Manipulation in Human Environments

Aaron Edsinger and Charles C. Kemp
Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139
Email: {edsinger@csail.mit.edu, cckemp@csail.mit.edu}

***Abstract*— Robots that work alongside us in our homes and workplaces could extend the time an elderly person can live at home, provide physical assistance to a worker on an assembly line, or help with household chores. In order to assist us in these ways, robots will need to successfully perform manipulation tasks within human environments. Human environments present special challenges for robot manipulation since they are complex, dynamic, uncontrolled, and difficult to perceive reliably.**

In this paper we present a behavior-based control system that enables a humanoid robot, Domo, to help a person place objects on a shelf. Domo is able to physically locate the shelf, socially cue a person to hand it an object, grasp the object that has been handed to it, transfer the object to the hand that is closest to the shelf, and place the object on the shelf.

We use this behavior-based control system to illustrate three themes that characterize our approach to manipulation in human environments. The first theme, *cooperative manipulation*, refers to the advantages that can be gained by having the robot work with a person to cooperatively perform manipulation tasks. The second theme, *task relevant features*, emphasizes the benefits of carefully selecting the aspects of the world that are to be perceived and acted upon during a manipulation task. The third theme, *let the body do the thinking*, encompasses several ways in which a robot can use its body to simplify manipulation tasks.[1]

Fig. 1. The humanoid robot Domo used in this paper.

I. Introduction

Robots that work alongside us in our homes and workplaces could extend the time an elderly person can live at home, provide physical assistance to a worker on an assembly line, or help with household chores. In order to assist us in these ways, robots will need to successfully perform manipulation tasks within human environments. Human environments present special challenges for robot manipulation since they are complex, dynamic, uncontrolled, and difficult to perceive reliably.

Addressing these issues is a focus of several active projects. The ARMAR project is investigating manipulation in human environments and has shown results including the bimanual opening of a jar [1]. Researchers working with the NASA Robonaut [2] have demonstrated a cooperative manipulation task where the robot employs a power drill to tighten lugnuts under human direction. Work at Fraunhofer IPA with Care-O-bot II has pursued fetch-and-carry tasks of everyday objects [3]. In addition, many groups are pursuing research on autonomous mobile manipulation in human environments [4], [5], [6].

[1]This work was sponsored by Toyota Motor Corporation: Autonomous Manipulation Capabilities for Partner Robots in the Home.

In this paper we present a behavior-based control system that enables a humanoid robot, Domo, to help a person place objects on a shelf. We describe the implementation of this control system and present a quantitative evaluation of its performance. We also use this behavior-based control system to illustrate three themes that characterize our approach to manipulation in human environments:

II. The Task

We present a behavior-based control system implemented on the humanoid robot Domo, pictured in Figure 1. For this work, Domo uses 29 DOF, a single camera, and Series Elastic Actuators [7], [8]. As shown in Figure 2, Domo takes objects that have been handed to it and places them on a shelf. Domo is able to physically locate the shelf, socially cue a person to hand it an object, grasp the object that has been handed to it, transfer the object to the hand that is closest to the shelf, and place the object on the shelf. By performing this task, the robot effectively extends the person's reach, allowing her to place objects in locations that might be difficult or uncomfortable to access without assistance. If this skill were combined with a mobile base, the person's effective reach could be dramatically extended. For an individual with serious physical limitations, this help might allow the person to maintain autonomy in everyday activities that would otherwise require help from another person. For example, an elderly person in a wheelchair might use a robot with this ability to put away common household objects, such as books and dishes.

III. Three Themes for Design

Three themes characterize our approach to manipulation in human environments. The first theme, *cooperative manipulation*, refers to the advantages that can be gained by having the robot work with a person to cooperatively perform manipulation tasks. The second theme, *task relevant features*, emphasizes the benefits of carefully selecting the aspects of the world that are to be perceived and acted upon during a manipulation task. The third theme, *let the body do the thinking*, encompasses several ways in which a robot can use its body to simplify manipulation tasks.

A. Cooperative manipulation

For at least the near term, robots in human environments will be dependent on people. Fortunately, people tend to be present within human environments. As long as the robot's usefulness outweighs the efforts required to help it, robot autonomy is unnecessary. Careful design can make robots intuitive to use, thereby reducing the effort required. For example, the initial version of the commercially successful Roomba relies on a person to occasionally prepare the environment, rescue it when it is stuck, and direct it to spots for cleaning and power. The robot and the person effectively vacuum the floor as a team, although the person's involvement is reduced to a few infrequent tasks that are beyond the capabilities of the robot.

By treating tasks that involve manipulation as a cooperative process, people and robots can perform tasks that neither one could perform as an individual. Researchers have looked at techniques for cooperative manipulation that physically couple the robot and the person, such as carrying an object together [9], [10]. Fewer researchers have investigated the use of social cues for cooperative manipulation [2].

