

Generalizing Demonstrated Manipulation Tasks

Nancy S. Pollard^{1,3} and Jessica K. Hodgins^{2,3}

¹ Brown University, Providence RI 02912, USA

² Carnegie Mellon University, Pittsburgh PA 15213, USA

³ ATR Human Information Science Laboratories, Kyoto 619-0288, Japan

Abstract. Captured human motion data can provide a rich source of examples of successful manipulation strategies. General techniques for adapting these examples for use in robotics are not yet available, however, in part because the problem to be solved by the robot will rarely be the same as that in the human demonstration. This paper considers the problem of adapting a human demonstration of a quasistatic manipulation task to new objects and friction conditions (Figure 1). We argue that a manipulation plan is similar to a demonstration if it involves the identical number of contacts and if the applied contact wrenches follow similar trajectories. Based on this notion of similarity, we present an algorithm that uses the human demonstration to constrain the solution space to a set of manipulation plans similar to the demonstration. Our algorithm provides guarantees on maximum task forces and flexibility in contact placement. Results for the task of tumbling large, heavy objects show that manipulation plans similar to a demonstration can be synthesized for a variety of object sizes, shapes, and coefficients of friction. Experimental results with a humanoid robot show that the approach produces natural-looking motion in addition to effective manipulation plans.



Fig. 1. (Top) Human demonstration of a tumbling task. (Bottom) The demonstration has been adapted to a new object geometry and to the robot kinematics.



Fig. 2. Demonstration is a straightforward way to indicate how a task should be performed. Even for a task as simple as tumbling an object from one face to another, a variety of strategies can be observed. These subjects and the subject in Figure 1 have placed their hands very differently while performing this task. All objects are being tumbled from the right to the left.

1 Introduction

Developing robot manipulation skills from human demonstrations is appealing because of the variety of strategies that can be observed in human motion and the ease of conveying a strategy through demonstration (Figure 2). To make human demonstrations a practical way of describing new robot skills, however, we would like to have a system that can take as input a single example and generate behavior similar to that example, gracefully managing differences between the task that was demonstrated and task that the robot must actually perform. Figure 1 shows the type of situation we would like to handle; a significant difference exists between the geometry of the object manipulated by the human demonstrator and the geometry of the object manipulated by the robot.

For quasistatic manipulation tasks, both contact positions and applied forces are important. For this type of task we propose that similarity between a manipulation plan and an example should be based on contact wrenches, i.e., the forces and torques applied to the object at the contact points during the task. In this paper, a demonstration is used to constrain the solution space so that a planner considers only solutions that are similar to that demonstration. Figure 3 shows some results. Compare the contact wrenches applied to the different objects at any instant in time.

The main idea is to use an example to define the number of contacts required for the task and to constrain the roles of those contacts. These constraints can be powerful, and we show how they can be used to define equivalence classes of manipulation plans. Our approach is first to convert the manipulation-planning problem into a set of grasp-planning problems whose solution is a sufficient condition for a manipulation plan to exist. An example-based algorithm is presented to solve this set of grasp-planning problems and




Fig. 3. (Top) The human demonstration is abstracted as an object trajectory and a sequence of contact points. A plausible contact force trajectory is estimated from this information. Snapshots are at intervals of 0.33s. (Middle, Bottom) Our paper presents an algorithm for adapting an example to new objects with known geometries. One goal is that contact wrenches in the new manipulation plans should be similar to those in the example.

select contact points. Given these contact points, a manipulation plan is produced by computing target forces to be applied at the selected contact points over time. Figure 4 shows a block diagram of this approach. Features of this approach include the following:

- **Computational efficiency.** Each step in the process is polynomial in all parameters, including the number of contacts.
- **Task flexibility.** Significant variation in object geometry and friction conditions can be supported.
- **Flexibility in contact placement.** Given the geometry of an object to be manipulated, independent regions of contact are identified so that placing one contact somewhere within each region is sufficient to guarantee that a manipulation plan exists. Examples of such regions are shown in Figure 8.
- **Preservation of force-related features of the example.** Contact wrenches are constrained to a region that contains the example (Figure 3).
- **Quality guarantees.** The regions in Figure 8, for example, were produced with the requirement that the maximum sum of contact normal

