# The UMass Mobile Manipulator UMan:
# An Experimental Platform for Autonomous Mobile Manipulation

Dov Katz, Emily Horrell, Yuandong Yang, Brendan Burns, Thomas Buckley
Anna Grishkan, Volodymyr Zhylkovskyy, Oliver Brock, Erik Learned-Miller
Department of Computer Science
University of Massachusetts Amherst

*Abstract*— **Research in Autonomous Mobile Manipulation critically depends on the availability of adequate experimental platforms. In this paper, we describe an ongoing effort at the University of Massachusetts Amherst to construct a hardware platform with redundant kinematic degrees of freedom, a comprehensive sensor suite, and significant end-effector capabilities for manipulation. In our research, we pursue an end-effector centric view of autonomous mobile manipulation. In support of this view, we are developing a comprehensive software suite to provide a high level of competency in robot control and perception. This software suite is based on a multi-objective, task-level motion control framework. We use this control framework to integrate a variety of motion capabilities, including task-based force or position control of the end-effector, collision-free global motion for the entire mobile manipulator, and mapping and navigation for the mobile base. We also discuss our efforts in developing perception capabilities targeted to problems in autonomous mobile manipulation. Preliminary experiments on our UMass Mobile Manipulator (UMan) are presented.**

## I. INTRODUCTION

Autonomous robots are beginning to address real-world tasks in unstructured and dynamic environments. Today these robots predominantly perform tasks based on mobility. However, the potential of augmenting autonomous mobility with dexterous manipulation skills is significant and has numerous important applications, ranging from in-orbit satellite servicing to elderly care and flexible manufacturing [3]. The successful deployment of autonomous robots that combine manipulation skills with mobility, however, still requires substantial scientific advances in a variety of research areas, including autonomous manipulation in unstructured environments, perception, and system integration. In support of research in these areas, we are developing and constructing the mobile manipulation platform *UMan* (UMass Mobile Manipulator). In this paper, we discuss the objectives for building such a platform, describe our approach of combining software and hardware to provide an adequate level of competency for research, and report on progress with preliminary research initiatives.

A platform adequate for research in autonomous mobile manipulation has to combine mobility with dexterous manipulation. A manipulator with a dexterous end-effector allows complex interactions with objects in the environment. And mobility extends the workspace of the manipulator, posing new challenges by permitting the robot to operate in unstructured environments. In such environments it is impossible to anticipate all scenarios and to pre-program the behavior of the robot. Manipulation in combination with mobility thus requires algorithmic approaches that are versatile and general. The variability and complexity of unstructured environments also require algorithmic approaches that permit the robot to continuously improve its skills from its interactions with the environment. Furthermore, the combination of mobility and manipulation poses new challenges for perception as well as the integration of skills and behaviors over a wide range of spatial and temporal scales.

Fig. 1. UMan, the UMass Mobile Manipulator

The mobile manipulation platform *UMan*, shown in Figure 1, has been designed to support our research in algorithms and control for autonomous mobile manipulation. Due to the focus on dexterous manipulation, *UMan*'s ability to perform physical work in its environment was of particular importance during the design process. Our objective was to build a hardware platform with redundant kinematic degrees of freedom, a comprehensive sensor suite, and significant end-effector capabilities for manipulation. Since the focus of our research is on algorithms and not on hardware design, we chose to use mostly off-the-shelf components in *UMan*'s construction.

*UMan* differs from related robotic platforms (see Section II) because of our end-effector centric, integrated view of hardware and software infrastructure. This combination of hardware and software results in an experimental platform that permits researchers to focus solely on the specification of end-effector behavior without having to worry about the motion of the remaining degrees of freedom. Furthermore, *UMan* is one of the few platforms currently available that combines mobility with dexterous manipulation capabilities. Most comparable platforms are limited in either dexterity or mobility. A more detailed review of related hardware platforms will be presented in the next section.

## II. RELATED PLATFORMS

Most current robotic platforms focus on one of two complementary aspects of autonomous mobile manipulation: bi-manual dexterous manipulation or bipedal mobility.

Bi-manual robots consist of a torso, two arms and dexterous hands. Examples from this category include UMass's Dexter [8], a bi-manual robot consisting of two commercial seven degree-of-freedom arms (Barrett's WAM [24]) and three-fingered hands. Dexter's arms are stiff with negligible inherent compliance, requiring active control of compliance using accurate force sensing. Its hands have four degrees of freedom each, and can afford various dexterous tasks. In contrast to Dexter, MIT's bi-manual robot Domo [6] employs series elastic actuators in its two arms, providing inherent compliance for safe interactions with objects in the environment. NASA's Robonaut [2], a humanoid torso designed to assist or replace astronauts, closely imitates the kinematics of human arms and hands. Robonaut has performed complex dextrous manipulation tasks, resembling tasks performed by astronauts in space. Other examples of a bi-manual manipulation platform include Clara at the Technical University Munich and Justin at DLR in Germany. These stationary, bi-manual platforms are designed for dexterous manipulation and benefit from having two arms that can cooperate in manipulation tasks. However, these platforms have limited workspace and cannot be deployed in large unstructured environments.