For the task described in this paper, the person hands Domo the object to be placed on the shelf. This cooperation avoids the challenges involved with having a robot autonomously select, locate, and grasp an object. Domo uses social cues to simplify this cooperation. If Domo sees a person, it reaches out toward the person with an open grasp, which communicates to the person that Domo is prepared to place an object on the shelf. Domo's outstretched, open hand gives a clear indication of where the person should place the object and biases the person to place the object within Domo's hand at a desirable orientation.

B. Task relevant features

Donald Norman's book *The Design of Everyday Things* [11], indicates that objects found within human environments are likely to have common structural features that simplify their use (see Figure 3). By developing behaviors that are matched to such structural features, we can simplify robot manipulation in human environments. For example, we have previously shown that the manipulation of a large set of human tools can be specified in terms of the tool's tip, such as the tip of a screwdriver [12].

Fig. 3. Donald Norman's "Coffeepot for Masochists". Many objects in human environments have been designed to match our physical and cognitive abilities. The design of a traditional coffeepot, for example, has evolved such that the pot and coffee can be easily controlled from the handle, the handle is matched to a human-scale power grasp, and the spout is positioned to accommodate perception and control of the spout during pouring. (Personal collection of D. A. Norman. Photograph by Norman. Reproduced with permission).

Many researchers treat robot manipulation as a planning problem performed with respect to the global state of the world [13], [14]. In contrast, our work is influenced by the work of researchers such as Jagersand, Platt and Grupen, Connell, and Brooks [15], [16], [17], [18], who make use of carefully chosen aspects of the world's state. Rather than attempting to reconstruct the world in its entirety, we focus the robot's sensory resources on elements of the world that are relevant to the current task, such as the contact surface of the hand, the edge of the shelf, or the grasp of an object. Other researchers have used task relevant features for manipulation, although typically with fiducial markers or simplified environments [19], [20]. Except for fiducial markers on the edge of the shelf, we have not altered the world to accommodate perception.

We define tasks in terms of behaviors that perform closed-loop control with respect to these features, so that at all times the robot has rich feedback about its performance and the ability to react to the unexpected. In particular, if one of these constantly monitored aspects of the world violates the requirements of the current behavior, the robot has fallback behaviors to which it can resort. Since a human is present, the robot can ask for help as a last resort.

C. Let the Body Do The Thinking

This theme bundles together a number of design strategies that make use of the robot's body to simplify manipulation. Research on robot locomotion has convincingly demonstrated the benefits of exploiting compliance and natural dynamics for robot control when the robot is in contact with the world [21], [22], [23]. Moreover, Williamson's work [24] on robot manipulation shows that similar strategies can be successfully applied to robot manipulation.

1) Human Form: As we have discussed within [25] and Norman has discussed in [11], human environments are well matched to the human body. Domo's human form allows

