Planning pre-grasp manipulation for transport tasks

Lillian Y. Chang*   Siddhartha S. Srinivasa[†]   Nancy S. Pollard*[‡]

*The Robotics Institute, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA. `lillianc@ri.cmu.edu`
[†]Intel Labs Pittsburgh, 4720 Forbes Avenue, Pittsburgh, Pennsylvania, USA. `siddhartha.srinivasa@intel.com`
[‡]Computer Science Department, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA. `nsp@cs.cmu.edu`

Abstract— **Studies of human manipulation strategies suggest that pre-grasp object manipulation, such as rotation or sliding of the object to be grasped, can improve task performance by increasing both the task success rate and the quality of load-supporting postures. In previous demonstrations, pre-grasp object rotation by a robot manipulator was limited to manually-programmed actions. We present a method for automating the planning of pre-grasp rotation for object transport tasks. Our technique optimizes the grasp acquisition point by selecting a target object pose that can be grasped by high-payload manipulator configurations. Careful selection of the transition states leads to successful transport plans for tasks that are otherwise infeasible. In addition, optimization of the grasp acquisition posture also indirectly improves the transport plan quality, as measured by the safety margin of the manipulator payload limits.**

I. INTRODUCTION

Even the seemingly simple manipulation task of fetching an object involves a sequence of multiple *action components*, including the reaching approach for grasp acquisition and object transport to the goal location (Fig. 1). Usually there are several ways to accomplish each action due to kinematic redundancy in the manipulator as well as freedoms in the task specification. Examples of task freedoms are the graspable regions on the object [1] and the allowable object poses at the goal location [2].

In grasping movable objects, *pre-grasp manipulation* is possible due to the additional task freedom of the object pose in the environment. Pre-grasp manipulation adjusts the object pose prior to grasp acquisition by actions such as sliding or pivoting the object on the support surface. Studies of human manipulation strategies have shown that people use pre-grasp manipulation to adjust the object pose in lifting tasks [3], enabling more robust task completion and higher quality grasping postures at the time of object lifting [4].

In this paper, we augment the typical action sequence of grasp acquisition followed by object transport with an additional preparatory action component. This results in a complete manipulation sequence of pre-grasp manipulation, grasp acquisition, and post-grasp transport, illustrated in Fig. 1. To achieve this sequence, a planner needs to automatically decide on the extent of pre-grasp manipulation. This decision

Fig. 1. A fetching task requires a complete manipulation plan including grasp acquisition and transport. We present a method for incorporating pre-grasp manipulation, specifically object rotation, in the manipulation plan. The object adjustment resulting from pre-grasp manipulation increases the task success rate and improves the quality of the transport plan.

is complicated by the high-dimensional configuration of the manipulator, as well as the fact that the pre-grasp manipulation goal affects not only the pre-grasp action but also the allowable goals and start states of the acquisition and transport components.

We present a method that efficiently incorporates pre-grasp rotation in a transport task plan by focusing on the selection of a single key point in the task plan. This key point is the grasping configuration at the time of object lifting, which is the transition state between grasp acquisition and post-grasp transport. This state determines the amount of pre-grasp rotation, and optimizing only the transition state manipulator configuration dramatically reduces the search space compared to optimizing the complete motion plan. This simplification makes planning pre-grasp manipulation practical, and we found that empirically it is often sufficient to improve the quality of the load-bearing postures during transport. In addition, the focus on the transition state makes our method independent of the actual planners used for the individual action components.

For object poses unreachable by direct grasping without pre-grasp interaction, our autonomous plans with object adjustment increase the task success rate. This improvement in task feasibility reproduces results shown previously only for manually-programmed plans. Furthermore, for object poses that were already reachable with direct grasping, the pre-grasp interaction can improve the quality of the manipulation plan with respect to the safety margin of the manipulator payload limits.

The remainder of this paper is structured as follows:

First, Section II discusses related work in manipulation path planning. Section III describes the action components composing the complete transport task sequence, and Section IV presents our method for selecting transition states with low payload costs to plan pre-grasp manipulation. Experimental validation in simulation and on a physical platform are presented in Section V. Finally, Section VI discusses the method's limitations and directions for future work.

II. RELATED WORK

This work builds upon multiple concepts from motion planning and manipulation. The previous literature presents several planners for finding feasible plans of individual manipulation components in isolation.