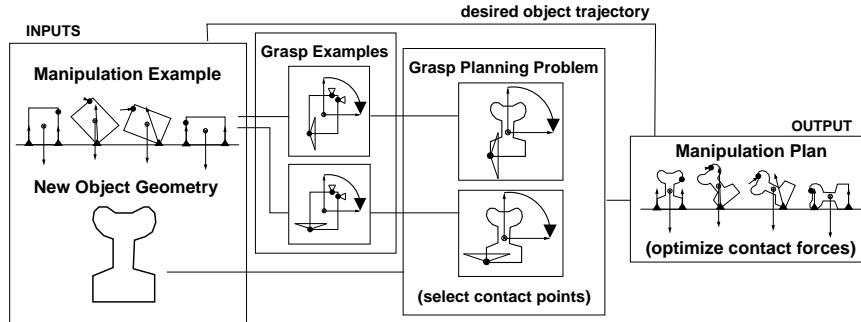


Fig. 4. Block diagram of the planner described in Section 3. The planner takes as input a manipulation demonstration and new object geometry. Contact selection is treated as an example-based grasp-planning problem. A trajectory for the new object is computed from the demonstration. Force computation for the new manipulation plan is accomplished by time sampling and optimization.

forces for manipulating the object should never be greater than twice that required for the example.

- For a humanoid robot, much of the **character of the original human motion** can also be preserved to create a natural-looking appearance (Figure 1).

This approach is suitable for manipulation tasks where the motion is not highly dynamic and where contacts with the object can be abstracted as non-sliding point contacts with friction. The active set of contacts may vary over the course of the motion.

2 Background

This paper draws from research in manipulation, manipulation planning, and grasp planning. Tasks such as tumbling [21] and pivoting [1] have been demonstrated for robot hands, and a number of researchers have explored manipulation planning for tasks similar to tumbling (e.g. [24] [7] [28] [13] [12] [29] [27]). Bicchi [4] provides an overview of grasp-planning research.

Of particular interest here is classification of contact regions. Trinkle and Paul [24] classify contact regions based on object response to squeezing and use these to plan motion to lift an object into a grasp, and Lynch [12] identifies regions where a contact point can be placed so that an object will be toppled by a moving fence. In the work closest to ours, Nguyen [14], Ponce and his colleagues [19] [18] [20] and Chen and Burdick [5] describe algorithms for computing independent contact regions for two-to-four-fingered grasps so that as long as one contact can be placed within each region, a force-closure grasp can be found. van der Stappen et al. [25] and Liu [11] describe how all force closure grasps can be computed for 2D grasps. Independent contact regions

could be extracted from the results of their algorithms. Our paper builds on many of these ideas to show how contact regions can be constructed, not from scratch as in previous work, but from an example, so that the contact regions represent a subclass of grasps similar to that example. The key algorithm is an adaptation of Pollard [16] from grasps to manipulation strategies and from frictionless contacts to contacts with friction.

Our work also draws on research in learning from demonstration. Previous work has used one or a combination of two basic approaches. The first is to extract a policy or control strategy directly from the demonstration [2] [10]. The second is to extract enough information about the goal and/or task that a planner or learning algorithm can compute such a policy [9] [8] [26] [22]. The two approaches can be combined. In this case, the planner somehow makes use of the detailed information observed in the human demonstration. For example, in learning a pendulum swingup task, Atkeson and Schaal [3] begin by tracking the hand trajectory of the human demonstrator. They note that merely tracking observed motion was not effective and present an iterative learning algorithm to successfully duplicate pendulum motion from the demonstration.

Our paper relies on a combined approach. In our case, the object trajectory and a set of contact points are extracted from the human demonstration. The example cannot be applied directly to a new object, because identical contact points may not be available; a new set of contact points must be found that are appropriate for the new object. Our approach differs from previous work in that the demonstration is not modified in a local way as in [3], but instead it is used to constrain the space of solutions considered by a global planner. This tactic is practical for tasks that are not highly dynamic. Our approach gives us many of the advantages of a global planner while keeping the results tied to an example and leading to a computationally efficient algorithm.

3 Generalizing a Manipulation Example

An example contains a trajectory for the manipulated object and a trajectory for each contact (Figure 3, top). Our problem is to adapt the example to a new object with known geometry, which requires identifying contact points and solving for contact force targets on the new object. The main idea pursued here is to convert the manipulation problem into a grasp-planning problem such that a solution to the latter is a sufficient condition for a solution to the former to exist. Section 3.1 describes how a manipulation problem is converted to a grasp-planning problem; Sections 3.2 and 3.3 describe how this problem is solved. The frictionless case is presented first, followed by an extension to contacts with friction. Section 3.4 details how these results are converted to a manipulation plan.

