulators are in use, and either ISO 3691-4 or the former EN 1175-1:1998 when robots navigate by keeping their arms static to prevent accidental contact. More recently, the American National Standards Institute (ANSI) published ANSI/Robotic Industries Association R15.08-1-2020 targeting safety requirements for industrial mobile robots. On the other hand, health-care applications may require ISO 13482, which conveys safety requirements for personal care robots.

Mobile Manipulation Hackathon

The Mobile Manipulation Hackathon was conceived to encourage participants to implement demonstrations that showcase the applicability of a wheeled robotic manipulator. The call was open to contributions from any field (e.g., learning by demonstration, grasp planning, and human–robot interaction) and domain (e.g., logistics, health care, and services), as long as submissions could be integrated into a predefined robotic platform to execute a mobile manipulation application. Participating teams proposed applications and demonstration scripts. Hackathon organizers evaluated and filtered the most promising and appropriate proposals to ensure that they fit the scope and purpose of the activity.

This approach is different than that in most other competitions, which are based on detailed task descriptions for participants to execute. In our experience, this method delivers overfitted and engineered solutions to specific projects that are not easily generalizable and therefore usually have little impact on associated research fields. In an open domain, such as mobile manipulation, we feel that this is ineffective. As an alternative, we support an open format in which teams can demonstrate their knowledge through tasks they propose.

Figure 2. The TIAGo wheeled manipulation platform.

Mobile Manipulation Platform

To ease and motivate hackathon participation, we proposed a common mobile manipulation platform, TIAGo (http://pal-robotics.com/robots/tiago). It is endowed with an arm that has seven degrees of freedom (DoF), a liftable torso, and a pan–tilt head equipped with a red–green–blue–depth (RGB-D) camera and stereo microphones (Figure 2). Participants benefited from completely ROS-based interfaces and a simulation environment to develop, in their own labs, an initial proposal for their demonstrations. Their applications were required to make effective use of a mobile robot (e.g., tasks could not be solved with only a fixed-base manipulator). Entrants could exploit publicly available ROS tutorials and demonstrations (http://wiki.ros.org/Robots/TIAGo). Applications developed through simulation were later implemented on a real robot, with the support of PAL Robotics researchers and engineers. The company sponsored the competition by lending three TIAGo robots that were available during the final event. In addition, selected teams were allowed to spend a week testing and tuning their demonstrations at the PAL Robotics' site during the month before the contest.

Competition Procedure

Hackathon submissions had to be prepared well in advance. With this in mind, we designed a procedure that gave teams enough time to develop their proof of concept and the organizers enough time to set up the selection process, which consisted of the following milestones:

- *Call for participation (early 2018)*: An announcement was distributed through several mailing lists, with descriptions of the hackathon scope, goals, procedures, and timeline.
- *Expression of interest (March 2018)*: Interested parties submitted a letter introducing their team and presenting their proposed application and demonstration, background, planned use of equipment, and so forth.
- *Feedback to teams (April 2018)*: Organizers provided suggestions about how to create a high-impact demonstration.
- *Entry submission (June 2018)*: Teams submitted a video and a short technical report explaining, in detail, their proposed demonstration and their original approach/technology. At this stage, simulations were allowed in the videos.
- *Announcement of finalists (July 2018)*: Six finalists were chosen. The selection criteria included development maturity, the novelty and relevance of specific components, and application relevance.
- *Support in Barcelona, Spain (September, 2018)*: The finalists were given an opportunity to test and tune their demonstrations for one week on a robot at PAL Robotics' headquarters.

- *Competition (1–5 October 2018)*: The final event took place during IROS, in Madrid, Spain. The event lasted three days, and two teams participated each day. Teams were given a whole day with the robots to prepare their demonstration, which was presented in the late afternoon. A committee of three international experts—Prof. Jeannette Bohg, Stanford University; Graham Deacon, Ocado Technology; and Prof. Weiwei Wan, Osaka University—evaluated the demonstrations according to novelty, academic merit, industrial merit, integration quality, and the impressiveness of the presentation. The winners were announced at the end of the third day.

Competition Results

Thirteen teams submitted entries. The teams came from India (two), Germany (two), Spain (three), Switzerland (one), Singapore (one), Japan (one), Brazil (one), Mexico (one), and the United States (one), and they proposed an extensive variety of applications, as listed in the following (some applications were proposed by several teams):

- imitation learning of manipulation tasks
- a robotic home assistant
- a robotic assistant in a hospital
- a robotic feeding assistant
- an autonomous mechanic assistant
- an autonomous librarian
- an autonomous bartender
- gardening applications
- item picking in logistics scenarios.

The six finalist teams are described in Table 1, and their demonstrations appear in Figure 3. A video overview of the competition is publicly available (www.youtube.com/watch?v=mt7JGXHb8jQ). Due to the high quality of the demonstrations, the jury decided to select two winners, teams TAMS and Robotics.SG. Their demonstrations are summarized in the following:

- *TAMS*: Members implemented a software system that converted TIAGo into a bartender, pouring drinks and cocktails to clients from behind a counter (Figure 4). The robot recognized a person sitting in a predefined location at the counter and approached to take an order. It instructed the person to point to a drink on a typical cocktail menu, detected the menu's pose on the table via keypoint detection, and extracted the person's fingertip via contour detection and heuristic filtering. It could tell if the person was trying to fool it by pointing to something that was not a drink. Once the desired drink was identified via deictic interaction, the robot proceeded to a separate table where liquor bottles were stored and created a composite manipulation plan to retrieve the required ingredients, transport them, and pour them into a transparent glass in front of the customer. The glass was identified using an infrared image from the RGB-D camera. A composite motion plan was generated to pour a specific amount (parameterized by duration) into the glass, without spilling during reaching motions.
- *Robotics.SG*: The robot was used to reshelve products that were returned to a convenience store (Figure 5). It picked up a tray that held the items, identified the objects inside the tray (verifying, as well, that the tray was not empty), and planned the required motions to put the items on a corresponding shelf. While setting up, the robot scanned and stored a map of the area, including the location of items on different shelves. The item identification was performed using a pretrained, learning-based perception approach, which also delivered object poses. The acquisition of images for training the perception system was performed through an in-house-developed rotatory platform to scan the shapes and textures of objects. Once an

Table 1. The finalist teams.

Team	Affiliation	Country	Number of Members	Demonstration
Homer	Koblenz University	Germany	Two	Imitation learning of human actions www.youtube.com/watch?v=Pf91wv2ddQE
Robotics.SG	Nanyang Technological University, Panasonic R&D Center Singapore, Hand Plus Robotics, and Panasonic Connected Solutions	Singapore	Six	Placing an item in an e-commerce warehouse www.youtube.com/watch?v=_3wZ3J6NWCc
IRI	Technical University of Catalunya/ Spanish National Research Council	Spain	Three	Adaptive robotic feeding assistance www.youtube.com/watch?v=dM9DoZ2z6To
PMM Tohoku	Tohoku University	Japan	Five	Dexterous liquid-pouring in a domestic situation
TAMS	Hamburg University	Germany	Five	TIAGo as a bartender www.youtube.com/watch?v=AOkhmyDtDfQ
IOC-AUDECO	Technical University of Catalunya/ Institute of Industrial and Control Engineering	Spain	Ten	TIAGo serving drinks www.youtube.com/watch?v=VocnVbh5Nq8

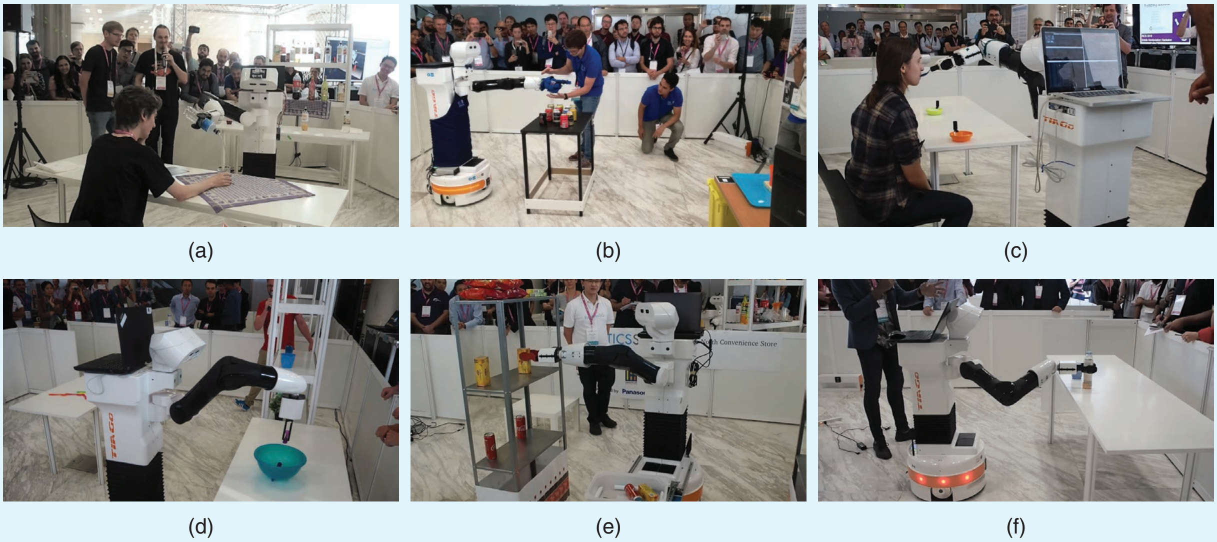

Figure 3. The finalist teams' demonstrations in the live competition. (a) TAMS. (b) IOC-AUDECO. (c) IRI. (d) Homer. (e) Robotics.SG. (f) PMM Tohoku.

Figure 4. The TAMS demonstration. (a) A user points at a menu to order a drink. (b) TIAGo moves to a bar to retrieve the liquor. (c) The robot pours the (real) liquor into a glass.

Figure 5. The Robotics.SG demonstration. TIAGo (a) retrieves a bin with returned items, (b) navigates a store by using a prerecorded map, and (c) places an item on the required shelf.

individual object pose was defined, a grasp motion was planned to pick up the item. Checkpoints were defined to verify whether a grasp was successful or not. The robot then navigated to the required shelf to place each item at its intended location.