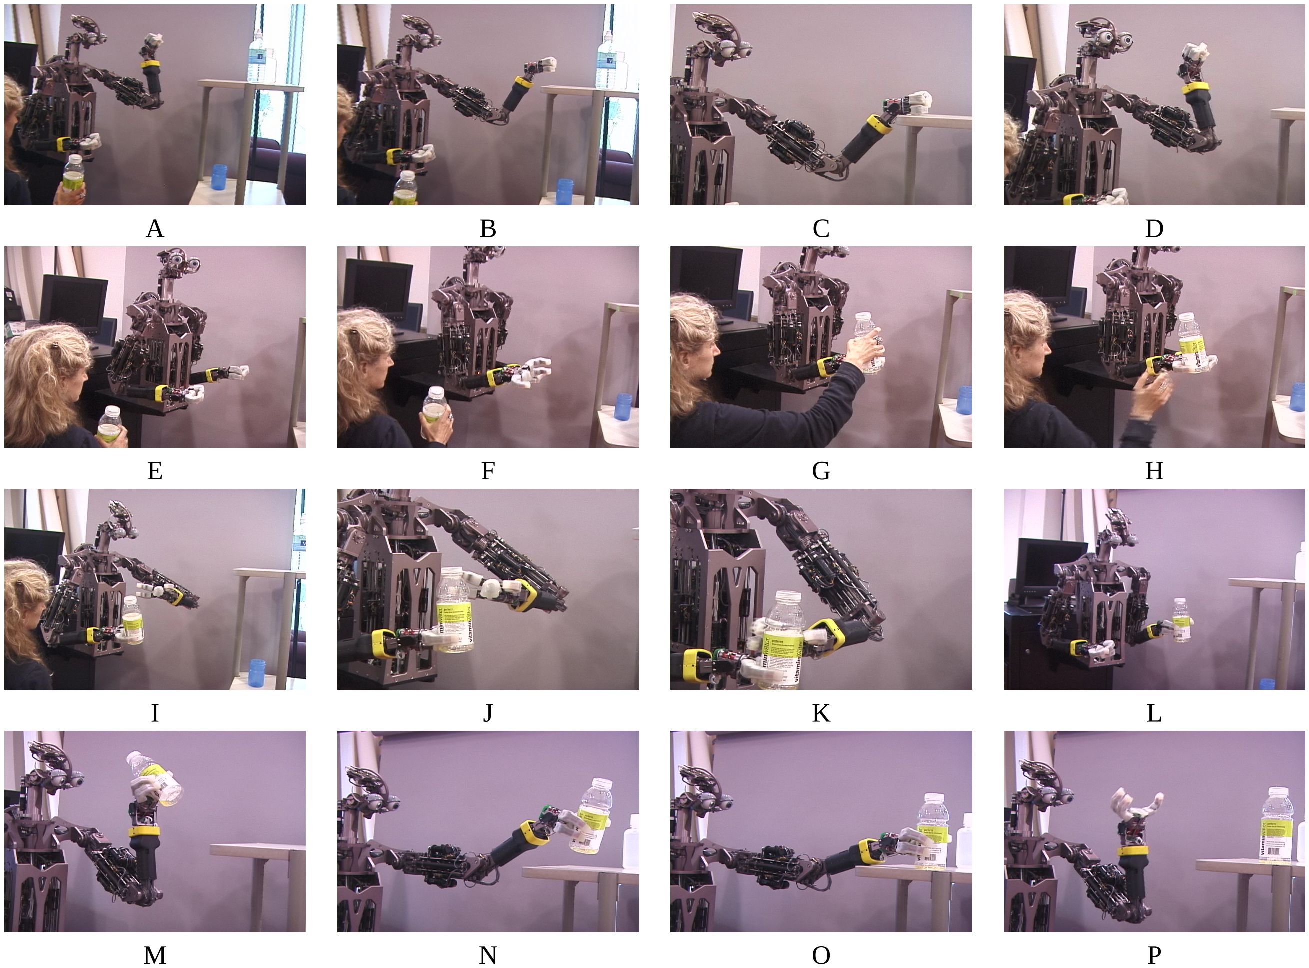

Fig. 2. Domo assists a person in placing a bottle on a shelf. (A-D) Hypothesis testing. The shelf is rolled up to Domo. It is visually detected and a visually guided reach in the direction of the shelf is performed. The arm compliantly lowers onto the shelf and its location is confirmed. The proprioceptive state of the arm is used to represent the location of the shelf in the world. (E-H) Cooperative interaction. A person is detected by Domo. Domo reaches towards the person, cueing her to place a bottle in its hand. In placing the bottle in Domo's hand, the person intuitively offers it in an appropriate orientation. A grasp reflex is triggered when contact forces are detected in the arm. A successful grasp is detected and the arm is lowered. (I-L) Transfer. The shelf is out of the person's reachable workspace but within the workspace of the robot's left arm. The left arm performs a visual reach to the bottle, using hand motion to detect and estimate the location of the hand's contact surface in the image. The left arm compliantly lowers on to the right. Contact forces detected in the right arm trigger a grasp reflex in the left hand. If both hands have a successful grasp, the right hand releases the bottle and the transfer is complete. (M-P)Placement. Domo reaches to just above the shelf surface using its previously estimated location. The arm and wrist are place in a compliant force mode and the bottle is lowered onto the shelf. The manipulator compliance and downward force allow the bottle to become aligned with the shelf. If successful shelf contact is detected, the bottle is released and the arm retracted.

it to intuitively cue the person with whom it is working. Domo's eye gaze, arm gesture, and open hand are similar in appearance to a human requesting an object. This can help communicate Domo's request and cue the appropriate response more effectively than a wholly alien body. Also, Domo's eyes and arms are high off of the ground, which simplifies perception and action relative to the shelf. In addition, Domo's hand is well suited to grasping everyday cylindrical objects, since it is approximately the size of an adult's hand and has a compliant exterior.

2) Compliance: Domo is mechanically distinctive in that it incorporates passive compliance and force sensing throughout its body [7]. Domo's 22 Series Elastic Actuators lower the mechanical impedance of the its arms, allowing for safe physical interaction with a person [26]. In addition, the angle of each joint can be controlled as a virtual spring, where the resting setpoint of the spring and the stiffness can be specified in real-time.

Compliance allows Domo to safely explore the world, such as when it finds the shelf with its hand. Domo also uses low stiffness control to take advantage of favorable natural dynamics, such as the tendency of an object to stably align itself with a flat surface. When transferring an object between its hands, compliance helps Domo to achieve a good grasp on the object without knocking the object out of its hands.