Motion planning for grasp acquisition and object transport actions searches for a collision-free path for the manipulator to reach a specified goal configuration [5]. A large body of work has focused on modelling the manipulator's configuration space (c-space) [6] in a preprocessing step to find the free-space regions [7] or connected paths for a static environment [8–10]. Recent work has also developed several variants of sampling-based planning techniques, such as Rapidly-exploring Random Trees (RRTs) [11], for path-planning in dynamic environments. These studies have focused on finding manipulation plans for reach-to-grasp motion [12–15], transport from a starting configuration where the object is already grasped [2, 16, 17], or sequences of alternating reaching and transport paths [18–20].

Another area in the motion planning literature focuses on manipulation actions to reconfigure objects in the workspace, which is relevant to the pre-grasp interaction component. Methods for planning non-prehensile pushing of objects have been investigated [21–23], and recent work also demonstrated push-planning methods on humanoid robots [24–26]. Other possible pre-grasp actions are toppling and tumbling [27], re-grasping [20, 28], and whole-body pivoting maneuvers [29–31]. In the previous work, the reconfiguration of the object on the support surface is specified as the primary task goal, and the interdependency with other components of a higher level task is not considered.

In this work, we investigate how these individual actions can be combined sequentially in an overall task to acquire and transport an object. Specifically, we incorporate pre-grasp interaction as a preparatory step to achieve successful manipulation for a broader range of task conditions. Previously pre-grasp interaction for robot manipulation was demonstrated using manually-programmed actions for a specific object [4]. We extend the previous work by presenting a method for autonomous planning of transport manipulation tasks. Similar to the approach in Hauser et al. [25], our method is based on the selection of good transitions between actions to obtain a successful sequence.

III. MANIPULATION TRANSPORT TASK

In this paper we focus on the class of manipulation tasks which involve transporting an object between start and goal locations. For example, a domestic service robot may be commanded to fetch an object for a person or re-organize items to clean up a room. The complete manipulation plan for the transport task is a sequence of three action components (Fig. 1):

1) pre-grasp manipulation – interaction to adjust the object pose on the support surface,
2) grasp acquisition – free-space reaching motion to position the end effector for the grasp, and
3) post-grasp transport – movement to a new configuration with the object fixtured to the end effector.

We assume that the object becomes fixtured after the grasp acquisition reach and before the transport by a hand-closing action pre-defined for the object.

In this work, we focus on the selection of the *transition states* between the components rather than the planning techniques for the individual component actions. Next we describe each of the component actions and the existing planners that we use in our framework.

A. Pre-grasp manipulation

The pre-grasp manipulation adjusts the object from its presented configuration in the environment to a new configuration prior to grasp acquisition. In general, this interaction encompasses any 6-DoF re-configuration of a rigid object and higher-DoF changes for articulated or deformable objects. Actions with 3-DoF planar displacements, such as sliding a box over a table or dragging a water jug on a refrigerator shelf, are natural candidates for moving objects resting on horizontal support surfaces.

We focus specifically on pre-grasp rotation, which adjusts the object's 1-DoF orientation θ in the plane of the support surface [3, 4]. This approach limits the additional complexity that the pre-grasp interaction contributes to the search space. Our experiments show that even adjusting only the orientation can dramatically change whether a manipulator can reach the side handles of household objects such as cooking pans and pitchers.

Our simplified model assumes that each object can be pivoted about a fixed rotation axis perpendicular to the plane of the support surface. The pre-grasp rotation from the initial object pose θ_i to the target orientation θ_t is synthesized using the planner presented by Diankov et al. [15] for manipulation with caging grasps, described briefly here. For each object model of constrained motion about the rotation axis, the pre-computation phase generates a set of caging grasps given the geometry of the object and the end effector. The planning phase searches for a path in the arm configuration space which moves the object from specified start to goal orientations using caging grasps.