Survey of the Competition

To compare the effort for the competition and its relationship to the research performed by each team, we distributed a survey via email to the finalists. The survey contained 19 questions, and the answers were provided in a free text format. The questionnaire was as follows:

1) *Team survey*:
   - team name
   - institution(s)
   - number of team members (including a breakdown by academic degree)
   - previous experience in competitions.
2) *Development process*:
   - Did you develop the system from scratch? (If not, provide a previous publication, if possible.)
   - What was the estimated time for developing the demonstration, in person months?
3) *Demonstration/system description*:
   - description of the demonstration
   - sensors used for the demonstration (tactile, vision, microphones, and so on)
   - hardware adaptations/additional tools for the demonstration
   - software framework
   - simulation tools
   - motion planner
   - external libraries/dependencies
   - robot autonomy (full or shared)
   - type of control
   - human interaction (none, tactile, voice, and so forth).
4) *Takeaways*:
   - Which components of the system caused you the most trouble during the competition?
   - Did you evolve the demonstration after the hackathon? (Include references to publications that resulted from the demonstration, if any.)
   - What was the most important lesson from your participation in the hackathon?

Team Survey

Among the finalists, five were university teams, while one combined institutions (a university, a research institution, and two companies). As a condition to enter the hackathon, we limited the number of team members to five; however, the survey reported that the number of contributors was between two and 10. All teams had some combination of Ph.D. degree holders and M.Sc. students, and some included supervisors (postdocs/professors), technicians, and undergraduates. Of the participants, 15% were postdocs and professors, 45% were doctoral students, 35% were M.Sc. and undergraduate students, and 5% were technicians. Four of the six finalist teams had experience from other competitions, including the European Robotics League, RoboCup, European Robotics Challenge, Amazon Robotics Challenge, World Robot Summit, DJI Mobile Manipulation Challenge, and Nvidia Jetson Challenge. However, previous experience was not a guarantee of success, as one of the two winning teams reported having none.

Development Process

All finalists based their demonstration on previous work, either scientific (papers and Ph.D. theses) or technological (platform/software components for other competitions). Four teams had at least one mobile manipulation platform in their labs. The estimated time for preparing the hackathon demonstration strongly depended on previous experience, ranging from one to nine full person months. Note, however, that this effort estimate is merely indicative, as it was recalled after the competition.

Demonstration/System Description

The demonstrations in the final round mixed interactive and noninteractive executions. They were fully autonomous and required human intervention only for solving certain problems (e.g., objects that were out of reach, self-localization failures, and collisions). The three noninteractive demonstrations focused on completing tasks that required some sequence of object perception, manipulation, and navigation. Team Homer demonstrated the autonomous picking and sorting of cutlery after a party (the objects were randomly placed on a table) using semantic scene reasoning, as the objects were not easily identifiable through only depth information. A guarded motion was employed to grasp the cutlery by first touching the table in a pregrasp pose and then closing the robot's fingers around an object. Suitable checkpoints were provided to verify whether the grasp was successful. The object was then placed in a bowl located on a different table. The process was repeated until the first table was clean. Robotics.SG showed a reshelving application, as described previously. The PMM Tohoku team exhibited a liquid-pouring task, including the detection of a transparent bottle and a container. This was based on simple segmentation techniques, fitting a plane to a table, removing it, and then fitting cylinders to the remaining clusters of points (which represented the bottle and a cup).

The other teams required some interaction with humans. Team IRI presented a robot capable of feeding people with disabilities in a safe and delicate manner. The demonstration used an Amazon Alexa 3G interface to request commands, e.g., a choice of food, and human detection to find and interact with a person. The robot transported food and placed it on a table in front of the person. An arm-mounted camera enabled the robot to detect whether the human was interested in eating (when the human looked toward the camera), and, when appropriate, it retrieved food with a spoon. Then, if the robot detected that the person opened his or her mouth, it fed the patient. The process continued until the person indicated that no more food was required. After this, the robot removed the food from the table (and politely said goodbye).

Team TAMS showed a bartending application, as already described. Team IOC-AUDECO offered a drink-serving application. In this case, the robot perceived drinks available on a cluttered table, and a human chose a desired beverage by using a tablet or keyboard. Then, the robot calculated a manipulation sequence to retrieve the drink; the plan included moving cans that were in the way. A randomized physics-based motion planner introduced in [40] was used for this purpose. This planner permits robot–object and object–object interactions such that, when there is no collision-free path toward an item to be grasped, no explicit high-level task reasoning is required. However, possible complex multibody dynamical interactions are evaluated using a physics engine and considered in the expansion of a sampling-based planner. In particular, the planner enhances the state validity checker, the control sampler, and the tree exploration strategy of the Kinodynamic Motion Planning by Interior–Exterior Cell Exploration (KPIECE) process [41].

The teams mainly based their demonstrations on the hardware and sensors available on TIAGo. IRI additionally required a 6-DoF force torque sensor to guarantee safety. It also developed its own special 3D-printed gripper adapters for assuring an easy and stable grasp on the cutlery. Apart from these upgrades, the capabilities of TIAGo for carrying out collision-free navigation and arm motion planning were used. IOC-AUDECO employed a four-fingered Allegro hand instead of the default two-finger gripper to provide more-advanced grasping. TAMS needed an additional high-definition webcam, mounted on top of TIAGo, to get an image with enough resolution to detect the desired drink from the menu. Singapore.SG added a portable table to the robot to support the tray with the returned items. Additionally, it modified the shelves so that their lower part was perceived as a solid obstacle by the laser scanner used for navigation (otherwise, the shelf would have been missed, as the four legs were thin).