3) Active Perception: By taking action, the robot can make task relevant perception easier. Domo selects postures in order to more easily view task relevant features, which is especially

Fig. 4. The robot's view as it transfers an object to its right hand to its left. The everyday environment includes the shelf, a person, as well as natural lighting and a cluttered background. The edge motion (right) of the hand is used to estimate the location of the contact surface (green). This estimate is then used to visually servo the contact surface in the scene. In addition, the robot actively adapts its posture such that it can more easily view the controlled features.

important for visual servoing. As shown in Figure 4, when the robot transfers an object from one hand to another, it moves its hands to a posture such that the contact surface of the hand can be readily detected and servoed. From this posture, the convex region of the hand's grip corresponds with a convex edge in the image. The robot also uses visual motion and active control of objects to simplify detection of task relevant features. For example, the motion of the robot's hand simplifies detection of the hand's convex contact surface.

Domo also uses its body to reduce uncertainty by making contact with the world. The robot uses an exploratory behavior to physically locate the surface of the shelf after getting a coarse visual estimate. The robot reaches above the shelf and then moves its arm down compliantly until contact. This is similar to what a person does when searching for a surface in the dark. Through this behavior, the robot also finds an arm posture that is likely to lead to success when placing the object on the surface. This circumvents the need for planning an arm trajectory and target posture.

IV. Implementation

A. Behavior Architecture

The robot's control is implemented as a hierarchical set of perceptual and motor behaviors. The behaviors run in a distributed, real-time architecture at $15-100hz$ on a 15 node Linux cluster. We have adopted a layered approach similar to that of Brooks[27] and Connell[17]. We couple constant perceptual feedback to many simple behaviors in order to increase the task robustness. For example, if a person removes the object from the robot's grasp at anytime during task execution, the active behavior will become inhibited and a lower-level behavior may attempt to reacquire the object or to smoothly bring the arm to a relaxed posture. If the arm trajectory is disturbed during a visually guided motion, a head behavior will automatically keep the hand in view. In this section we first describe the overall algorithm and then describe some of the behaviors in more detail.

B. Algorithm

As shown in Figure 5, the task can be decomposed into four behaviors: *ShelfTest*, *CollaborativeGrasp*, *Transfer*, and *Place*.

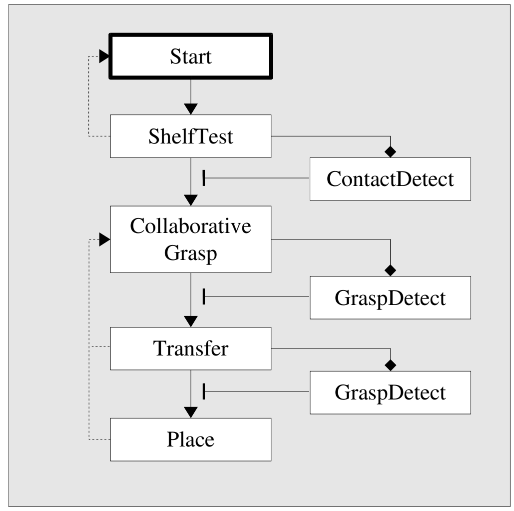

Fig. 5. A high level description of the task implementation. The robot assists a person in placing objects on a shelf using four behaviors: *ShelfTest, CollaborativeGrasp, Transfer*, and *Place*. In the behavior diagram, bold indicates the default behavior. Transitions (arrows) occur contingent on perceptual feedback (bars). For a given behavior, the robot takes actions to increase the likelihood and robustness of a desired perceptual feature (diamond). Exceptions from the expected feedback result in a reset transition (dashed line). *ShelfTest* detects and verifies the location of a useable surface. If the robot is not grasping an object, *CollaborativeGrasp* utilizes natural social cues to gain human assistance in grasping an object. Depending on the shelf location, *Transfer* can optionally allow the robot to pass the object to the other hand. Finally, if the robot is holding an object in the hand nearest to the shelf, *Place* allows it to put the object on the shelf.

These behaviors can be further decomposed into a shared set of perceptual and motor behaviors, as shown in Figures 6,7,8,and 9. Our algorithm is as follows:

1) *ShelfTest* (Figure 6). Detect and verify the location of a useable surface.
   a) Visually identify a candidate flat surface (through fiducials).
   b) Reach out along a ray that is above the front edge using visual servoing.
   c) Use compliant force control to move the hand down and make contact with the surface.
   d) Detect contact (or lack of) with the shelf.
   e) Store the posture, prior to descent, that led to success, or try again.
2) *CollaborativeGrasp* (Figure 7). Transfer the object from a person to the robot.
   a) Detect the person.
   b) Reach to the person, cueing them to offer an object at the appropriate place and orientation.
   c) Detect the interaction forces created by the object being placed in the hand.
   d) Form a power grasp and detect its success.
   e) If no success, relinquish control.
3) *Transfer* (Figure 8). Transfer the object from one hand to the other.
   a) Achieve a posture which brings both hands into the