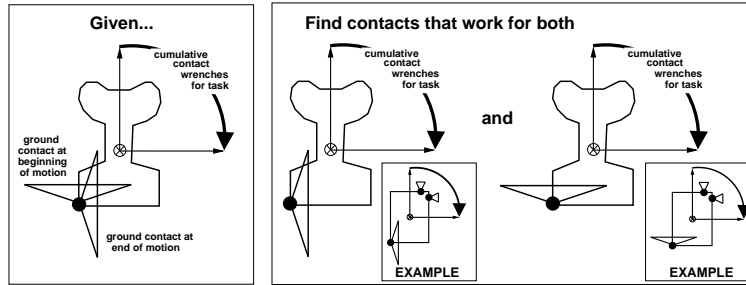


Fig. 5. Grasp-planning problem. Given new object geometry, task wrenches, and a pivot point, shown on the left, find contact points on the new object such that both grasps on the right can achieve all task wrenches. The small insets show the example grasps used as a reference.

3.1 From Manipulation Problem to Grasp-Planning Problem

We wish to convert a manipulation problem where the task is a function of time to a grasp-planning problem where time is not relevant. A first step is to model the task. For the quasistatic problem of tumbling the object 90 degrees to the left, task wrenches are pure forces through the object’s center of mass, spanning a 90 degree range as shown in Figure 5.

The next step is to solve for a grasp that can resist these task wrenches. For tumbling, the ground contact at the pivot point can be used as one contact; however, the range of forces available at the ground contact changes with time. If we consider this problem to be two-dimensional, then the ground contact is an edge-vertex contact. Friction cone constraints at this contact are a linear combination of the two extremes at the start and end of the motion. An example of friction cones available at extremes of a tumbling task is shown in Figure 5, left. A sufficient condition for a manipulation plan to exist is to find a set of contact points which, when combined with either extreme ground friction cone, will produce a grasp capable of resisting all task wrenches (Figure 5, right). If the same set of contact points works for each extreme separately, it will also work for intermediate configurations of the ground contact.

To solve the grasp-planning problem posed in Figure 5, we need to find a set of contact points that, for example, produces a force closure grasp with each of the two ground friction cones. If this were the only goal, the demonstration would be unnecessary. The goal that we have, however, is different; we wish contact wrench profiles in the new manipulation plan to be similar to those estimated from the demonstration.

Example grasps extracted from the demonstration can be used to constrain the search for grasps that achieve this goal. The example grasps are created by collecting all contact points observed in the manipulation example and creating grasps from these contact points plus the two extreme ground

friction cones. Example grasps for the demonstration explored in this paper are shown as insets in Figure 5. Each of these grasps is capable of achieving all required task wrenches. If this were not true, the task would have to be segmented into subtasks, each covering a smaller range of object motion.

3.2 Placement of Frictionless Contacts

The grasp-planning problem illustrated in Figure 5 must be solved to select a discrete set of contact points. Solving this problem requires several steps. First, independent contact regions are computed for each extreme ground friction cone. Second, these regions are intersected to create independent contact regions that will work for either grasp. Third, discrete contact points are chosen within the regions. After regions have been identified, intersecting them is trivial, and contact points can be selected as desired by the user (e.g., in the centers of the largest connected region or to achieve a kinematic goal such as similarity to the original human motion). The difficult part of the problem is to identify a set of independent contact regions. This section describes how regions are computed for a given ground contact friction cone extreme.

When the grasp is frictionless, the algorithm from Pollard [16] can be used to adapt a grasp to new object geometries by designing wrench space regions around the contact wrenches of the example grasp and then projecting these regions onto a new object. This approach is reviewed below, and Section 3.3 describes how it can be extended to contacts with friction.

Equivalence Classes of Contact Configurations. From a given example grasp, we wish to define equivalence classes of contact configurations on a new object. The information extracted from the example grasp is the particular role ascribed to each contact within that grasp. The idea is that if contacts on a new object can be found to play the same roles, the new grasp is in a way “the same as” the example. This idea can be quantified with the following construction.

Let $\mathbf{g}_i = [\hat{f}_i \quad \gamma(r_i \times \hat{f}_i)]^T$ be the set of contact wrenches for each contact i of the example grasp, formed from unit normal contact forces \hat{f}_i and resulting torque about the object’s center of mass, $(r_i \times \hat{f}_i)$, where r_i is the vector from object’s center of mass to the contact point. Torque is scaled by some suitable factor γ , for example the distance from the object center of mass to the pivot point.