B. Grasp acquisition

The grasp acquisition component consists of the reaching motion which positions the end effector at a desired grasp pose $g \in SE(3)$ relative to the object pose. We assume that the object pose θ resulting from any pre-grasp rotation remains fixed throughout the grasp acquisition action. The end-effector shape is fixed in an open-hand preshape during
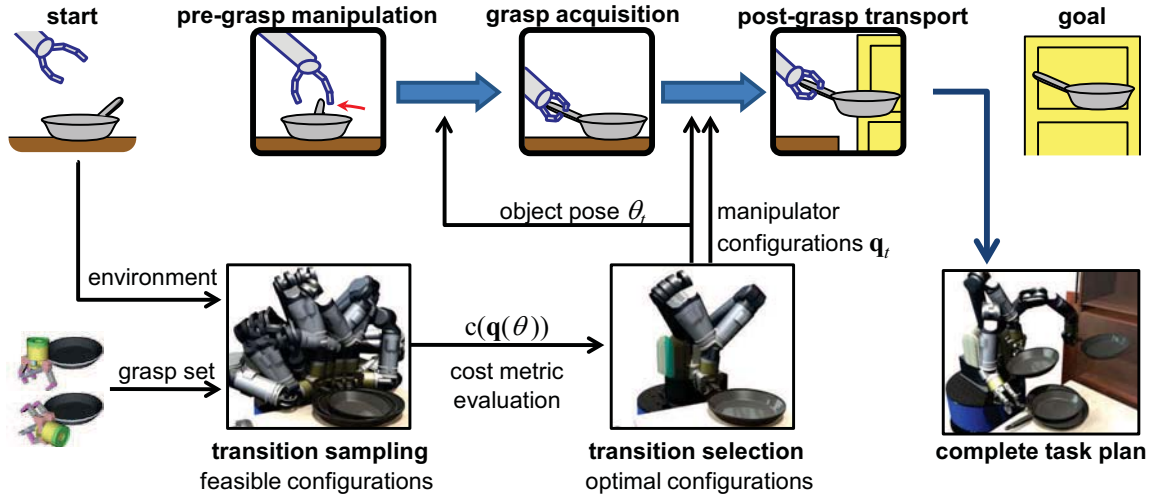
Fig. 2. Overview of the transition selection method for planning a complete object transport task. The specification of the individual action components requires selection of the target object orientation θ_t and the manipulator configuration $\mathbf{q}_t$ at the grasp acquisition point. The object orientation at grasp acquisition also determines the target orientation of the pre-grasp rotation. First, candidate manipulator configurations are generated by sampling IK solutions of the arm over multiple possible object poses, and only feasible non-colliding configurations are retained. Then, a single object pose is selected by evaluating the associated manipulator configurations by the safety margin relative to the manipulator payload limits. This selection results in a single target object pose θ_t and multiple low-cost manipulator configurations $\mathbf{q}_t(\theta_t)$ for the key grasp acquisition point before post-grasp transport.

reaching. Once the desired end-effector pose in the world frame $g(\theta)$ is reached, the hand is closed to fixture the object to the end effector.

We assume that a grasp set is pre-defined for the object. There are often multiple grasps of an object that are appropriate for a task, which adds flexibility for finding a feasible manipulation plan. A grasp set may be a group of individual grasp end-effector poses (points in the 6-DoF task-space) and/or as regions in task-space [2]. For planning the grasp acquisition component, we use a bi-directional RRT planner which allows multiple configuration goals to be specified as either c-space goals for the manipulator configuration or task-space goals of the end effector [32].

C. Post-grasp transport

The post-grasp transport brings the object from the grasp acquisition pose to the transport goal location. The object is fixtured to the end effector during the transport component such that there is no re-grasping. This is the only component action in the sequence where the object weight is supported by the manipulator rather than an environmental structure.

Often it is desirable to limit the object transport path to maintain nearly-upright orientation of the object, in order to avoid spilling or unsettling the contents of a container object. Even for the transport of empty or solid objects, maintaining a nearly-upright orientation may increase the predictability or natural appearance of a manipulator operating in a space shared by humans. We use the planner by Berenson et al. [32] to plan transport paths with the constraint to maintain, within a specified tolerance, nearly-upright object orientation. The transport goal is specified as a task-space region of allowable object poses in the environment.

IV. TRANSITION STATE SELECTION FOR PLANNING

Fig. 2 provides an overview of our approach for synthesizing a complete manipulation plan for the transport task. Two variables compose the key transition point at the time of grasp acquisition and object lifting:

1) Object pose θ_t – The target object orientation at acquisition, following any pre-grasp rotation. This determines the allowable poses $g(\theta_t)$ in the world frame for the end effector to grasp the object.
2) Manipulator configuration $\mathbf{q}_t$ – The robot arm configuration at acquisition. This state is also the start input to the post-grasp transport component.