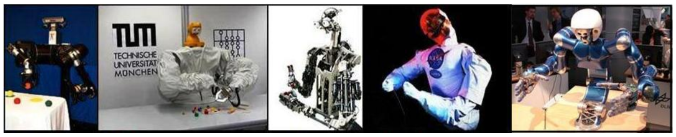

Fig. 2. Bi-manual robots: Dexter (UMass Amherst), Clara (Technical University Munich), Domo (MIT), Robonaut (JSC, NASA), Justin (DLR, Germany)

The development of legged robots, such as Honda's ASIMO [17], has focused on issues of bipedal locomotion. Legs seem to be better suited than wheels for locomotion in human environments. Due to the focus on locomotion, the manipulation capabilities of these platforms are often limited. For example, ASIMO's hands only possess a single degree of freedom and are not able to perform complex manipulation tasks. Other examples are Sony's QRIO [19], Kwada Industries' HRP2 [12], Waseda University Tokyo's WABIAN RIII [26], and the Toyota Partner Robot [25]. All of these platforms have brought about significant advances in bipedal locomotion and mechanical design, but have had limited impact in the area of autonomous mobile manipulation.

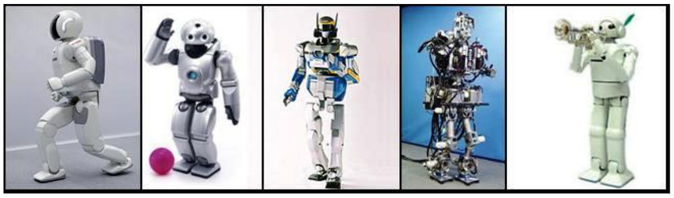

Fig. 3. Examples of legged humanoid robots: Asimo (Honda), QRIO (Sony), HRP-3 (Kawada Industries), WABIAN RIII (Waseda University Tokyo), Toyota Partner Robot

Ideally, the advances in bi-manual manipulation and locomotion could be combined to provide competent experimental platforms for autonomous mobile manipulation. Such platforms could navigate in complex, human indoor and outdoor environments to perform dexterous manipulation tasks. Unfortunately, such integration is not straight-forward. Most commercial manipulators are heavy, large, and power hungry. Consequently, highly mobile robots, such as legged mobility platforms, cannot easily accommodate dexterous manipulators. To overcome this challenge, most existing mobile manipulation platforms rely on less flexible mobility solutions. One of the earliest mobile manipulators, the Stanford Assistant Mobile Manipulation (SAMM) [10] consists of a holonomic Nomadic XR4000 base and a PUMA 560 manipulator arm, equipped with a parallel-jaw gripper. The mobile manipulation platform ARMAR [21] at the University of Karlsruhe consists of two arms mounted on a wheeled mobile base, two parallel-jaw grippers, and a stereo vision head. The Stanford AI Robot (STAIR) uses a modified Segway RMP in tractor mode for mobility; a five degree-of-freedom manipulator with a parallel-jaw gripper provides manipulation capabilities. Other examples of mobile manipulators can be found at Örebro University (PANDI1), at the Centre for Autonomous Systems at the Royal Institute of Technology, and at the Robotics and Automation Laboratory at Michigan State University.

## III. HARDWARE PLATFORM

*UMan* consists of a holonomic mobile base with three degrees of freedom, a seven degree-of-freedom manipulator arm, and a four degree-of-freedom hand.

Our main objective was to support dexterous manipulation on a mobile platform. To achieve this objective for a single, integrated system, it was not possible to optimize every aspect of the platform; instead, we had to carefully balance different constraints. We strove to provide adequate end-effector capabilities for a wide range of dexterous manipulation tasks. We considered mobility as additional degrees of freedom in service to manipulation, rather than as an objective itself. We therefore choose a mode of mobility that maximally supports manipulation without imposing additional constraints.

Fig. 4. Examples of mobile manipulators: ARMAR (University of Karlsruhe), PANDI1 (Örebro University), SAMM (Stanford University)

At the same time, we focused on the use of commercially available components. This facilitates recreating our results, as well as encourages standardization in robotics hardware. We justify this choice with the following observation: NASA researchers at JSC were able to perform dexterous manipulation tasks by teleoperating Robonaut. In particular, they teleoperated Robonaut to perform complex maintenance tasks on a model of the Hubble space telescope [1]. These experiments demonstrate that existing hardware is in principle capable of performing complex dexterous manipulation tasks. Consequently, we believe that algorithmic aspects currently represent the most important challenge in autonomous mobile manipulation.

### A. Actuation

*UMan*'s mobility is provided by a modified Nomadic XR4000 mobile base. Its four casters are dynamically decoupled [5] to provide holonomic motion, which facilitates a unified control scheme for degrees of freedom associated with mobility and manipulation. This unified control scheme leverages all degrees of freedom to be able to exploit kinematic redundancies when performing multi-objective behavior (see Section IV). Other modes of mobility, such as legged locomotion or dynamically stable mobility platforms (Segway RMP), result in additional challenges pertaining to the coordination of mobility and manipulation. While these challenges represent interesting research problems by themselves, our objective was to build a hardware platform that maximally facilitates research in autonomous manipulation. Therefore, we chose a holonomic mobility platform that supports the coordination of mobility and manipulation.