In terms of software, the developments were primarily based on ROS since all robot interfaces were tightly integrated with this framework. Simulations and visualizations generally employed Gazebo. All teams created specialized modules for certain tasks, and some relied on additional libraries. IRI used OpenFace (https://cmusatyalab.github.io/openface/) for face recognition and OpenPose (https://github.com/CMU-Perceptual-Computing-Lab/openpose) for person detection. IOC-AUDECO implemented planning in clutter by using the Kautham Project (https://sir.upc.edu/projects/kautham/). TAMS used the MoveIt Task Constructor (https://github.com/ros-planning/moveit_task_constructor), developed by team members and fully integrated in ROS, to define and plan actions consisting of multiple interdependent subtasks. Robotics.SG incorporated You Only Look Once (YOLO) (https://github.com/pjreddie/darknet/wiki/YOLO:-Real-Time-Object-Detection) for object perception, which was trained via images obtained with a self-built acquisition system [42]. Team Homer reused custom mapping and navigation tools (https://github.com/homer-robotics) developed for other competitions. They used the Mask Region-Based Convolutional Neural Network (Mask_RCNN) (https://github.com/matterport/Mask_RCNN) for object detection and segmentation, which, combined with planar surface segmentation, helped to detect the cutlery.

For control, most teams relied on the open-loop, position-based execution of planned sequences followed by a verification stage using TIAGo's sensors (joint encoders and vision) to decide if a plan was executed as intended. IRI employed a force-based control loop to govern the robotic arm while the feeding action was in progress. Team Homer integrated continuous current measurements into the grasping approach to detect contact with a table. Interestingly, no team used visual servoing techniques for controlling the manipulation actions. This indicates the focus on restricted scenarios that had quasi-static assumptions or that explicitly required human cooperation.

Takeaways

We asked the teams to identify the most troublesome components of their demonstration. Each team could identify any number of challenging areas; Figure 6 summarizes the responses. The most problematic area was object detection. Interestingly enough, in a survey performed with the participants of the Amazon Picking Challenge [7], perception was also identified as the most difficult aspect. Different techniques were employed by the teams for object detection and pose estimation: featured-based schemes, CAD models and surface textures, learning-based detection and estimation, and registration based on a fusion of depth and RGB data. In some cases, challenges came from the detection of transparent objects (bottles and glasses). Localization of the mobile base was ranked as the second most challenging area. To cope with it, for instance, IOC-AUDECO relied on Aruco markers to enhance the robustness of the table localization. Robotics.SG wrapped paper around the shelf legs to facilitate the mapping, navigation, and localization of the mobile base.

We were also interested in finding out if the experience gained during the hackathon was exploited afterward. Of the four teams that provided an answer to this question, three indicated that they evolved some of the components from the demonstration to create a more advanced lab version (IOC-AUDECO and TAMS) or reused solutions for a new competition (Homer). The TAMS demonstration, for instance, was transferred to a different platform, a PR 2 robot, thus showing the generality of the team's solution (https://www.youtube.com/watch?v=8S2MvKNbwmM). Three teams (IRI, IOC-AUDECO, and TAMS) indicated that some of the components were further developed and already published or submitted for publication as scientific papers. IRI has been able to transfer the knowledge gained from the force loop controller in the feeding task to a scenario involving bimanual cloth manipulation [43]. IOC-AUDECO continued to develop task and motion planning for mobile manipulation executions [44]. TAMS advanced object perception from the competition to detect and reconstruct transparent items [45].

Finally, we asked the teams to tell us the most important lesson they learned from the hackathon. IRI highlighted the need for further supervision during the demonstration execution. Its members reported that, as a lesson learned, their current exhibitions are now carefully designed to accommodate double-check control at different levels. In this line, IOC-AUDECO identified the need for more robust error detection and recovery strategies to resume tasks and recover from unexpected situations. TAMS highlighted the benefits of integrating independent components in a unified demonstration and recognized the need for intensively testing each component beforehand to avoid the more difficult job of debugging the overall execution. Team Homer appreciated the benefit of having on-site robotic platforms for implementing the demonstration outside its lab, thus reducing funding needs and transportation/insurance costs. Also, its members noted the benefits of having a common robotic platform for increasing the comparability of results across multiple research groups.

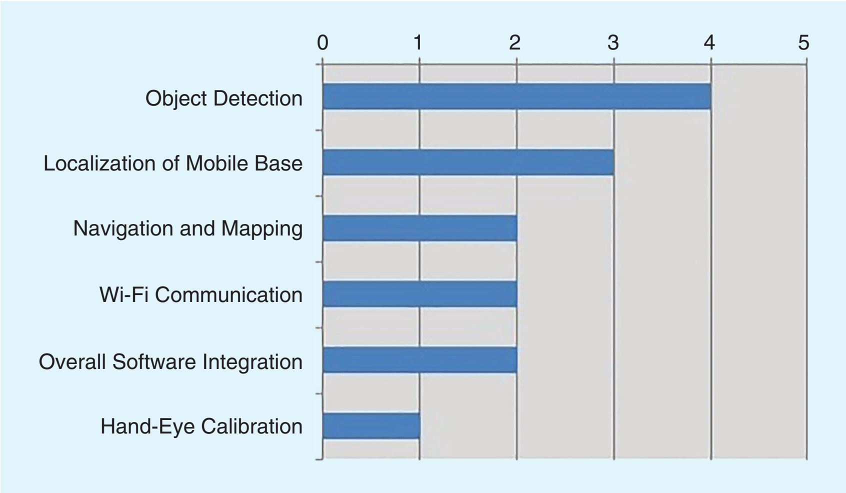

Figure 6. Challenging areas in the hackathon demonstrations.