As shown in the Fig. 2, the desired object pose for grasp acquisition θ_t also determines the target goal for the pre-grasp rotation action. Both θ_t and $\mathbf{q}_t$ must be selected automatically, since they are not specified by the overall task command to move the object to the final goal location.

Our method evaluates the candidate manipulator configurations $\mathbf{q}$ over multiple candidate object poses θ according to the payload safety margin. A summary score determines the target orientation θ_t based on the associated set of manipulator configurations $\mathbf{q}_t(\theta_t)$.

A. Configuration cost metric: payload safety margin

First we describe the cost metric for evaluating candidate transition states of the manipulator configuration. The cost c of a configuration $\mathbf{q}$ is inversely related to the manipulator payload p supportable at that configuration:

$$c(\mathbf{q}) = \frac{1}{p(\mathbf{q})} \qquad (1)$$

and has possible values in the range $[0, \infty]$. The overall manipulator payload p is computed from the load limits of the individual joints, τ_{max}, as follows. The joint torques $\hat{\tau}$

necessary to statically support a unit payload force $\hat{f}$ are computed from the Jacobian of the manipulator pose:

$$\hat{\tau}(\mathbf{q}) = J(\mathbf{q})^{\mathrm{T}}\hat{f} = J(\mathbf{q})^{\mathrm{T}}\begin{bmatrix} 0 \\ 0 \\ -1 \end{bmatrix}. \qquad (2)$$

The individual joint torque limits $\tau_{\max}$ are normalized by $\hat{\tau}$ to determine the maximum payload at the end effector $\mathbf{p}_{\max}(\mathbf{q})$ per joint j. The maximum payload of the configuration is limited by the minimum value over all the joints:

$$p(\mathbf{q}) = \min_j(\mathbf{p}_{\max j}(\mathbf{q})) = \min_j\left|\frac{\tau_{\max j}(\mathbf{q})}{\hat{\tau}_j(\mathbf{q})}\right|. \qquad (3)$$

The payload has possible values in the range $[0, \infty]$. For example, a configuration with a maximum payload of $p(\mathbf{q}) = 5$ N has a cost of $c(\mathbf{q}) = 0.2$ N^{-1}.

In essence, the cost metric reflects the safety margin of the object weight relative to the maximum payload. A low cost for a manipulator pose indicates a high safety margin for any constant object weight. It is desirable to maintain comfortable safety margins for a transport task in order to reduce wear and tear on the actuators, as well as to avoid exceeding payload limits due to errors in object weight estimation.

B. Sampling of candidate transition states

Before planning, multiple configurations $\mathbf{q}$ are sampled over multiple object orientations θ for the grasp acquisition point.

Candidate target object poses are generated in a straight-forward fashion by sampling the 1-DoF object orientation at even intervals. Object orientations in collision with environment obstacles are discarded.

At each sampled object orientation θ, corresponding grasp poses of the end effector $\mathbf{g}(\theta)$ are computed from the object-specific grasp set. A small finite grasp set can be completely sampled. For grasps specified as a continuous task-space region with respect to the object frame [2], the region can be sampled at discretized intervals for candidate grasps. The corresponding end-effector pose in the world frame $g(\theta)$ is computed from these grasps g expressed relative to the object pose θ.

The candidate manipulator configurations are IK solutions $\mathbf{q}(\mathbf{g}(\theta))$ that achieve the desired end-effector pose in the world frame. For the redundant manipulator in our implementation, we evaluate a finite set of IK solutions provided by an iterative solver in the OpenRAVE simulation software [33]. Only collision-free manipulator configurations are retained and evaluated according to the payload cost metric.

The result of the pre-planning sampling is a set of manipulator configurations $\mathbf{q}(\mathbf{g}(\theta))$ for each possible object orientation θ and their associated costs $\mathbf{c}(\mathbf{q}(\mathbf{g}(\theta)))$, or $\mathbf{c}(\theta)$.