The XR4000 mobile platform was specifically designed for mobile manipulation. Its power system allows untethered operation of *UMan* for several hours. The wheels and the frame of the mobile base are designed to be very stiff and damped so that accurate position control of the mobile manipulator is feasible. This advantageous property reduces the ability to navigate on uneven terrain, a cost we were willing to pay since we do not think it fundamentally limits our ability to perform our research. Finally, the base is sized to be able to contain adequate computational resources (see Section III-B) and sensors (see Section III-C).

A Barrett Technologies Whole Arm Manipulator (WAM) with seven anthropomorphic degrees of freedom (three in the shoulder, one in the elbow, three in the wrist, see figure 5) together with the three-fingered Barrett hand provide *UMan*'s dexterous manipulation capabilities. The new lightweight design of the WAM has very low power consumption and is thus well-suited for mobile applications. All electronics for the control of the arm are built into the arm itself, facilitating its integration with a mobile platform. The WAM provides good dynamic performance and torque sensing in its joints. It is thus capable of using all of its links to perform force-controlled manipulation tasks.

Fig. 5. Barrett Technologies Whole Arm Manipulator (WAM)

The three-fingered Barrett hand (see figure 6) can flex any of its three-link fingers individually. A fourth degree of freedom in the hand permits switching between an enveloping grasp to grasp with an opposing thumb. While this hand is clearly less dexterous than the human hand, it provides significantly more dexterity than the parallel-jaw grippers that can be found on most mobile manipulators today. Robonaut mounted on a Segway RMP represents a notable exception.

Fig. 6. Barrett Technologies three-fingered hand

### B. Computation

*UMan*'s mobile base houses two single-board PCs with Pentium 4 2.4GHz CPUs (see figure 7). One of these PCs is dedicated to real-time control of the base and the manipulator arm. It is running the real-time operating system QNX [16].

The other PC is running Linux and is dedicated to higher-level tasks, such as vision processing, navigation, motion planning, and task planning. Both computers are connected via a dedicated wired Ethernet link. Wireless Ethernet connects both computers to the off-board computing resources. Should additional on-board resources become necessary, we could place laptops inside the cargo bays of the XR4000 mobile base.

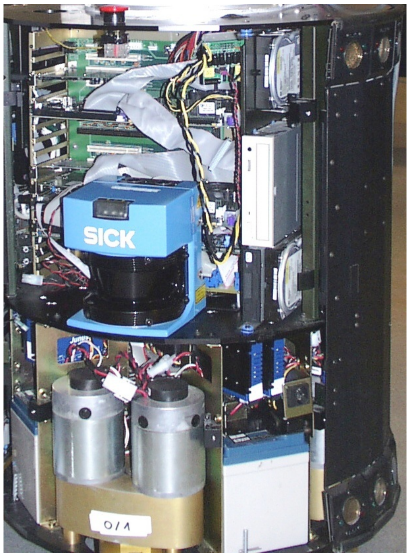

Fig. 7. UMan's interior

### C. Sensing

We view sensing as one of the most critical and under-explored aspects in autonomous mobile manipulation. Consequently, we are equipping UMan with a rich suite of sensors.

*UMan* can navigate in the environment using an onboard SICK LMS200 laser range finder. The SICK's field of view is 180 degrees, with an angular resolution of 1 degree and a distance resolution of 10mm with a range of 80m. A 180 degree scan of the environment can be completed 75 times per second.

Visual input is received from a Sony XCD710CR color camera mounted on the wrist. The camera has an IEEE-1394 (Firewire) interface, and was chosen for its superior image quality. It can produce 30 frames per second at a resolution of 1024 by 768 pixels, and has many features (i.e. partial scans, various resolutions and color encodings). By controlling the position and orientation of its arm, *UMan* can collect visual data and generate 3D images far beyond the capabilities of a simple stationary pan-tilt vision system. Despite the complication in linking the vision system and the manipulator, this combination is able to look behind obstructions in the field of view and can easily generate multiple views of the same object. In the future, more cameras will be installed to further enhance *UMan*'s visual perception.

We plan to mount an additional laser range finder (Hokuyo's URG-04LX [9]) on *UMan*'s wrist. Hokuyo's sensor is relatively light (160g), compact (L50xW50xH70mm), and has low power requirements (2.5W). Those characteristics make it suitable to be carried on a robotic arm. Moreover, the URG-04LX has a 240 degree field of view, and provides an accuracy of 10mm with an angular resolution of 0.36 degrees. Complementing our vision sensors with a movable laser range finder will allow us to perform sensor fusion of complementary information about the environment.