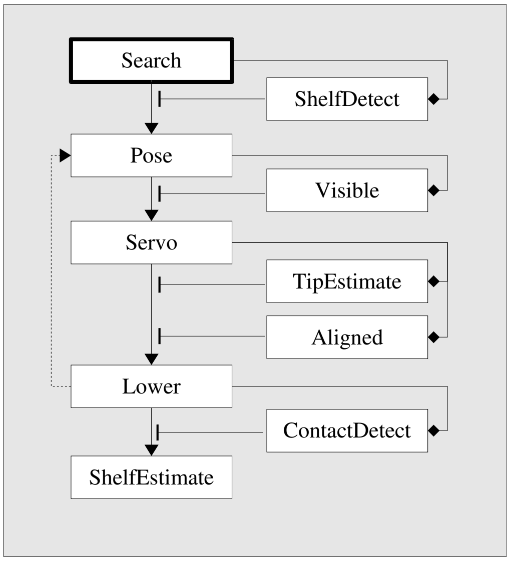

Fig. 6. *ShelfTest*. The robot detects and verifies the location of a useable surface. See Figure 5 for the diagram notation. When the robot has not yet detected a shelf, it engages in periodic visual search around the room (*Search*), increasing the likelihood of detecting a shelf. When *ShelfDetect* finds a hypothesized shelf, the robot visually fixates the location and moves its nearest hand to the surface into the image (*Pose*) until *Visble* is true. The hand is then visually servoed (*Servo*) to a visual target just above the shelf edge. The arm is also extended to a fixed depth along the camera's optical axis. The visual motion of its arm is used to detect the tip of the hand and improve the precision of the servoing process (*TipEstimate*). When the hand has achieved its target (*Aligned*), the hand is lowered onto the shelf using compliant force control *(Lower)*. If stable contact is detected (*ContactDetect*), the posture, prior to descent, that led to success is stored (*ShelfEstimate*). Otherwise, the robot may try the process again.

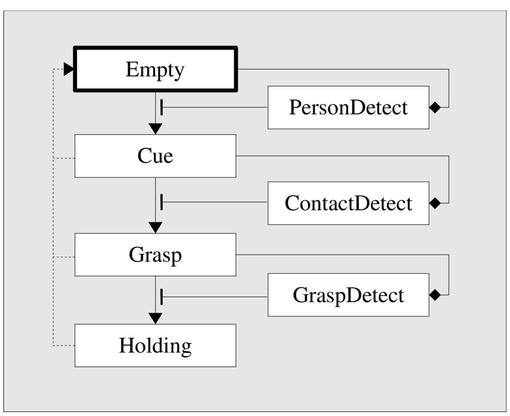

Fig. 7. *CollaborativeGrasp*. The robot leverages natural social cues for human assistance in grasping an object. See Figure 5 for the diagram notation. When its hands are empty (*Empty*), the robot engages in visual search, increasing the likelihood that a person is detected in the scene. When *PersonDetect* signals that a person is waving their hand, the robot cues the person to hand it an object. For the *Cue* behavior, the robot reaches towards the person, its hand open, at an appropriate time, location, and orientation such that the person intuitively reads the gesture as a request. The stiffness of the arm is also lowered. These actions increase the likelihood that a person will place an appropriately oriented object in the robot's hand and that the interaction forces of the placement can be sensed. When interaction forces are sensed at the hand (*ContactDetect*), the robot executes a power-grasp. If *GraspDetect* signals a successful grasp, the robot remains in the *Holding* state until *GraspDetect* is no longer true. Through continual perceptual feedback and by placing the person within the feedback loop, the robot is able to grasp objects in the environment.

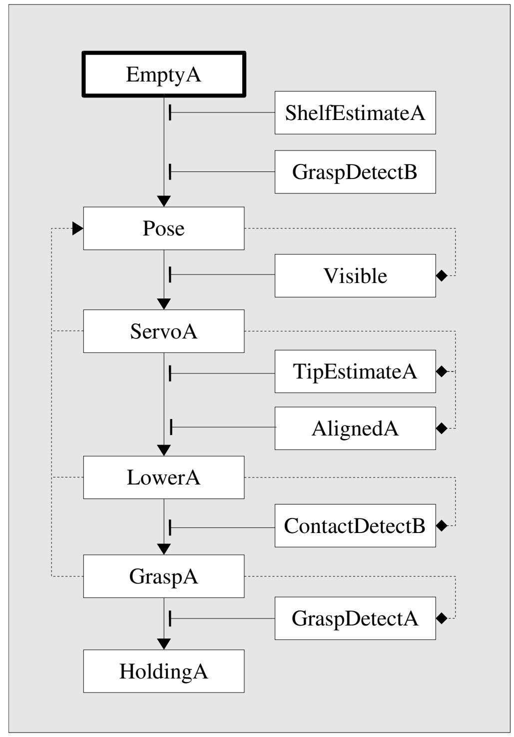

Fig. 8. *Transfer*. The robot transfers an object between hands. See Figure 5 for the diagram notation. If hand *A* is not holding an object (*EmptyA*), hand *B* is, and a shelf exists nearest hand *A* (*ShelfEstimateA*), then the robot moves its two arms to a stereotype pose (*Pose*). The pose is selected such that both the object and the contact surface of hand *A* are visible (*Visible*).The contact surface of Hand *A* is then visually servoed (*ServoA*) to a fixed offset from hand *B* (*AlignA*). The visual motion of its arm is used to detect the contact surface and improve the precision of the servoing process (*TipEstimateA*). Hand *A* is lowered onto hand *B* using compliant force control *(LowerA)*. The stiffness of arm *B* is also lowered, increasing the likelihood of detecting the contact between the two arms (*ContactDetectB*). Finally, the behavior is successful if hand *A* executes a stable power grasp on the object (*GraspDetectA*) and hand *B* releases the object. Otherwise, the robot may try the process again.