Form the convex hull of \mathbf{g}_i and the zero wrench. Represent it as a collection of hyperplanes, each expressed as an outward pointing normal \mathbf{n}_j and distance from the wrench space origin d_j :

$$\{ [\mathbf{n}_1, d_1], [\mathbf{n}_2, d_2], \dots, [\mathbf{n}_H, d_H] \} = CH\{ \mathbf{0}, \mathbf{g}_1, \mathbf{g}_2, \dots, \mathbf{g}_N \} \quad (1)$$

There are N contact points in the example grasp and H facets in the convex hull. An example set of contact points and the convex hull of those points in




Fig. 6. An example isolates roles of the individual contacts as regions in wrench space. This example corresponds to a frictionless grasp of a circular disc with four evenly spaced contacts. (Left) Contact wrenches \mathbf{g}_i and the convex hull formed from those wrenches. (Center) Move hyperplanes so that they just contain the task wrench space, which is the black shaded area in the figure, oriented toward forces in the y direction. (Right) Scale hyperplanes by β/F_{max} , and highlight unit force wrench space regions that satisfy Equation 5. The bottom two contacts are associated with larger regions due to the anisotropy of the task.

a two-dimensional wrench space (a frictionless grasp of a 2D circular disc) is shown in Figure 6, left.

Index set ρ_i identifies hyperplanes associated with contact i :

$$\rho_i = \{j : (\mathbf{g}_i \cdot \mathbf{n}_j) = d_j\} \quad (2)$$

For example, in Figure 6 left, $\rho_1 = \{1, 2\}$, because \mathbf{g}_1 is contained in hyperplanes 1 and 2.

Adjust each hyperplane along its normal until it just contains all wrenches required for the given task. For a task that can be represented as a collection of points \mathbf{s}_k in wrench space:

$$d'_j = \max\left(0, \max_k (\mathbf{s}_k \cdot \mathbf{n}_j)\right) \quad (3)$$

where $(\mathbf{s}_k \cdot \mathbf{n}_j)$ is the inner product of wrench space vectors \mathbf{s}_k and \mathbf{n}_j . An example of this construction is shown in Figure 6 center. The task is the black shaded area, dominated by a range of forces in the positive y direction.

The maximum sum of contact normal forces required for this task, F_{max} , is then:

$$F_{max} = \max_{j=1}^H \frac{d'_j}{d_j} \quad (4)$$

The value of F_{max} is 1.41 in the example of Figure 6, because the topmost hyperplane must be scaled by this factor to contain the entire set of task wrenches within its interior.

Any new grasp, defined as the set of contact wrenches \mathbf{g}'_i , is in equivalence class β if it satisfies the following collection of inequalities for each contact i :

$$\left(\mathbf{g}'_i \cdot \mathbf{n}_j\right) \geq \frac{\beta}{F_{max}} d'_j \quad \forall j \in \rho_i \quad (5)$$

This equation states that wrench \mathbf{g}'_i is in the exterior or exactly on all of the hyperplanes that contained point \mathbf{g}_i in the convex hull formed from the example grasp, after these hyperplanes have been adjusted to fit the task and scaled by β/F_{max} . Contact wrench \mathbf{g}'_i is meant to correspond directly to original contact wrench \mathbf{g}_i . In other words, the two contacts are intended to play the same roles in their respective grasps. Parameter β allows limits to be placed on the sum of contact normal forces. Setting β to 0.5, for example, guarantees that in the new grasp the sum of contact normal forces will never be greater than twice that required in the example if the task wrenches are the same as those in the example [15]. Figure 6 right shows an example where β is 0.5.

Valid contact regions can now be “painted” onto a new object by sampling the object surface, computing the contact wrench associated with each surface point and testing whether Equation 5 holds. A different region will be computed for each contact of the original grasp. These regions are independent. As long as the new grasp has one contact in each region, we can guarantee that the grasp can exert the task wrenches with a sum of contact normal forces no greater than $1/\beta$ times those of the example grasp. Practically this guarantee is useful for keeping robot joint torques within reasonable limits.