Tactile sensing is the heart of dexterous manipulation. Unfortunately, robotic arms with inherent force compliance are not commercially available. The lack of force compliance makes the safe operation of the manipulator arm more difficult and calls for a more complex control architecture. Manipulator arms with inherent force compliance exist. However, since our research focuses on algorithmic aspects, we rely on off-the-shelf hardware and use two sets of force sensors in conjunction with precise position feedback from the arm/hand to achieve compliant control. The first set of sensors, mounted on the fingertips of the Barrett hand, consists of three ATI Nano17 6-axis force/torque sensors. The ATI Nano17 has a resolution of 1/160N, for forces varying between -12N to +12N. For torques varying between -120Nmm to +120Nmm, it has a resolution of 1/32Nmm. The second sensor, mounted on the wrist, is an ATI Gamma 6-axis force/torque sensor. This sensor has a resolution of 1/640N for forces varying between -32N to +32N. For torques varying between -2.5Nm to +2.5Nm, it has a resolution of 1/1000Nm.

*UMan*'s wireless network adapter (see Section III-B) allows it to benefit from off-board sensing to complement its onboard sensors. When operating in sensor-enabled environments, *UMan* can collect data from various external sources and use it to increase its knowledge and understanding of the environment in which it operates.

## IV. MULTI-OBJECTIVE TASK-LEVEL CONTROL

Our research in autonomous mobile manipulation is primarily concerned with interactions between *UMan*'s end-effector and the environment. To facilitate this research, we chose an end-effector-centric control framework. The operational space framework for robot control [13] allows us to define manipulation tasks in terms of end-effector motion and forces, thus abstracting from the kinematics and dynamics of the underlying robot mechanism. At the same time, it does not prevent us from exploiting the mechanism's kinematic and dynamic characteristics through other methods of control to achieve desired performance metrics. Moreover, this framework allows treating several different parts of the mechanism as end-effector, and prioritizing the tasks addressed by each end-effector. In this section we describe the task-level control scheme implemented on *UMan* and our method of combining several behaviors to exploit *UMan*'s redundancy.

Task-level control [13] is a convenient and powerful method of generating multi-objective behavior for robotic systems. Rather than specifying joint trajectories, this framework permits direct control of the manipulator's end-effectors, greatly facilitating programming for kinematically redundant robots, such as *UMan*. Task-level control also permits the task-consistent execution of subordinate behaviors, exploiting nullspace projections. Given an end-effector task, expressed as a force/torque vector $F_{\text{task}}$ acting on the end-effector, and given an arbitrary subordinate behavior, expressed as a vector of joint torques $\Gamma_0$, we can determine the torque $\Gamma$ to achieve task and subordinate behavior as follows:

$$\Gamma = J_{\text{task}}(q)\, F^T_{\text{task}} + N^T_{\text{task}}(q)\, \Gamma_0, \tag{1}$$

where $N^T_{\text{task}}$ represents a projection into the nullspace of the end-effector Jacobian $J_{\text{task}}$. This projection ensures that the subordinate behavior will not alter task behavior, i.e., it will result in task-consistent motion.

This principle of nullspace projections can be extended to cascade an arbitrary number of hierarchical behaviors [18]. If behavior $i$ results in torque $\Gamma_i$, the torque

$$\Gamma = \Gamma_1 + N^T_1(q) \quad \left(\Gamma_2 + N^T_2(q) \quad \left(\Gamma_3 + N^T_3(q) \quad (\ldots)\right)\right) \tag{2}$$

combines these behaviors in such a way that behavior $i$ does not affect behavior $j$ if $i > j$. In equation 2, $N^T_i$ is the nullspace projection associated with the task Jacobian of behavior $i$. Here, we adopt the more compact notation of the control basis [11] to describe such cascaded nullspaces. We associate a control primitive $\phi_i$ with each torque $\Gamma_i$. If a control primitive $\phi_i$ is executed in the nullspace of the control primitive $\phi_j$, we say that $\phi_i$ is performed subject to $\phi_j$, written as $\phi_i \lhd \phi_j$. We can now rewrite equation 2 as

$$\ldots \lhd \phi_3 \lhd \phi_2 \lhd \phi_1 . \tag{3}$$

This task-level framework with nullspace projections serves as the underlying control scheme for *UMan*. Using this approach, other research groups have already demonstrated sophisticated multi-objective behavior on simulated humanoid robots [18]. Our research effort based on this control framework will focus on the control primitives themselves as well as on their automatic composition and sequencing to achieve robust, autonomous manipulation behavior in unstructured and changing environments.

## V. GLOBAL PLANNING AND CONTROL OF ROBOT MOTION

The task-level control framework described in the previous section is well-suited for the specification of task-based forces and position constraints at the end-effector. However, in the context of autonomous mobile manipulation, end-effector tasks have to be performed in dynamic and unstructured environments. In such environments, obstacles can interfere with task execution. The task-level control framework is thus not sufficient to provide a complete abstraction for the end-effector. In addition to the control-based framework described in the previous section, our platform also has to include the ability to perform globally task-consistent motion for manipulation tasks, i.e., motion that fulfills the position and force constraints imposed by the task as well as the global motion constraints imposed by obstacles in the environment.

We have developed the elastic roadmap framework [27] to augment the task-level control framework from Section IV with the ability to perform globally task-consistent and collision-free motion. The elastic roadmap framework translates global workspace connectivity information into a series of potential functions. These potential functions collectively represent an approximated navigation function for the robot's task. The navigation function in turn represents a global motion plan that is capable of responding to feedback from the environment. The motion plan integrates seamlessly with our task-level control framework. By combining the global plan with the control primitives that generate task-behavior, task constraints as well as global obstacle avoidance constraints can be maintained throughout the execution of a task, even in unstructured and dynamic environments.