Discussion and Outlook

In this section, we discuss the lessons learned after organizing the hackathon and the outlook for similar future events.

Hackathon Structure

The Mobile Manipulation Hackathon challenged the community to show integrated demonstrations that exploited the benefits of a mobile robotic manipulation platform. This required the development and integration of components at different levels, e.g., perception, navigation, and localization; grasp and manipulation planning; and human–robot interaction. Teams were free to propose a demonstration script, and they used this opportunity to showcase their latest developments in those fields, as opposed to a fixed task. This was a key difference of our hackathon when compared with competitions that have predetermined challenges. We believe both types of events are beneficial to the field: competitions with a fixed task provide a clearer picture of progress in a particular area, while competitions with an open structure, such as ours, are useful to understand a variety of possible applications. Therefore, we encourage the community (and we intend) to organize both types in the future.

Use of a Fixed Demonstrator Platform

The opportunity to use a unified hardware/software platform based on ROS provided a chance to compare multiple approaches. A solid software and simulation framework enabled teams to remotely develop their demonstration, thus reducing the time required for physical integration in the robotic system. However, we recognize that the basic tools for fast prototyping and quick debugging still need to be enhanced to facilitate the integration of full systems within a few days. In terms of the competition, it was greatly beneficial to have the robots on site, thus relieving teams of the burden of worrying about transportation costs, insurance, basic setups, and infrastructure and enabling them to focus on the pure development process.

From the perspective of robot manufacturers, the hackathon was a great opportunity to gather valuable feedback from both experienced and new users of the robots, which will help to improve how the next generation of devices will be conceived. The research community can also benefit from this kind of competitions to identify tools, libraries, and frameworks that could accelerate the implementation of real-world applications with complex robots, such as mobile manipulators. As an example of this, one prospect for mobile manipulators is the adoption, in the coming years, of ROS 2, which will provide better and more efficient data distribution among processes and support the coordination of multiple robots, security, and real-time control, among others.

Applications of Mobile Manipulation

Mobile manipulators are becoming increasingly available and have a huge potential to provide cost-effective solutions in different scenarios, including, for instance, industrial automation, manufacturing, logistics, health care, teleassistance, and crop harvesting. In many scenarios, robots will replace humans in dull, dirty, dangerous, and difficult tasks, such as bomb disposal and handling biological samples, as demanded now in times of pandemic. But, as we saw during the hackathon, a huge potential also lies in collaborative applications, where robots try to efficiently share their workspace and physically cooperate with humans in a delicate manner. Pouring liquid into a glass, serving a drink, and feeding a person are clear examples of this. More interesting and complex applications with autonomous bimanual, rigid, and deformable object manipulation tasks can be considered if more than one mobile manipulator is simultaneously used or if a dual-arm mobile manipulator is employed.

Further Technical Advances Required

Because mobile manipulators are complex systems that encompass different areas, they benefit from advances in fundamental topics, including perception, localization, and navigation as well as overall software integration and reliability, which we also identified as critical topics in our competition results (Figure 6). Some of the challenges are platform-dependent, including, for instance, robustness in communication (robust and reliable wireless communications are required), the integration of third-party hardware and software, and kinematics (e.g., simplicity to obtain a closed-form inverse-kinematics solution). On the other hand, some issues can be considered general mobile manipulation difficulties, including the following:

- *Localization*: precise location procedures within a robot's environment
- *Perception*: the robust identification of objects and estimation of their poses, using different sensors, including hand-held cameras for visual servoing purposes
- *Grasping*: the automatic determination of grasp configurations, taking a scene into account
- *Motion planning*: the capacity for planning collision-free motions as well as those that require contact to perform push actions
- *Task planning*: the automatic determination of the sequence of actions to perform a manipulation task, perhaps including regrasping actions, and the need to simultaneously consider motion planning
- *Reasoning*: the need for reasoning capabilities to understand a situation and accordingly tune all previously stated issues
- *Failure detection and recovery*: using reasoning capabilities for failure detection and the selection of recovery strategies.

If robots are to enter more complex scenarios, such as warehouses, grasping and manipulation must be greatly improved, as robots must show the ability to handle a huge variety of products in terms of size, weight, textures, and rigidity, all of which are located in different types of containers, bins, and shelves, especially in densely packed and cluttered facilities. This requires the further integration of tactile sensing and visual servoing and, in general, the fusion of multiple sensing modalities to enhance robot awareness.

Since the competition called for system-level demonstrations, successful execution depended on multiple components running simultaneously. Inevitably, failure rates multiply in such complex scenarios, and success requires a heightened awareness of failure sources and the handling of nonprototypical situations. In other words, platform reliability must be enhanced, and systems must be endowed with advanced error detection and recovery capabilities. The speed of execution is also a pending topic. During the hackathon, the robots took several minutes to perform actions that a human could do in a matter of seconds. Autonomy while working on battery power was not an issue, as the demonstrations were relatively short (fewer than 10 min for a full run), but it will be critical in real applications, where the robots must be available for extended periods of time. The proper exploitation of whole-body coordination to simultaneously employ a mobile base and a manipulator while performing a task is also required [46]. This has implications not only in terms of how to effectively use multiple DoFs and the redundancy of these platforms but also in terms of standardization and certification, which are essential to guarantee safety, especially in human–robot collaborations.