Figure 7 shows an example of a 2D, four-contact, frictionless grasp with a 3D wrench space. The task is to slowly rotate the object 90 degrees to the left, and wrenches required to support the object during this task are pure forces through the object’s center of mass spanning a 90 degree range. After the grasp is processed as described in Equations 1 through 4, the set of valid wrenches for each contact, which is defined by Equation 5, can be computed and displayed. Contact force magnitude is arbitrary, and so it is restricted to one to allow contact wrenches to be plotted in a 2D space of force angle vs. torque in the center plots. Wrenches corresponding to frictionless point contact on the edges of any new object can be plotted in this same space. They appear as vertical lines, because force angle does not change as the contact point moves along the edge. Valid contact regions on the object occur when these vertical lines intersect the contact regions. The figures on the right show contact regions for the example object and an object with more complex geometry.

3.3 Placement of Contacts with Friction

The approach just described is not adequate for the example shown in Figure 1, or for the large majority of manipulation tasks because friction is




Fig. 7. Results of generalizing a four contact grasp in the frictionless case. (Left) The example grasp. (Middle) Valid wrenches for each contact, along with the vertical lines corresponding to frictionless point contact on the edges of the new object. (Right) Independent contact regions on the new object. Parameter β is 0.5 for these examples. The box is $0.32m \times 0.44m$, γ is 1, the second object is approximately the same size as the box.

important to the success of the strategy. With friction, what is required is that for each extreme of the original friction cone, *some* wrench within the new friction cone can be constructed to play the same role in the grasp. This problem can be solved using linear optimization as follows.

A friction pyramid approximation represents contact wrench \mathbf{g}_i as a linear combination of L extremes, where \mathbf{g}_i and its extremes $\mathbf{g}_{i,l}$ all have unit magnitude force components in the direction of the contact normal.

$$\mathbf{g}_i = \sum_{l=1}^L \alpha_l \mathbf{g}_{i,l}, \quad \sum_{l=1}^L \alpha_l = 1, \quad \alpha_l \geq 0, \quad l = 1, \dots, L \quad (6)$$

Form the convex hull of all extreme contact wrenches $\mathbf{g}_{i,l}$ and the zero wrench, and express this hull as a set of hyperplanes $[\mathbf{n}_j, d_j]$ as in the frictionless case:

$$\{ [\mathbf{n}_1, d_1], [\mathbf{n}_2, d_2], \dots, [\mathbf{n}_H, d_H] \} = CH\{ \mathbf{0}, \mathbf{g}_{1,1}, \dots, \mathbf{g}_{N,L} \} \quad (7)$$

Let index set $\rho_{i,l}$ be defined as follows:

$$\rho_{i,l} = \{ j : (\mathbf{g}_{i,l} \cdot \mathbf{n}_j) = d_j \} \quad (8)$$

Note that ρ has two indices: i indexes the original contact, and l indexes one of the samples of the friction pyramid at that contact.

Find d'_j and F_{max} as in Equations 3 and 4. Any new grasp, defined as a set of contact wrenches \mathbf{g}'_i , with friction pyramid samples $\mathbf{g}'_{i,m}$ is in equivalence

class β if the following set of expressions can be solved for all combinations of contacts $i = 1, \dots, N$ and friction cone extremes $l = 1, \dots, L$:

$$\left[\left(\sum_{m=1}^L \alpha_m \mathbf{g}'_{i,m} \right) \cdot \mathbf{n}_j \right] \geq \frac{\beta}{F_{max}} d'_j \quad \forall j \in \rho_{i,l} \quad (9)$$

$$\sum_{m=1}^L \alpha_m = 1 \quad (10)$$

$$\alpha_m \geq 0, \quad m = 1, \dots, L \quad (11)$$

The first expression states that for contact i in the example grasp, for each friction cone extreme l , proposed contact i on the new object must be able to generate a wrench that plays the same role in the grasp. The second expression maintains unit magnitude for the force component of \mathbf{g}'_i in the normal direction, and the third expression keeps all \mathbf{g}'_i within their friction cones.

For example, a proposed new grasp can be tested for membership in equivalence class β by solving this linear system using the simplex method. To do this, replace β with unknown parameter $\beta_{i,l}$ and maximize $\beta_{i,l}$ for each combination of indices i and l . The proposed grasp is in equivalence class β if:

$$\min_{i=1}^N \left(\min_{l=1}^L \beta_{i,l} \right) \geq \beta \quad (12)$$

As in the frictionless case, valid contact regions can be “painted” on the new object by sampling the object surface, attempting to solve the system in Equations 9 through 11 for a single contact i , for all friction cone extremes l , and checking that $\sum_l \beta_{i,l} \geq \beta$. All of the examples in Section 4 use this algorithm.