More detail about the elastic roadmap framework is provided in our RSS 2006 paper [27]. In Section VI-B, we show the simulation results from that paper for motion generation based on the elastic roadmap framework. We are currently implementing the elastic roadmap framework on *UMan* and plan to duplicate the simulation results on the real platform. Moving from simulation to the real world introduces new difficulties. Physical sensors, for once, are bound to be limited and inaccurate. Consequently, addressing the integration is an important aspect of motion generation in this application domain.

The multi-objective control scheme described in Section IV also facilitates the inclusion of readily available software for mobile robot navigation. We choose to integrate the SLAM package CARMEN [4], which can build a map of the robot's environment using laser scans, and can navigate the robot to a desired position while avoiding obstacles [23], [22].

## VI. PRELIMINARY RESULTS

### A. Hardware Integration

One of the most important challenges in constructing an experimental platform for autonomous mobile manipulation is the integration of heterogeneous hardware and software components [3]. In this section we describe *UMan's* evolution during the hardware integration phase.

*UMan* is constructed from off-the-shelf components. The first major milestone for platform integration was the simultaneous real-time control of the XR4000 mobile base and the Barrett WAM with the system running on battery power. For this test, we had *UMan* follow a circular trajectory on the floor while moving its arm through its range of motion. The circular trajectory is performed by the base without changing orientation, demonstrating its holonomic motion capabilities (see http://www.cs.umass.edu/∼dubik for a video). We have implemented joint space control and operational space control at rates up to 1KHz. Operational space control will permit the exploitation of the mechanisms kinematic redundancy using multi-objective control based on nullspace composition (see Section IV).

In the next step, we began the integration of the sensor suite and in particular the linking of the robot arm with our vision camera to perform visual servoing (see Figure 8). The camera (mounted on the wrist) provides color images to a simple color tracking algorithm. The algorithm requested arm movements (and therefore camera movements) to center a blue ball on the camera's image plane. The control code was executed in at high frequency on the real-time PC, while the vision processing was executed at a much lower frequency on the Linux-based PC. This simple test will be the basis of future visual servoing tasks. As discussed in section IV, this task will be performed in parallel with other tasks, by taking advantage of *UMan*'s kinematic redundancy.

Fig. 8. Visual servoing: Sony XCD710CR camera mounted on UMan's wrist, UMan's arm tracking the blue ball in an attempt to center it in the image plane, the velocity vector commanded by the vision system

### B. Elastic Roadmaps

The elastic roadmap framework has been implemented and tested in simulation [27]. Here, we report only on the high-level results to illustrate how the elastic roadmap framework is capable of supporting our end-effector-centric view in the context of autonomous mobile manipulation.

Most work in global motion planning only considers end-effector placement tasks. These tasks are specified by an initial and a final configuration for the entire manipulator. Constraints that need to be maintained *during* the motion itself are rarely considered. In our end-effector-centric view of autonomous mobile manipulation, it oftentimes is sufficient to consider only the end-effector placement, irrespective of the manipulator configuration that achieves that placement. In addition, a large number of tasks in autonomous mobile manipulation require that end-effector constraints be maintained throughout the entire motion. All of these requirements are met by the elastic roadmap framework. In all three experiments shown in Figure 9, the robot is performing a task that consists of end-effector placement in conjunction with task-constraints that need to be maintained throughout the placement motion. The experiments show that the elastic roadmap framework is able to achieve global real-time planning and replanning for our mobile manipulation platform, even in dynamic environments.

## VII. FUTURE RESEARCH

Autonomous mobile manipulation is a relatively young discipline within robotics. It combines a wide variety of research areas, ranging from force control to mechanism design to computer vision. A workshop on autonomous mobile manipulation brought together researchers from the field to identify a roadmap for research in this area [3]. Here, we briefly describe some of the areas where we will focus our efforts.

### A. Perception

Vision systems have been used to augment the abilities of both manipulators and mobile platforms since their inception. However, there are many fundamental capabilities which could benefit dramatically from a bidirectional flow of information between vision and manipulation, and which are still relatively unexplored.

Visually analyzing the environment is a precondition for successful manipulation. More surprising perhaps is the effect that manipulation can have on visual perception. For example, the difficulty of recognizing an object in computer vision often stems from the pose, or specific orientation, of an object. Moreover, specularities (illusions caused by highly reflective surfaces) can fool an observer into misinterpreting the visual data. Generating more data about the scene by simple manipulation (e.g., poking an object of interest, see [7]) can dramatically improve the performance of the visual system.

The correlation between manipulation and visual perception is striking. However, very little research has taken place on the interaction between the two, probably because of the rarity of having manipulation, mobility, and vision on the same platform. We intend to endow *UMan* with manipulation capabilities such as turning a light switch to effect the lighting conditions, and moving objects around to gain new perspective.