C. Transition state evaluation

The transition selection limits the grasp acquisition to low-cost configurations that load the joints less relative to the payload limits. Because the object remains stationary during grasp acquisition, the selected manipulator configurations must all correspond to the same object orientation.

A summary score s evaluates a set of manipulator configurations at a given pre-grasp object orientation θ. The score considers only the best manipulator configurations in the feasible set according to a threshold percentile t of the individual payload costs. We prefer many goal configurations with modest payload costs over only a few goal configurations with exceptionally-low costs, such that the grasp acquisition planner can quickly find a solution among many candidates instead of failing to find a few particular solutions. Thus sets with a large number of low-cost configurations should have lower summary scores. We achieve this by normalizing the threshold percentile cost by the number of candidate manipulator configurations N:

$$s(\theta) = \frac{\text{percentile}(\mathbf{c}(\theta), \mathrm{t})}{N(\theta)}. \qquad (4)$$

The manipulator configurations meeting the threshold cost $\mathbf{q}_t(\theta)$ for the lowest scoring set are selected as the target postures for grasp acquisition. The associated object orientation

$$\theta_t = \arg\min_\theta s(\theta) \qquad (5)$$

is selected as the target object pose for pre-grasp rotation. These target transition variables θ_t and $\mathbf{q}_t$ determine the subgoals necessary for planning the individual action components described in Section III.

V. VALIDATION

A. Simulation experiments

The approach was tested on multiple transport task scenarios, shown in Fig. 3. The manipulator has a 7-DoF configuration, and the pose of the robot base in the world frame is fixed per example problem. The example tasks involve large household objects with handles, although pre-grasp rotation applies to any object which is not rotationally-symmetric. The handle axis of the skillet pan object is horizontal, while the kettle and watering can have vertical handle axes.

The object orientation θ is sampled at 10-degree intervals for 36 possible object poses. For each separate grasp region defined for the object, 3 end-effector poses are sampled within the bounds of the 6-D pose region. For the objects in our example scenarios, this corresponded to choosing a set of grasps per region that only differ in the orientation around the object handle axis.

In the evaluation of the object transport cost, the percentile threshold was set to $t = 10$ such that only the manipulator configurations $\mathbf{q}_t(\theta_t)$ with the lowest 10% costs for the target pose θ_t were selected as goals for the grasp acquisition component. When the pre-grasp manipulation planner fails to find a feasible plan to rotate the object, the object remains in its presented pose, and direct grasping is then attempted for object acquisition. In this case, the top 10% configurations for the initial presented object pose $\mathbf{q}_t(\theta_i)$, which have already been computed in the pre-planning phase, are selected

Fig. 3.   Transport task scenarios for validation experiments. (a) Putting the skillet pan away in the cabinet. (b) Moving the skillet pan from the table to the counter. (c) Moving the kettle from the table to the counter. (d) Bringing the watering can from the counter to the plant.
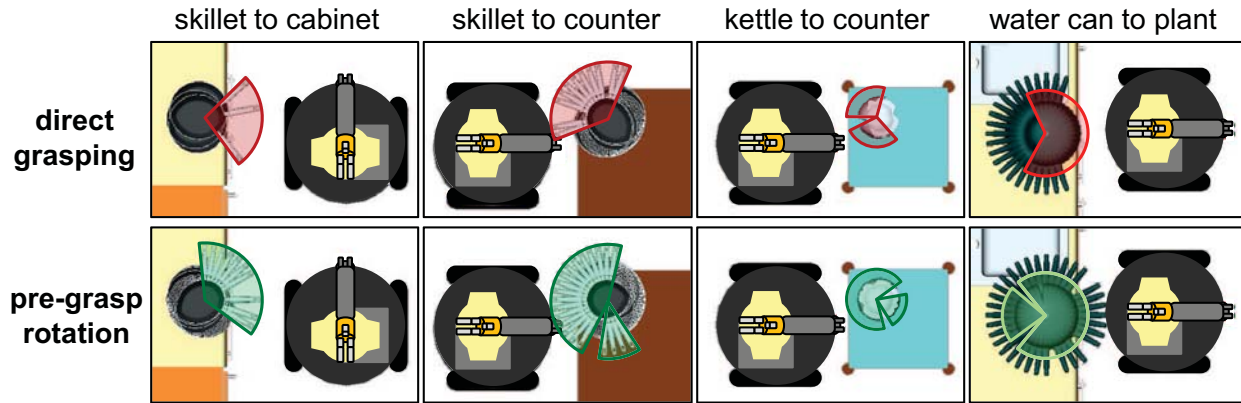