3.4 From Grasp-Planning Results to Manipulation Results

Finally, the manipulation strategy must be extracted from the grasp representation of the problem. Timing is obtained for the new motion by reusing the joint angle trajectories captured from the human actor at their original speed. This decision preserves the dynamics of the robot motion at the possible expense of the dynamics of the manipulated object.

The orientation trajectory of the new object, $\theta'(t)$, can be derived from the orientation trajectory of the example object, $\theta(t)$. Our goal was to duplicate orientation of the example as much as possible, while ensuring that the critical point where the object’s center of mass is just above the pivot occurs at the same time in each motion. This critical point is important because it is the




Fig. 8. (Left) Valid aspect ratios for two new objects are represented by a range of angles from pivot point to center of mass. (Right) The pivot point is marked with a solid triangle, and valid independent contact regions for the two contacts are shown. The hand/object coefficient of friction was set to 0.5, and object/ground coefficient of friction was set to 6. (The table was covered with foam, producing a high coefficient of friction.) Quality parameter β was set to 0.5. Note that these results are independent of scale. Length parameter L in the left hand diagram can be set to any value to reflect the actual size of the object to be manipulated.

point at which applied torque about the pivot point changes direction. The following expression is used:

$$\theta'(t) = \left[\psi + \frac{t}{T}(\phi - \psi) \right] \left(\theta(t) - \frac{\pi}{2} \right) + \frac{\pi}{2} \quad (13)$$

where angle θ is measured as the angle from horizontal of the vector from the pivot point to the object's center of mass so that the critical point occurs at $\theta = \pi/2$. Parameters ψ and ϕ are the required scale factors at the beginning and end of the motion, determined from object geometry, and T is the time required for the entire motion.

Forces are estimated by sampling time, computing object orientation from Equation 13, and optimizing contact forces for that orientation and the given contact positions. To optimize contact forces, the sum of normal forces applied at the hand/object contacts is minimized, ignoring the magnitude of the ground reaction force. A conservative estimate of the hand/object coefficient of friction is used, as solutions optimized in this way tend to fall on the extremes of the friction cone.

4 Results

Figures 8 and 9 show contact regions computed for new object geometries by generalizing from the given example. The sizes of the independent contact regions vary with geometry, aspect ratio (Figure 8), coefficient of friction




Fig. 9. (Left) Contact regions grow larger as the hand/object coefficient of friction (μ) increases. (Right) Contact regions grow larger as the acceptable amount of contact force ($1/\beta$ times the example) increases.



Fig. 10. Comparison of human and humanoid motion. Frames are spaced every 0.33s (top) and 0.67s (bottom). Object weight is between 4kg and 5.4kg.

(Figure 9 left), and maximum allowed contact force (Figure 9 right). Note especially in Figure 8 that the example can be adapted to any scale. The tumbling demonstration of Figure 1 could be used to develop a plan to manipulate a much smaller object with the fingertips of a hand.

Contact forces are computed through time sampling and optimization, as outlined in Section 3.4. Figure 3 shows results for two objects. Despite dramatic changes in object geometry, contact wrench trajectories are similar.

Figures 1 and 10 show comparisons of the original human motion to motion of the Sarcos humanoid robot at ATR [6] tumbling two objects: one similar to that manipulated by the human subject and one less similar. Videos can be found at <http://www.cs.brown.edu/people/nsp/papers/WAFRO2.html>. Implementation details can be found in the Appendix. The tumbling task was executed approximately a dozen times for the box without failures when the initial box placement was within a few centimeters of the expected position. This task was executed at half the speed of the human performance to reduce dynamic effects. Tumbling for the object in Figure 1 was slowed to quarter speed to obtain reliable performance.

5 Discussion

This paper describes an approach to adapting a single example of a quasistatic manipulation task to various object geometries and friction conditions. Significant flexibility was obtained both in geometry of the object and in choice of contact locations. Experimental results with a humanoid robot show that the resulting strategies can be executed reliably. They also retain some of the natural feel of the original human motion.

One natural question is “why bother with the example motion?” After the problem has been abstracted to the level of contact positions and forces, it lends itself easily to such alternative approaches as global optimization of contact configurations over time. Our informal observation is that globally optimal solutions frequently do not match strategies actually used by human subjects, perhaps because we have not adequately captured the optimization function, if any, in use by the human subject. We believe by combining planning and example-based approaches, some good features of both can be captured. Using analytic generalization gives us valid contact regions, which provide flexibility to avoid joint limits or other potential problems such as areas where the object is slippery or fragile. Because contact regions are constructed from an example, contact placement is often very similar to that of the example, which may be advantageous when the robot and the human demonstrator have similar kinematics and workspaces. By design of the algorithm, contact wrenches are constrained to be within some region containing those of the example, and so we have a much stronger argument for similarity of the results to the example than do planners which consider task goal only.