### B. Manipulation

We plan to develop elementary manipulation skills to allow *UMan* to interact with its environment in interesting ways. One such skill is grasping. Grasping has traditionally been performed by a rigorous analysis of grasp geometry [14] or by haptically probing the grasped object until form/force closure is reached [15]. However, many interesting objects have been designed to be grasped by humans. We will develop so-called "ballistic" or "optimistic" grasp controllers. These controllers exploit visual clues obtained from object features to instantiate various parameters of relatively low-feedback grasping behavior. During the execution of such behavior, visual and haptic feedback can be used to transition between reaching and grasping behaviors, for example, or to detect failure. Our grasp controllers are called optimistic, because they assume that successful grasping behavior can be initiated based on perceived object features. This stands in contrast with approaches that employ elaborate analysis to ensure success. Instead, these behaviors leverage relevant properties of objects in the environment.

### C. Tool Use

Tool use represents one of the hallmarks of human and animal intelligence. Creating a robot that can use a known tool, learn how to use new tools, and choose the best tool for a given task is an important and challenging problem in autonomous mobile manipulation [20]. We will develop methods that enable *UMan* to extract kinematic models of objects in the environment. These methods will rely on a multi-modal stream of sensor information, including exploratory

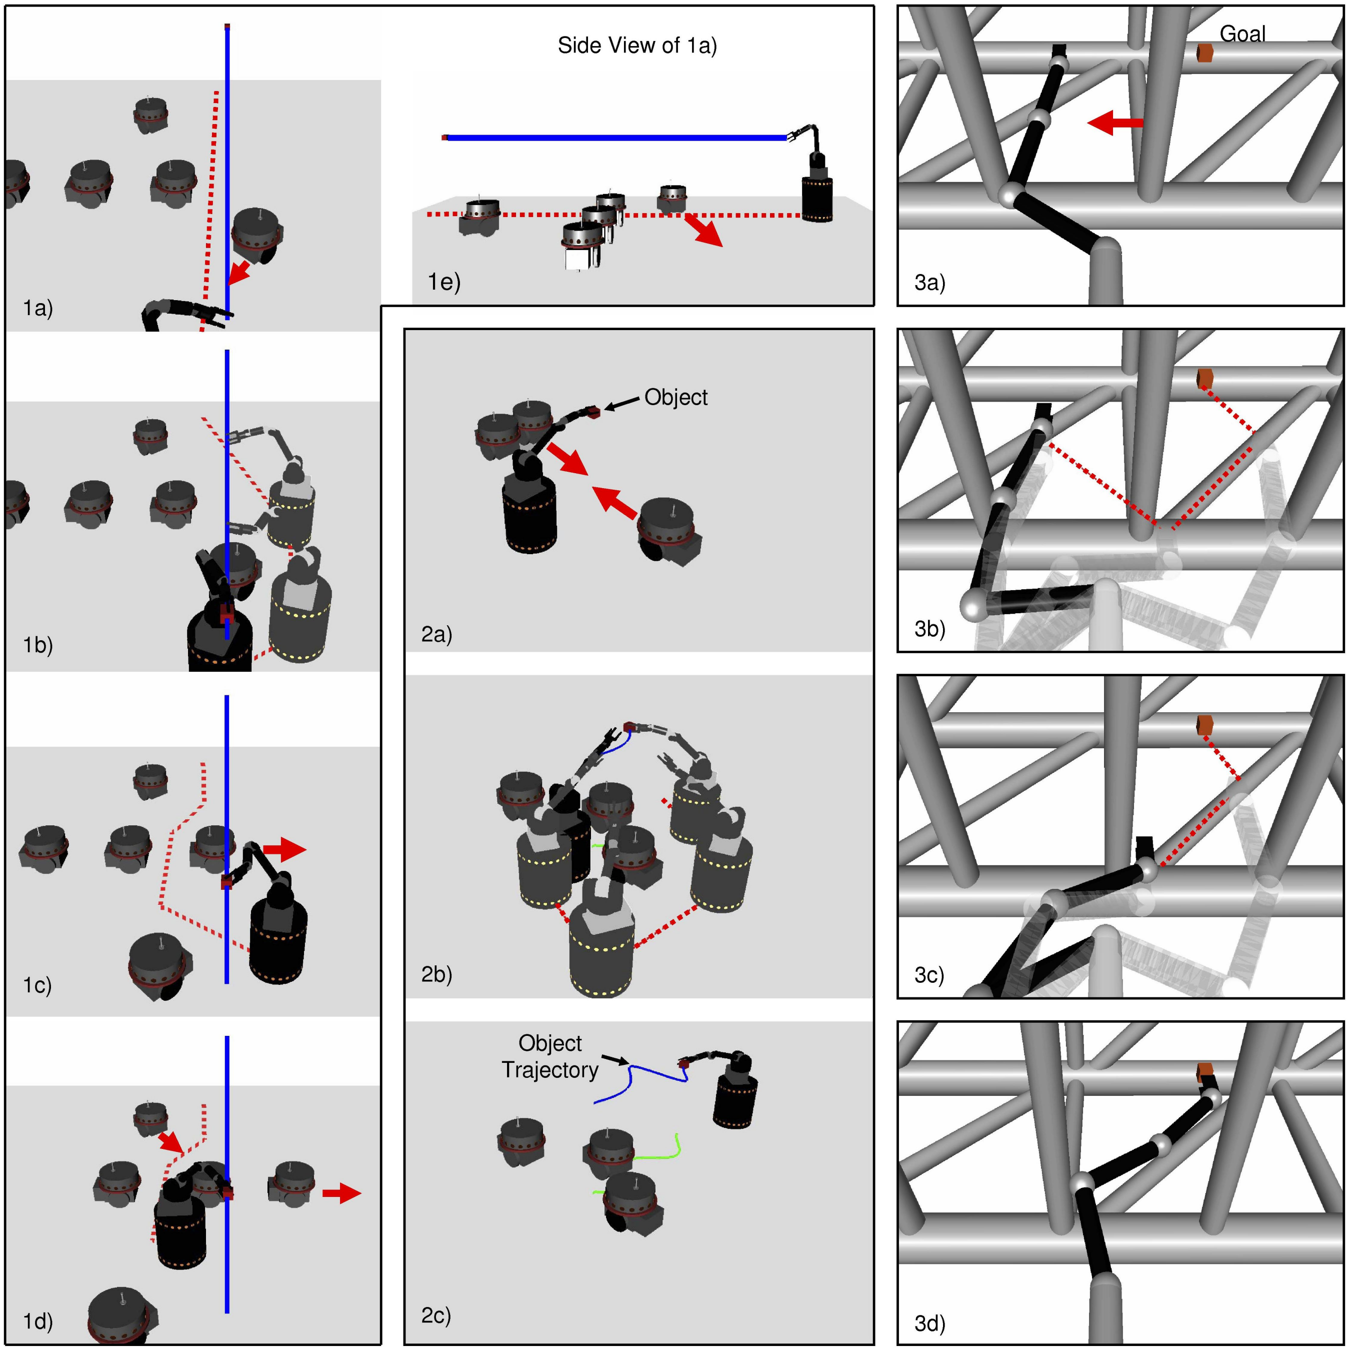

Fig. 9. Globally task-consistent motion with the elastic roadmap framework. The task is specified solely in terms of end-effector behavior and the elastic roadmap framework ensures that collisions are avoided and that task behavior is resumed, should it be interrupted by changes in the environment. In all of the figures, lighter versions of the robot represent milestones of the elastic roadmap that are part of the current solution paths. The connectivity of these milestones is indicated by a dashed line. The darker robots represent the actual robot in motion. The direction of motion of obstacles is indicated by arrows. *Experiment 1:* Images 1a) and 1e) show two perspectives of the same scene. The robot follows the line with its end-effector, while moving obstacles invalidate the solution. The elastic roadmap framework repeatedly generates global, task-consistent motion plans until the robot reaches the goal. *Experiment 2:* The task consists of following an object moving