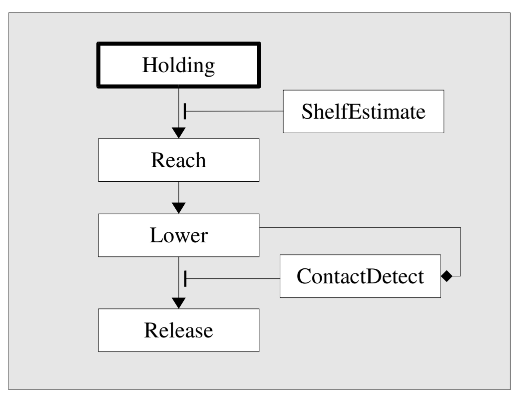

Fig. 9. *Place*. The robot an object from its hand to a shelf. See Figure 5 for the diagram notation. If a shelf exists (*ShelfEstimate*) nearest a grasping hand (*Holding*), the robot reaches to the previously learned posture that places the hand just above the shelf. The object is then lowered onto the shelf using compliant force control (*Lower*). The passive and active compliance in the wrist and hand allow the object to naturally align into a stable pose. If contact with the shelf is detected (*ContactDetect*), the grasp is released and the arm is retracted, leaving the object on the shelf.

visual field-of-view.

  b) Visually servo the inner contact surface of the empty hand to be just above the other hand.
  c) Compliantly lower the empty hand through force control until the other hand detects the interaction forces.
  d) Form a power grasp and detect its success.
  e) If no success, back out and retry the servo.
  f) If both hands are grasping, release the older grasp.

4) *Place* (Figure 9). Transfer the object from the hand to the shelf.
  a) Achieve the previously learned posture that places the hand just above the shelf.
  b) Descend compliantly for a fixed duration using force control. Allow the object to self-align using compliance.
  c) Release and retract.

C. Behaviors

1) Lower: The arm is controlled by specifying a virtual force vector at the hand, lowering the arm until contact is made. Virtual forces are converted into commanded torques using the arm Jacobian, while compensating for gravity-induced torques. If a large hand displacement occurs (no contact), the behavior is inhibited.

2) PersonDetect: Visual detections of faces [28], skin color, and foreground motion [12] are used to detect when a person is present and waving their hand. A face detection and repeated skin motion (waving) at an expected hand location is presumed to be a cue for the robot.

3) GraspDetect: The grasp aperture is used to estimate both the diameter and existence of a held object during a force-controlled power grasp. The grasp aperture is difficult to estimate analytically due to finger collisions, complex kinematics, and passive compliance of the robot skin. A map was learned between the four finger joint angles and the grasp aperture using support vector regression (SVR) with a Gaussian RBF kernel. The training data was gathered through robot power grasps on 5 cylindrical objects of known diameters between 25 and 75mm. The orientation of the object was varied by ± 30 degrees and the proprioceptive resting state was recorded for a total of 50 trials. Consistent power grasp apertures above a threshold diameter are signaled as a stable grasp detection.

4) ShelfDetect: Fiducial markers attached to the leading edge of the shelf are used for an initial detection of the shelf. We expect to use a non-fiducial based approach in future work. The position and location of the shelf surface is not assumed, but is determined with the *ShelfTest* behavior.

5) ContactDetect: We use two methods to detect when the hand makes contact with the world [29]. First, the wrist and arm are held motionless with zero joint stiffness and gravity compensation. External interaction forces from a person or the robot's other arm are detected as non-zero accelerations at the hand. The second method does not assume zero acceleration and joint stiffness. A point-mass dynamic model is used to measure the error between the predicted and measured joint-torques during force controlled movements. Contact is detected when the error exceeds a conservative threshold.

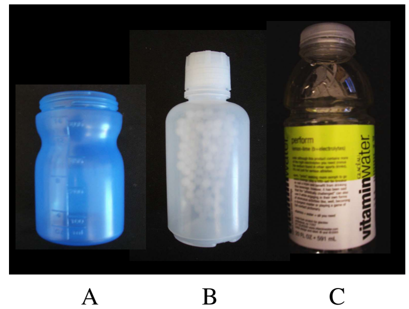

Fig. 10. The three bottles used in our experiments.