Fig. 4.   Comparison of plan success over different initial object orientations. The qualitative annotations highlight the general range of the handle orientations for which a transport plan was found. In the last example problem, the watering can spout extends further than the handle, thus the highlighted region is opposite the spout spokes.
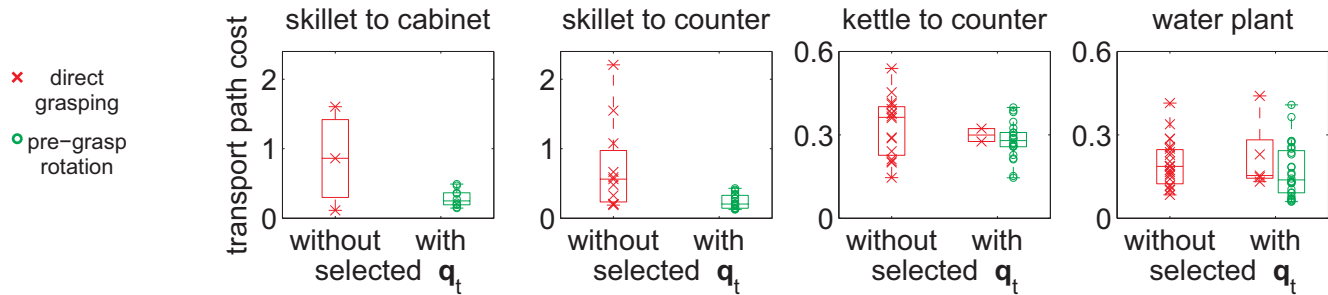

Fig. 5.   Comparison of transport path cost for the successful plans. Costs of the individual trials are denoted by × and ○, and the box plot overlay indicates the median and quartile values for each column. On the left of each plot, the costs of the direct grasping approach result from transport paths starting from the random, unrestricted goal configurations $\mathbf{q}$ of the grasp acquisition plan. On the right are the costs using the proposed approach, where the grasp acquisition goals are limited to selected configurations $\mathbf{q}_t$ with optimal costs. For the proposed approach, the results on the right are grouped by whether a pre-grasp rotation plan was found first or, if not, a direct grasping plan was found for the selected configurations. In the first two example problems, all trials with a successful transport plan incorporate pre-grasp rotation, so there were no direct grasping trials from the proposed approach with selected configurations. Note the difference in cost metric scales on the vertical axis.
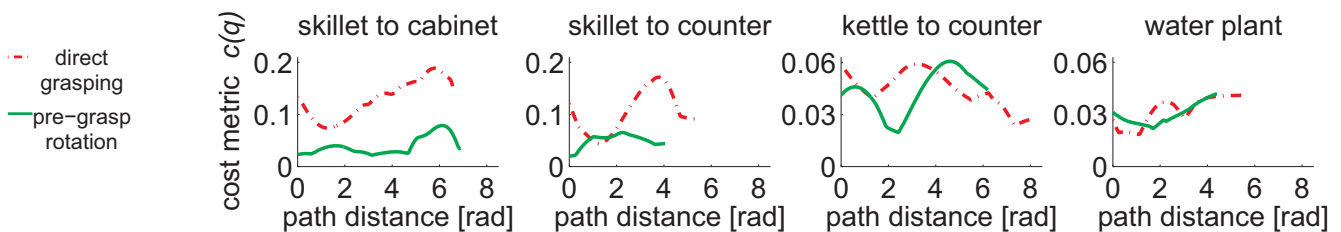

Fig. 6.   Evolution of the cost metric for the trials with median transport path cost. The path distance on horizontal axis is the Euclidean distance of the joint angles in the 7-DoF c-space. In the first two example problems, there is a large range of cost values, and paths which are initialized at a low cost configuration remain low cost. In the last two example problems, there is less variation in the cost of the initial configuration but the c-space path length for the pre-grasp rotation plan is shorter. Note the difference in cost metric scales on the vertical axis.

as the goals for grasp acquisition. The pre-grasp rotation planner fails after a maximum of 200 node expansions for the Randomized A* search of the caging grasp planner by Diankov et al. [15].