These issues (i.e., multimodal perception, manipulation planning and reasoning, system-level integration, speed of execution, and whole-body coordination) continue to be the main challenges in mobile manipulation systems, as observed during other recent competitions [10], [11], [13]. A more recent development is the introduction of competitions that focus on learning-based approaches, e.g., the Real Robot Challenge by the MPI-IS in 2020. This follows the general trend of merging robotics and artificial intelligence, but these competitions currently focus on manipulation-only tasks, as opposed to mobile manipulation. Far more challenging than fixed-based manipulation, mobile manipulation holds the potential to be a disruptive advance in robotics for applications at multiple levels, from industrial and home environments to health care. Open challenge hackathons/competitions targeting mobile manipulation would continue to serve the field and the community in the future.

Acknowledgments

The authors would like to thank the IROS 2018 Organizing Committee and, in particular, Prof. Carlos Balaguer, general chair, for supporting the development of this hackathon. We would also like to thank Prof. Jeannette Bohg, Dr. Graham Deacon, and Prof. Weiwei Wan for their great support serving as judges for the final competition. Additionally, we thank Ocado Technologies and the University of Leeds for sponsoring the awards and PAL Robotics for providing the TIAGo platforms and facilitating training at their premises. And, naturally, we recognize and greatly appreciate the interest and effort of all participating teams in the Mobile Manipulation Hackathon.

Mehmet Dogar received funding from the U.K. Engineering and Physical Sciences Research Council, under grant EP/P019560/1. Antonio Morales and the Universitat Jaume I Robotic Intelligence Laboratory were partially funded by the Ministerio de Economía y Competitividad (grant DPI2017-89910-R) and by Generalitat Valenciana (grant PROMETEO/2020/034). The work developed by IOC-AUDECO was partially supported by the Spanish government through the DPI2016-80077-R project. IRI was supported by the European Research Council through the European Union Horizon 2020 Program, under grant 741930 (Cloth Manipulation Learning From Demonstrations), and the Spanish State Research Agency through the María de Maeztu Seal of Excellence (grant MDM-2016-0656). TAMS was partially funded by the German Research Foundation and the National Science Foundation of China through the Crossmodal Learning project, grant TRR-169.

References

[1] G. Briscoe and C. Mulligan, "Digital innovation: The hackathon phenomenon," Creative Works London, London, U.K., 2014. [Online]. Available: http://www.creativeworkslondon.org.uk/wp-content/uploads/2013/11/Digital-Innovation-The-Hackathon-Phenomenon1.pdf

[2] E. H. Trainer, A. Kalyanasundaram, C. Chaihirunkarn, and J. D. Herbsleb, "How to hackathon: Socio-technical tradeoffs in brief, intensive collocation," in *Proc. ACM Conf. Computer-Supported Cooperative Work Social Comput.*, 2016, pp. 1118–1130. doi: 10.1145/2818048.2819946.

[3] H. Kienzler and C. Fontanesi, "Learning through inquiry: A global health hackathon," *Teach. Higher Educ.*, vol. 22, no. 2, pp. 129–142, 2017. doi: 10.1080/13562517.2016.1221805.

[4] J. R. Byrne, K. O'Sullivan, and K. Sullivan, "An IoT and wearable technology hackathon for promoting careers in computer science," *IEEE Trans. Educ.*, vol. 60, no. 1, pp. 50–58, 2016. doi: 10.1109/TE.2016.2626252.

[5] I. Nourbakhsh et al., "The winning robots from the 1993 robot competition," *AI Mag.*, vol. 14, no. 4, pp. 51–51, 1993. doi: 10.1609/aimag.v14i4.1068.

[6] S. Thrun et al., "Stanley: The robot that won the DARPA grand challenge," *J. Field Robot.*, vol. 23, no. 9, pp. 661–692, 2006. doi: 10.1002/rob.20147.

[7] N. Correll et al., "Analysis and observations from the first Amazon picking challenge," *IEEE Trans. Autom. Sci. Eng.*, vol. 15, no. 1, pp. 172–188, 2016. doi: 10.1109/TASE.2016.2600527.

[8] F. Von Drigalski et al., "Robots assembling machines: Learning from the World Robot Summit 2018 Assembly Challenge," *Adv. Robot.*, vol. 34, nos. 7–8, pp. 408–421, 2020.

[9] Y. Sun et al., "Robotic grasping and manipulation competition: Task pool," in *Robotic Grasping and Manipulation Challenge*, Y. Sun and J. Falco, Eds. Cham: Springer-Verlag, 2016, pp. 1–18. doi: 10.1007/978-3-319-94568-2.

[10] Z. Han et al., "Towards mobile multi-task manipulation in a confined and integrated environment with irregular objects," in *Proc. IEEE Int. Conf. Robot. Autom. (*ICRA), 2020, pp. 11,025–11,031. doi: 10.1109/ICRA40945.2020.9197395.

[11] J.-B. Yi and S.-J. Yi, "Mobile manipulation for the HSR intelligent home service robot," in *Proc. Int. Conf. Ubiquitous Robots (UR)*, 2019, pp. 169–173. doi: 10.1109/URAI.2019.8768782.

[12] M. Matamoros et al. "Robocup@home 2019: Rules and regulations." RoboCup. https://athome.robocup.org/wp-content/uploads/Rulebook2019GermanOpen.pdf (accessed Mar. 10, 2020).

[13] J. Huth, "Taking robots shopping," *Nature Mach. Intell.*, vol. 1, no. 11, pp. 545–545, 2019. doi: 10.1038/s42256-019-0118-0.