6) TipEstimate: The position of a fixed point within the hand's coordinate frame is estimated online through self-generated motion. The point is detected as the fastest moving convex feature in the image. This typically corresponds with the distal end of the robot's kinematic chain, allowing the robot to predict the visual location of its finger, palm, or the tip of a grasped object, even under perceptual uncertainty due to poor hand-eye calibration and unknown objects [12].

7) Servo: The tip of the kinematic chain, as estimated by TipEstimate, is servoed to a visual target within the 2D image plane. Reaching in depth is achieved by kinematically servoing the tip along the camera's optical axis. The visual motion generated by the servoing can be used to adapt TipEstimate online. For example, when performing a transfer between hands, the robot's palm is visually servoed to the object.

V. Results

We tested Domo's performance on the task over 18 trials with two subjects, where each trial lasted approximately one minute. A trial consisted of the subject handing Domo a bottle, Domo transferring the bottle to its other hand and then placing it on the shelf. One trial is depicted in Figure 2. Each subject performed 3 trials on each of the 3 empty bottles shown in Figure 10. The bottles vary in diameter from $40-75$mm and length from $100-200$mm. For each subject, the shelf remained stationary and the *ShelfTest* behavior executed only once at the start of the experiment. We measured success using the following criteria:

1) *CollaborativeGrasp*: Stable grasp after transfer of the bottle from the person to the robot.
2) *Transfer*: Stable grasp after transfer of the bottle between hands.
3) *PlaceX*: Bottle X was left on the shelf.
4) *StandX*: Bottle X was left on the shelf standing upright.

As seen in Figure 11, Domo was largely successful at the task for the given objects. One subject was experienced in working with the robot at this task and consequently achieved a higher success rate. Failures were typically a result of insecure grasps being formed during the object transfer phase. Variability in the subject's placement of the object in the

	CollaborativeGrasp	Transfer	PlaceA	StandA	PlaceB	StandB	PlaceC	StandC
Subject 1	9/9	9/9	3/3	3/3	3/3	3/3	3/3	2/3
Subject 2	9/9	8/9	2/3	2/3	3/3	2/3	2/3	1/3

Fig. 11. Task results for 18 trials with two subjects.

robot's hand tended to be amplified by the transfer operation. In the future, we hope to improve robustness through active perception of the grasped object during the transfer, as well as more active control when the object is lowered onto the shelf.

VI. Discussion

Placing objects on flat surfaces is an important component for a variety of everyday tasks for which people desire help. Human environments are dominated by flat surfaces upon which people place objects. Stocking goods, setting the table, arranging a product display, placing a part onto a conveyor belt, and putting dishes away are all examples of tasks that typically involve placing objects on a flat surface. Our approach takes steps towards performing these more general tasks in a cooperative way.

Motion planning systems have achieved impressive results performing sophisticated manipulation tasks in simulations that assume full knowledge of the state of the world and the robot. Unfortunately, for now and the foreseeable future, sensor and actuation technologies will force a robot to perform tasks using uncertain, piecemeal views of its body and the world. Methods must be developed that help robots overcome these uncertainties and the unexpected events that result. Within the domain of robot navigation, researchers have addressed similar issues by explicitly modeling uncertainty with statistical methods, [30], typically with respect to state representations in the form of 3d maps and the robot's pose. Related statistical methods will almost certainly play a role in addressing the challenges of manipulation in human environments, but they are only one piece to the puzzle. Without appropriate high-level design, explicit models of uncertainty will be ineffective or unnecessarily complex.

The themes we highlight in this paper, and the behavior based building blocks we used to achieve the cooperative manipulation task, should be applicable to new manipulation tasks in human environments. So far, we have only used a statistical learning technique to improve our grasp detector. We are currently seeking to integrate learning and statistical representations of uncertainty into other parts of our system in order to further improve robustness and autonomy.

References

[1] T. A. R. Zllner and R. Dillmann, "Programming by demonstration: Dual-arm manipulation tasks for humanoid," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2004)*, Sendai, Japan, September 28 - October 2, 2004 2004.

[2] W. Bluethmann, R. Ambrose, A. Fagg, M. Rosenstein, R. Platt, R. Grupen, C. Brezeal, A. Brooks, A. Lockerd, R. Peters, O. Jenkins, M. Mataric, and M. Bugajska, "Building an Autonomous Humanoid Tool User," in *Proceedings of the 2004 IEEE International Conference on Humanoid Robots*. Santa Monica, Los Angeles, CA, USA.: IEEE Press, 2004.

[3] B. Graf, M. Hans, and R. D. Schrft, "Development of a next generation robotic home assistant." *Autonomous Robots*, vol. 16, no. 2, pp. 193–205, 2004.

[4] O. Khatib, B. O. C. K. Yokoi, K., and A. Casal, "Robots in human environments: Basic autonomous capabilities," *International Journal of Robotics Research*, vol. 18, no. 684, 1999.