This approach will work for any object where Equations 9 through 11 can be solved for all contacts. These equations provide a mapping between a contact index and a set of acceptable object features, although this mapping is not in a practical form at present. Expressing these equations as sets of acceptable object features would be useful for evaluating the generality of a given strategy, for collecting/evaluating a motion database, and for selecting between strategies when faced with a new manipulation problem. We are working toward a practical mapping between wrench space regions and object features.

Computation times for the examples in this paper are on the order of milliseconds. However, sampling the object surface to find good contact regions will be more expensive for truly 3D objects. We are working on a projection algorithm along the lines of Ponce et al. [20] to address this problem.

More complex manipulation tasks will almost certainly require more structure and extensions to the techniques presented here. Manipulation tasks with many contact transitions (e.g., remove a hand and place it somewhere else) may need to be segmented into subtasks. Manipulation tasks with sliding contacts, which are common in human manipulation of objects, cannot be addressed with this approach as it stands. Complex contacts (e.g., contacts

in concavities and area contacts) seem to be similar to contact with friction but have not yet been explored.

We are excited about the results to-date and anticipate that the basic approach will prove general: that quasistatic manipulation can be treated as a single task domain, a library of example motions collected within this domain, and these examples adapted to new situations as needed, allowing us to take better advantage of the interesting variety of strategies exhibited in human behavior.

Acknowledgments

We would like to thank Mike Erdmann and Matt Mason for helpful discussions on this problem, Darrin Bentivegna, Gordon Cheng, Marcia Riley, and Stefan Schaal for assistance with the robot experiments, and Chris Atkeson for assistance in both areas. The robot and support for this work were provided by ATR Human Information Science Laboratories. Additional support was provided for Pollard by NSF grant CCR-0093072 and for Hodgins by NSF grant IIS-0196089.

Appendix: Implementation Details

The human demonstration was captured using a commercially available optical motion capture system from Vicon [23]. We used a standard marker set with 35 14mm markers for whole body motion, supplemented with two additional markers on each hand and three markers on the box. The additional markers made it possible to extract hand/object contact points during the motion.

The Sarcos robot at ATR (DB) [6] was used to demonstrate the example and new manipulation plans. To obtain joint trajectories for the robot, the raw marker data from the motion capture session was mapped to a human skeleton, scaled to the robot degrees of freedom and joint limits as described in [17], and adjusted using inverse kinematics to achieve and maintain contact points at desired locations on the object.

The control system used for these examples is shown in Figure 11. The input trajectory was specified as a set of kinematic parameters and a set of parameters relating to contact between the robot and the object. Specifically, inputs are robot joint angles and derivatives ($\theta_{DES}, \dot{\theta}_{DES}, \ddot{\theta}_{DES}$), box orientation ($\theta_{BOX,DES}, \dot{\theta}_{BOX,DES}$), contact forces (f_{ff}), and the ratio of contact force magnitudes to torque about the box pivot point (f_{ff}/τ_{BOX}). Two low-gain PD servos are used for control, one that uses visual feedback to maintain box state and one that uses position sensing at the joints to track the joint angle trajectory. The servo associated with box state controls the magnitude of feedforward contact forces to implement a torsional spring and damper about the box pivot point. Out-of-plane orientation errors are ignored. Correction




Fig. 11. Controller for the tumbling task. All inputs on the left may have to be adjusted to adapt the input trajectory to a new object and/or new friction conditions. This paper focuses on adjusting contact positions and contact force trajectories f_{ff} .

for such errors would make it possible to manipulate “thinner” objects (i.e., objects with smaller dimension along the axis of rotation).