[14] R. Awad, F. Caccavale, and A. J. van der Meer, "Evaluation and selection activities in EuRoC: Innovations and lessons learned," in *Bringing Innovative Robotic Technologies from Research Labs to Industrial End-users*, F. Caccavale, C. Ott, B. Winkler, and Z. Taylor, Eds. Cham: Springer-Verlag, 2020, pp. 15–34. doi: 10.1007/978-3-030-34507-5_2.

[15] F. Ruggiero, V. Lippiello, and A. Ollero, "Aerial manipulation: A literature review," *IEEE Robot. Automat. Lett.*, vol. 3, no. 3, pp. 1957–1964, 2018. doi: 10.1109/LRA.2018.2808541.

[16] M. A. Roa, K. Nottensteiner, A. Wedler, and G. Grunwald, "Robotic technologies for in-space assembly operations," in *Proc. Symp. Adv. Space Technol. Robot. Automat. (ASTRA)*, 2017. [Online]. Available: https://robotics.estec.esa.int/ASTRA/Astra2017/1.%20Wednesday%2021%20June/5B%20Orbital%20Robotics/S.B5_14.50_Roa.pdf

[17] A. Gomez-Tamm, V. Perez-Sanchez, B. Arrue, and A. Ollero, "SMA actuated low-weight bio-inspired claws for grasping and perching using flapping wing aerial systems," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, 2020, pp. 8807–8814. doi: 10.1109/IROS45743.2020.9341741.

[18] S. Hamaza, I. Georgilas, G. Heredia, A. Ollero, and T. Richardson, "Design, modeling, and control of an aerial manipulator for placement and retrieval of sensors in the environment," *J. Field Robot.*, vol. 37, no. 7, pp. 1224–1245, 2020. doi: 10.1002/rob.21963.

[19] A. Ollero et al., "The AEROARMS project: Aerial robots with advanced manipulation capabilities for inspection and maintenance," *IEEE Robot. Automat. Mag.*, vol. 25, no. 4, pp. 12–23, 2018. doi: 10.1109/MRA.2018.2852789.

[20] M. Zhao, F. Shi, T. Anzai, K. Okada, and M. Inaba, "Online motion planning for deforming maneuvering and manipulation by multilinked aerial robot based on differential kinematics," *IEEE Robot. Automat. Lett.*, vol. 5, no. 2, pp. 1602–1609, 2020. doi: 10.1109/LRA.2020.2967285.

[21] D. Youakim, P. Ridao, N. Palomeras, F. Spadafora, D. Ribas, and M. Muzzupappa, "MoveIt!: Autonomous underwater free-floating manipulation," *IEEE Robot. Automat. Mag.*, vol. 24, no. 3, pp. 41–51, 2017. doi: 10.1109/MRA.2016.2636369.

[22] S. Sivcec, J. Coleman, E. Omerdic, G. Dooley, and D. Toal, "Underwater manipulators: A review," *Ocean Eng.*, vol. 163, no. 1, pp. 431–450, 2018. doi: 10.1016/j.oceaneng.2018.06.018.

[23] A. Kheddar et al., "Humanoid robots in aircraft manufacturing: The Airbus use cases," *IEEE Robot. Autom. Mag.*, vol. 26, no. 4, pp. 30–45, 2019. doi: 10.1109/MRA.2019.2943395.

[24] E. Krotkov et al., "The DARPA robotics challenge finals: Results and perspectives," *J. Field Robot.*, vol. 34, no. 2, pp. 229–240, 2017. doi: 10.1002/rob.21683.

[25] J. Schuler, *Integration von Förder-und Handhabungseinrichtungen*, vol. 104. New York: Springer-Verlag, 2013.

[26] C. Weisbin, B. Burks, J. Einstein, R. Feezell, W. Manges, and D. Thompson, "HERMIES-III: A step toward autonomous mobility manipulation and perception," *Robotica*, vol. 8, no. 1, pp. 7–12, 1990. doi: 10.1017/S0263574700007268.

[27] T. Lueth, U. Nassal, and U. Rembold, "Reliability and integrated capabilities of locomotion and manipulation for autonomous robot assembly," *J. Robot. Auton. Syst.*, vol. 14, nos. 2–3, pp. 185–198, 1995. doi: 10.1016/0921-8890(94)00029-2.

[28] J. Cameron, D. MacKenzie, K. Ward, R. Arkin, and W. Book, "Reactive control for mobile manipulation," in *Proc. IEEE Int. Conf. Robot. Automat. (ICRA)*, 1993, pp. 228–235. doi: 10.1109/ROBOT.1993.292181.

[29] O. Khatib, K. Yokoi, K. Chang, D. Ruspini, R. Holmberg, A. Casal, and A. Baader, "Force strategies for cooperative tasks in multiple mobile manipulation systems," in *Proc. Int. Symp. Robot. Res.*, 1995, pp. 333–342. doi: 10.1007/978-1-4471-1021-7_37.

[30] M. Hvilshøj, S. Bøgh, O. S. Nielsen, and O. Madsen, "Autonomous industrial mobile manipulation (AIMM): Past, present and future," *Ind. Robot., An Int. J.*, vol. 39, no. 2, pp. 120–135, 2012. doi: 10.1108/01439911211201582.

[31] J. Kuehnle et al., "6D object localization and obstacle detection for collision-free manipulation with a mobile service robot," in *Proc. IEEE Int. Conf. Adv. Robot.*, 2009, pp. 1–6.