on an unknown trajectory. The following task is achieved based on force control. Moving obstacles force the robot to suspend the force control task, loosing contact with the object. The elastic roadmap framework computes a path to re-attain the force control task, shown in image 2c). This image also shows the trajectory taken by the object and its projection onto the floor. *Experiment 3:* A stationary robot, operating under a moving truss, reaches for a goal location while maintaining a constant orientation with its end-effector. The sequence of images shows how the goal can be reached. Continued motion by the truss will repeatedly force the robot to move away from the goal location to avoid collision. The elastic roadmap framework repeatedly generates motions such as those shown in images 3a)–3c) to re-attain the goal location.

behavior based on force-compliant motion primitives. The ability to obtain a kinematic and dynamic model of objects in the environment, we believe, represents one of the necessary preconditions for learning the function of tools and objects, learning how to use them, and learning how objects and tools can be employed in new situations.

Initially, we intend to employ the basic gasping skills described in the previous section, together with exploratory behavior, to obtain the kinematic models of a variety of door handles from *UMan*'s multi-modal sensor stream. We regard door handles as a very simple type of tool, since it only possesses a single degree of freedom (in most cases) and is affixed to the environment. We hope to generalize these methods to other tools and objects, such as pliers, drawers, doors, latches, lids, and later to more complicated mechanisms, such as those found in mechanical vises.

## VIII. CONCLUSION

Autonomous mobile manipulation in many ways can be viewed as a "Grand Challenge" for robotics and related disciplines. The successful deployment of autonomous robots to perform complex manipulation tasks in unstructured environments will not only require significant advances in a variety of areas of robotics and computer vision, it will also require that these new technologies be integrated into a single, robust, and autonomous robotic platform. In this paper, we described our efforts towards building such a mobile manipulation platform. Our mobile manipulator, *UMan*, was designed to support research in autonomous mobile manipulation. We have combined a number of off-the-shelf components to construct a hardware platform with redundant kinematic degrees of freedom, a comprehensive sensor suite, and significant end-effector capabilities for manipulation.