[5] S. Caselli, E. Fantini, F. Monica, P. Occhi, and M. Reggiani, "Toward a mobile manipulator service robot for human assistance." [Online]. Available: citeseer.ist.psu.edu/caselli03toward.html

[6] Y. Yang and O. B. E. Roadmaps, "Globally task-consistent motion for autonomous mobile manipulation," in *Proceedings of Robotics: Science and Systems*, Philadephia, USA, August, 2006.

[7] A. Edsinger-Gonzales and J. Weber, "Domo: A Force Sensing Humanoid Robot for Manipulation Research," in *Proceedings of the 2004 IEEE International Conference on Humanoid Robots*. Santa Monica, Los Angeles, CA, USA.: IEEE Press, 2004.

[8] G. A. Pratt and M. M. Williamson, "Series elastic actuators," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-95)*, vol. 1, Pittsburg, PA, July 1995, pp. 399–406.

[9] B. Waarsing, M. Nuttin, and H. V. Brussel, "Introducing robots into a human-centred environment - the behaviour-based approach," in *Proceedings of the 4th International Conference on CLAWAR*, May 2001, pp. 465–470.

[10] K. Kyong, S. Freedman, M. Mataric, M. Cunningham, and B. Lopez, "A hands-off physical therapy assistance robot for cardiac patients," in *9th International Conference on Rehabilitation Robotics*, July 2005, pp. 337– 340.

[11] D. A. Norman, *The Design of Everyday Things*. New York: Doubleday.

[12] C. C. Kemp and A. Edsinger, "Robot manipulation of human tools: Autonomous detection and control of task relevant features," in *In Submission to: 5th IEEE International Conference on Development and Learning (ICDL-06)*, Bloomington, Indiana, 2006.

[13] T. Simeon, J. P. Laumond, J. Cortes, and A. Sahbani, "Manipulation planning with probabilistic roadmaps," *IJRR*, vol. 23, no. 7-8, pp. 729–746, 2004.

[14] M. Saha and P. Isto, "Motion planning for robotic manipulation of deformable linear objects," in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2006.

[15] M. Jagersand and R. Nelson, "Visual Space Task Specification, Planning and Control," in *Proceedings of the IEEE International Symposium on Computer Vision*, 1995, pp. 521–526.

[16] R. Platt, A. H. Fagg, and R. Grupen, "Manipulation gaits: Sequences of grasp control tasks," in *Proceedings of the 2004 IEEE Conference on Robotics and Automation (ICRA)*, April 2004.

[17] J. Connell, "A behavior-based arm controller," *IEEE Transactions on Robotics and Automation*, vol. 5, no. 5, pp. 784–791, December 1989.

[18] R. A. Brooks, *Cambrian Intelligence*. Cambridge, MA: MIT Press, 1999.

[19] J. H. Piater and R. A. Grupen, "Learning appearance features to support robotic manipulation," *Cognitive Vision Workshop*, 2002.

[20] N. Pollard and J. Hodgins, "Generalizing Demonstrated Manipulation Tasks," in *Proceedings of the Workshop on the Algorithmic Foundations of Robotics (WAFR '02)*, December 2002.

[21] M. H. Raibert, *Legged robots that balance*. Cambridge, MA, USA: Massachusetts Institute of Technology, 1986.

[22] Fast and R. H. R. via Shape Deposition Manufacturing, "J. g. cham and s. a bailey and j. e clark and r. j full and m. r cutkosky," *The International Journal of Robotics Research*, vol. Volume 21 Issue 10, 2002.

[23] R. L. Tedrake, "Applied optimal control for dynamically stable legged locomotion," Ph.D. dissertation, Massachusetts Institute of Technology, 2004.

[24] M. Williamson, "Robot arm control exploiting natural dynamics," Ph.D. dissertation, Massachusetts Institute of Technology, 1999.

[25] R. Brooks, L. Aryananda, A. Edsinger, P. Fitzpatrick, C. Kemp, U.-M. O'Reilly, E. Torres-Jara, P. Varshavskaya, and J. Weber, "Sensing and manipulating built-for-human environments," *International Journal of Humanoid Robotics*, 2004.

[26] M. Zinn, O. Khatib, B. Roth, and J. Salisbury, "Playing it safe: human-friendly robots," *IEEE Robotics & Automation Magazine*, vol. 11, no. 2, pp. 12–21, June 2004.

[27] R. A. Brooks, "A robust layered control system for a mobile robot," vol. RA-2, pp. 14–23, Apr. 1986.

[28] P. Viola and M. Jones, "Robust real-time object detection," vol. 57, no. 2, pp. 137–154, 2004.

[29] A. Edsinger-Gonzales, "Developmentally guided ego-exo force discrimination for a humanoid robot," in *Fifth International Workshop on Epigenetic Robotics*, Nara, Japan, 2005.

[30] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. MIT Press, 2005.