References

1. Y. Aiyama, M. Inaba, and H. Inoue. Pivoting: A new method of grasplless manipulation of object by robot fingers. In *Proc. IEEE/RSJ Intl. Conference on Intelligent Robots and Systems*, 1993.
2. H. Asada and Y. Asari. The direct teaching of tool manipulation skills via the impedance identification of human motion. In *Proc. IEEE Intl. Conference on Robotics and Automation*, 1988.
3. C. G. Atkeson and S. Schaal. Robot learning from demonstration. In *International Conference on Machine Learning*, 1997.
4. A. Bicchi. Hands for dexterous manipulation and robust grasping: A difficult road toward simplicity. *IEEE Transactions on Robotics and Automation*, 16(6):652–662, 2000.
5. I-M. Chen and J. W. Burdick. Finding antipodal point grasps on irregularly shaped objects. *IEEE Transactions on Robotics and Automation*, 9(4):507–512, 1993.
6. DB. <http://www.his.atr.co.jp/cyh/>.
7. M. Erdmann. An exploration of nonprehensile two-palm manipulation. *International Journal of Robotics Research*, 17(5), 1998.
8. S. B. Kang and K. Ikeuchi. Toward automatic robot instruction from perception—Temporal segmentation of tasks from human hand motion. *IEEE Transactions on Robotics and Automation*, 11(5):670–681, 1995.
9. Y. Kuniyoshi, M. Inaba, and H. Inoue. Learning by watching: Extracting reusable task knowledge from visual observation of human performance. *IEEE Transactions on Robotics and Automation*, 10(6):799–822, 1994.

10. S. Liu and H. Asada. Transferring manipulative skills to robots: Representation and acquisition of tool manipulative skills using a process dynamics model. *Journal of Dynamic Systems, Measurement, and Control*, 114:220–228, June 1992.
11. Y.-H. Liu. Computing n-finger form-closure grasps on polygonal objects. *International Journal of Robotics Research*, 19(2):149–158, February 2000.
12. K. M. Lynch. Toppling manipulation. In *Proc. IEEE Intl. Conference on Robotics and Automation*, 1999.
13. K. M. Lynch and M. T. Mason. Dynamic nonprehensile manipulation: Controllability, planning, and experiments. *International Journal of Robotics Research*, 18(1):64–92, January 1999.
14. V. Nguyen. Constructing force-closure grasps. *International Journal of Robotics Research*, 7(3):3–16, June 1988.
15. N. S. Pollard. Parallel methods for synthesizing whole-hand grasps from generalized prototypes. Technical Report AI-TR-1464, MIT AI Lab., January 1994.
16. N. S. Pollard. Synthesizing grasps from generalized prototypes. In *Proc. IEEE Intl. Conference on Robotics and Automation*, April 1996.
17. N. S. Pollard, J. K. Hodgins, M. J. Riley, and C. G. Atkeson. Adapting human motion for the control of a humanoid robot. In *Proc. IEEE Intl. Conference on Robotics and Automation*, May 2002.
18. J. Ponce and B. Faverjon. On computing three-finger force-closure grasps of polygonal objects. *IEEE Transactions on Robotics and Automation*, 11(6):868–881, 1995.
19. J. Ponce, D. Stam, and B. Faverjon. On computing force-closure grasps of curved two-dimensional objects. *International Journal of Robotics Research*, 12(3):263–273, June 1993.
20. J. Ponce, S. Sullivan, A. Sudsang, J.-D. Boissonnat, and J.-P. Merlet. On computing four-finger equilibrium and force-closure grasps of polyhedral objects. *International Journal of Robotics Research*, 16(1):11–35, February 1997.
21. N. Sawasaki, M. Inaba, and H. Inoue. Tumbling objects using a multi-fingered robot. In *Proc. 20th ISIR*, 1989.
22. S. Schaal, C. G. Atkeson, and S. Vijayakumar. Scalable techniques for non-parametric statistics for real-time robot learning. *Applied Intelligence*, (in press).
23. Vicon Motion Systems. <http://www.vicon.com/>.
24. J. C. Trinkle and R. P. Paul. Planning for dexterous manipulation with sliding contacts. *International Journal of Robotics Research*, 9(3):24–48, June 1990.
25. A. F. van der Stappen, C. Wentink, and M. H. Overmars. Computing immobilizing grasps of polygonal parts. *International Journal of Robotics Research*, 19(5):467–479, May 2000.
26. R. M. Voyles, J. D. Morrow, and P. K. Khosla. Towards gesture-based programming: Shape from motion primordial learning of sensorimotor primitives. *Robotics and Autonomous Systems*, 22:361–375, 1997.
27. M. T. Zhang and K. Goldberg. Gripper point contacts for part alignment. *IEEE Transactions on Robotics and Automation*, in press.
28. R. Zhang and K. Gupta. Automatic orienting of polyhedra through step devices. In *Proc. IEEE Intl. Conference on Robotics and Automation*, 1998.
29. T. Zhang, K. Goldberg, G. Smith, R.-P. Berretty, and M. Overmars. Pin design for part feeding. *Robotica*, 19(6):695–702, 2001.