We believe that autonomous mobile manipulation requires tight integration of a broad range of capabilities and skills. Thus, our efforts to provide a comprehensive software platform that enables competent, robust, and autonomous task execution on our mobile manipulator. This software suite includes the elastic roadmap approach, an efficient approach to globally task-consistent motion that combines multi-objective, task-level control with real-time computation of global motion plans. We reported on our efforts to integrate one of the freely available software packages for mobile robot mapping and navigation with our multi-objective, task-level control scheme. In addition, we described our initial attempts to develop perceptual capabilities that are tailored to the requirements of autonomous mobile manipulation. Finally, we gave an overview of planned research activities to be supported by our combined hardware and software platform *UMan*.

## ACKNOWLEDGMENTS

This work is supported in part by the National Science Foundation (NSF) under grants CNS-0454074, IIS-0545934, IIS-0546666, and CNS-0552319, and by QNX Software Systems Ltd. in the form of a software grant. We are grateful for this support. The authors are greatly indebted to Bob Holmberg for providing software for the control of the XR4000. Also, we are very grateful for Bob's invaluable advice and help during many stages of this project.

## REFERENCES

[1] R. Ambrose. Experiments with a teleoperated robonaut on a model of the Hubble space telescope. Personal communication, 2005.

[2] R. O. Ambrose, S. R. Askew, W. Bluethmann, , and M. A. Diftler. Robonaut: NASA's space humanoid. *IEEE Intelligent Systems and Applications, Special Issue on Humanoids*, 15(4):57–63, 2000.

[3] O. Brock and R. Grupen. Final report for the NSF/NASA Workshop on Autonomous Mobile Manipulation (AMM). http://www-robotics.cs.umass.edu/amm/results.html, November 2005.

[4] CARMEN. http://carmen.sourceforge.net/.

[5] K.-S. Chang, R. Holmberg, and O. Khatib. The augmented object model: Cooperative maniplulation and parallel mechanism dynamics. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 470–475, San Francisco, USA, 2000.

[6] A. Edsinger-Gonzales. Design of a compliant and force sensing hand for a humanoid robot. In *Proceedings of the International Conference on Intelligent Manipulation and Grasping*, 2004.

[7] P. Fitzpatrick and G. Metta. Grounding vision through experimental manipulation. *Philosophical Transactions of the Royal Society: Mathematical, Physical, and Engineering Sciences*, 361(1811):2165–2185, 2003.

[8] L. M. G. Goncalves, R. Grupen, A. A. F. Oliveira, D. Wheeler, and A. Fagg. Attention and categorization: Basis for cognition in a humanoid robot. *IEEE Intelligent Systems and Applications, Special Issue on Humanoids*, 2000.

[9] Hokuyo. http://www.hokuyo-aut.jp/products/urg/urg.htm.

[10] R. Holmberg and O. Khatib. Development and control of a holonomic mobile robot for mobile manipulation tasks. *International Journal of Robotics Research*, 19(11):1066–1074, 2000.

[11] M. Huber and R. A. Grupen. A feedback control structure for on-line learning tasks. *Robotics and Autonomous Systems*, 22(3-4):303–315, 1997.

[12] Kawada Industries. http://www.kawada.co.jp/global/ams/hrp_2.html.

[13] O. Khatib. A unified approach to motion and force control of robot manipulators: The operational space formulation. *International Journal of Robotics and Automation*, 3(1):43–53, 1987.

[14] M. T. Mason. *Mechanics of Robotic Manipulation*. MIT Press, 2001.

[15] R. Platt, A. H. Fagg, and R. A. Grupen. Nullspace composition of control laws for grasping. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Lausanne, Switzerland, 2002.

[16] QNX. http://www.qnx.com.

[17] Y. Sakagami, R. Watanabe, C. Aoyama, S. Matsunaga, N. Higaki, and K. Fujimura. The intelligent ASIMO: system overview and integration. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, volume 3, pages 2478–2483, 2002.

[18] L. Sentis and O. Khatib. Synthesis of whole-body behaviors through hierarchical control of behavioral primitives. *International Journal of Humanoid Robots*, 2(4):505–518, 2005.

[19] Sony. http://www.sony.net/SonyInfo/QRIO/.

[20] A. Stoytchev. Behavior-grounded representation of tool affordances. In *Video Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2005.

[21] K. B. Tamim Asfour and R. Dillmann. The humanoid robot armar: Design and control. In *Proceedings of the IEEE International Conference on Humanoid Robots*, 2000.

[22] S. Thrun. Robotic mapping: A survey. In G. Lakemeyer and B. Nebel, editors, *Exploring Artificial Intelligence in the New Millenium*. Morgan Kaufmann, 2002.

[23] S. Thrun, D. Fox, W. Burgard, and F. Dellaert. Robust monte carlo localization for mobile robots. *Artificial Intelligence*, 128(1-2), 2000.

[24] W. T. Townsend and J. K. Salisbury. Mechanical design for whole-arm manipulation. *Robots and biological systems: towards a new bionics?*, pages 153–164, 1993.

[25] Toyota. http://www.toyota.co.jp/en/special/robot/.

[26] Wabian. http://www.uc3m.es/uc3m/dpto/IN/dpin04/Wabian03.html.

[27] Y. Yang and O. Brock. Elastic roadmaps: Globally task-consistent motion for autonomous mobile manipulation. In *Proceedings of Robotics: Science and Systems (RSS)*, Philadelphia, USA, August 2006.