[32] J. Bohren et al., "Towards autonomous robotic butlers: Lessons learned with the PR2," in *Proc. IEEE Int. Conf. Robot. Automat. (ICRA)*, 2011, pp. 5568–5575. doi: 10.1109/ICRA.2011.5980058.

[33] M. A. Diftler et al., "Robonaut 2-the first humanoid robot in space," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2011, pp. 2178–2183. doi: 10.1109/ICRA.2011.5979830.

[34] B. Graf, U. Reiser, M. Hägele, K. Mauz, and P. Klein, "Robotic home assistant Care-O-bot 3—product vision and innovation platform," in *Proc. IEEE Workshop Adv. Robot. Social Impacts*, 2009, pp. 139–144. doi: 10.1109/ARSO.2009.5587059.

[35] S. S. Srinivasa et al., "HERB: A home exploring robotic butler," *Auton. Robots*, vol. 28, no. 1, pp. 5–20, 2010. doi: 10.1007/s10514-009-9160-9.

[36] C. Borst et al., "Rollin'Justin-mobile platform with variable base," in *Proc. IEEE Int. Conf. Robot. Automat. (ICRA)*, 2009, pp. 1597–1598. doi: 10.1109/ROBOT.2009.5152586.

[37] T. Asfour et al., "ARMAR-III: An integrated humanoid platform for sensory-motor control," in *Proc. IEEE-RAS Int. Conf. Humanoid Robots*, 2006, pp. 169–175. doi: 10.1109/ICHR.2006.321380.

[38] M. Quigley et al., "ROS: An open-source Robot Operating System," in *Proc. ICRA Workshop Open Source Softw.*, 2009, p. 5.

[39] "World Robotics 2019 Service Robots," International Federation of Robotics, Frankfurt, 2019.

[40] Muhayyuddin, M. Moll, L. Kavraki, and J. Rosell, "Randomized physics-based motion planning for grasping in cluttered and uncertain environments," *IEEE Robot. Automat. Lett.*, vol. 3, no. 2, pp. 712–719, 2018. doi: 10.1109/LRA.2017.2783445.

[41] I. Sucan and L. E. Kavraki, "A sampling-based tree planner for systems with complex dynamics," *IEEE Trans. Robot.*, vol. 28, no. 1, pp. 116–131, 2012. doi: 10.1109/TRO.2011.2160466.

[42] Z. Chong et al., "An innovative robotic stowing strategy for inventory replenishment in automated storage and retrieval system," in *Proc. Int. Conf. Control, Automat., Robot. Vision (ICARCV)*, 2018, pp. 305–310. doi: 10.1109/ICARCV.2018.8581338.

[43] I. Garcia-Camacho et al., "Benchmarking bimanual cloth manipulation," *IEEE Robot. Automat. Lett.*, vol. 5, no. 2, pp. 1111–1118, 2020. doi: 10.1109/LRA.2020.2965891.

[44] S. Saoji and J. Rosell, "Flexibly configuring task and motion planning problems for mobile manipulators," in *Proc. Annu. Conf. Emerg. Technol. Factory Automat. (ETFA)*, 2020, pp. 1285–1288. doi: 10.1109/ETFA46521.2020.9212086.

[45] P. Ruppel, M. Görner, N. Hendrich, and J. Zhang, "Detection and reconstruction of transparent objects with infrared projection-based RGB-D cameras," in *Proc. Int. Conf. Cognit. Syst. Inform. Process. (ICCSIP)*, 2020.

[46] K. Harada and M. A. Roa, "Manipulation and task execution by humanoids," in *Humanoid Robotics: A Reference*, A. Goswami and P. Vadakkepat, Eds. Dordrecht: Springer-Verlag, 2019, pp. 1633–1655. doi: 10.1007/978-94-007-6046-2_31.

Máximo A. Roa, Institute of Robotics and Mechatronics, German Aerospace Center, Wessling, 82234, Germany. Email: maximo.roa@dlr.de.

Mehmet Dogar, University of Leeds, Leeds, LS2 9JT, U.K. Email: m.r.dogar@leeds.ac.uk.

Jordi Pages, PAL Robotics, Barcelona, 08005, Spain. Email: jordi.pages@pal-robotics.com.

Carlos Vivas, PAL Robotics, Barcelona, 08005, Spain. Email: carlos.vivas@pal-robotics.com.

Antonio Morales, Robotic Intelligence Laboratory, Universitat Jaume I, Castellon, 12071, Spain. Email: morales@uji.es.

Nikolaus Correll, Department of Computer Science, University of Colorado Boulder, Boulder, Colorado, 80309, USA. Email: nikolaus.correll@colorado.edu.

Michael Görner, University of Hamburg, Hamburg, 22527, Germany. Email: goerner@informatik.uni-hamburg.de.

Jan Rosell, Institut d'Organització i Control, Universitat Politècnica de Catalunya, Barcelona, 08028, Spain. Email: jan.rosell@upc.edu.

Sergi Foix, Institut de Robòtica i Informàtica Industrial, Universitat Politècnica de Barcelona, Barcelona, 08028, Spain. Email: sfoix@iri.upc.edu.

Raphael Memmesheimer, Computational Visualistics, University of Koblenz-Landau, Koblenz, 56073, Germany. Email: raphael@uni-koblenz.de.

Francesco Ferro, Pal Robotics, Barcelona, 08005, Spain. Email: francesco.ferro@pal-robotics.com.