We compare our method to a typical direct grasping approach which does not include pre-grasp manipulation. In the direct grasping approach, there is no pre-planning evaluation of candidate configurations and the grasp acquisition is always attempted for the presented object pose without adjustment. The goal for the grasp acquisition is specified as the entire task-space region of allowable end-effector poses relative to the object. There are no restrictions on the quality of the grasp acquisition configuration.

For both methods, if no grasp acquisition plan is found, then no transport plan is attempted. The planner parameters for maximum iterations, maximum time, and smoothing iterations were identical for both approaches. The maximum planning time was limited at 60 seconds for both the grasp acquisition RRT planner and the object transport RRT planner with constraints [32]. The tolerance for upright object orientation was ±10 degrees for the pan object and ±6 degrees for the kettle and watering can.

We evaluate the successful plans according to the payload cost metric and the path length. Only the post-grasp transport path is evaluated since it is the only component action where the manipulator supports the weight of the object. Because the planner output may have nodes which are unevenly spaced, the paths are first re-sampled by linear interpolation in the 7-DoF c-space. The total path cost is the integral of the safety margin cost metric over the path length, which we approximated by the sum of the costs over all the re-sampled nodes weighted by the inter-node path distances.

B. Simulation results

For all four example problems, incorporating pre-grasp rotation increased the range of initial task conditions where successful transport plans were found (Fig. 4). In general, direct grasping without pre-grasp rotation was only successful when the object was presented in the environment with the handle oriented toward the robot. In the third example problem, direct grasping sometimes failed even when the kettle handle was oriented directly towards the robot, which required the end effector to be too close to the robot base. Pre-grasp manipulation augmented the performance of direct grasping by adjusting objects toward the target pose with the optimal set of manipulator configurations.

In addition to increasing task feasibility, our approach of optimizing the single grasp acquisition point also indirectly improved the path cost of the overall transport component for some scenarios. Fig. 5 shows the path cost results, grouped by whether pre-grasp rotation or direct grasping was used for acquisition. The greatest gains occurred for the first two example problems for transporting the skillet pan. The grasp set for the skillet had the largest range of possible end-effector poses around the handle, such that there was large variation in the path costs for the random manipulator configurations selected by direct grasping. In

Pre-grasp rotation plans	Increase in planning time relative to direct grasping	
	Median time	Maximum time
Offline precomputation		
Poses reachable by direct grasping	16%	44%
Poses unreachable by direct grasping	29%	56%
Online sampling		
Poses reachable by direct grasping	214%	137%
Poses unreachable by direct grasping	153%	128%

the kettle transport example, the path costs using the optimal configuration set were lower on average and less varied than those from direct grasping. In the plant watering example, pre-grasp rotation with selected configurations resulted in the lowest path costs, but the range of costs were similar to those from direct grasping.

The difference in the path cost results is further illustrated in Fig. 6, which shows examples of the cost metric evolution per configuration along the motion path. For the skillet pan, paths starting in a low-cost configuration tended to remain in low-cost configurations over the transport path. For the other example problems, there was smaller variation in configuration costs, but the overall path cost is lower partially due to shorter path length.

The improvement of the proposed method does require more planning time due to the additional component of pre-grasp manipulation. Successful direct grasping plans with only grasp acquisition and transport required a median and maximum planning time of 41 seconds and 81 seconds, respectively, over all the example problems. Table I shows the percent increase in total planning times for pre-grasp rotation, grouped according to whether the object pose was reachable by direct grasping for the same initial object pose. When sampling and evaluation of transition states are completed in an offline pre-computation stage, the median planning time increased by 16% and 29% for pre-grasp rotation when the object pose was, respectively, reachable and unreachable by direct grasping. Maximum planning times increased by 44% and 56%, indicating the difficulty of finding a feasible plan for challenging task conditions. When the pre-planning sampling and evaluation are considered part of online planning, the median and maximum planning times both increase over 100%. This overhead for online sampling of candidate configurations could be decreased by sparser sampling of object poses and grasp choices. Note also that the increased planning time is less relevant when the object was unreachable with direct grasping, because the success in finding a feasible plan is prioritized over the expense of the pre-grasp interaction.
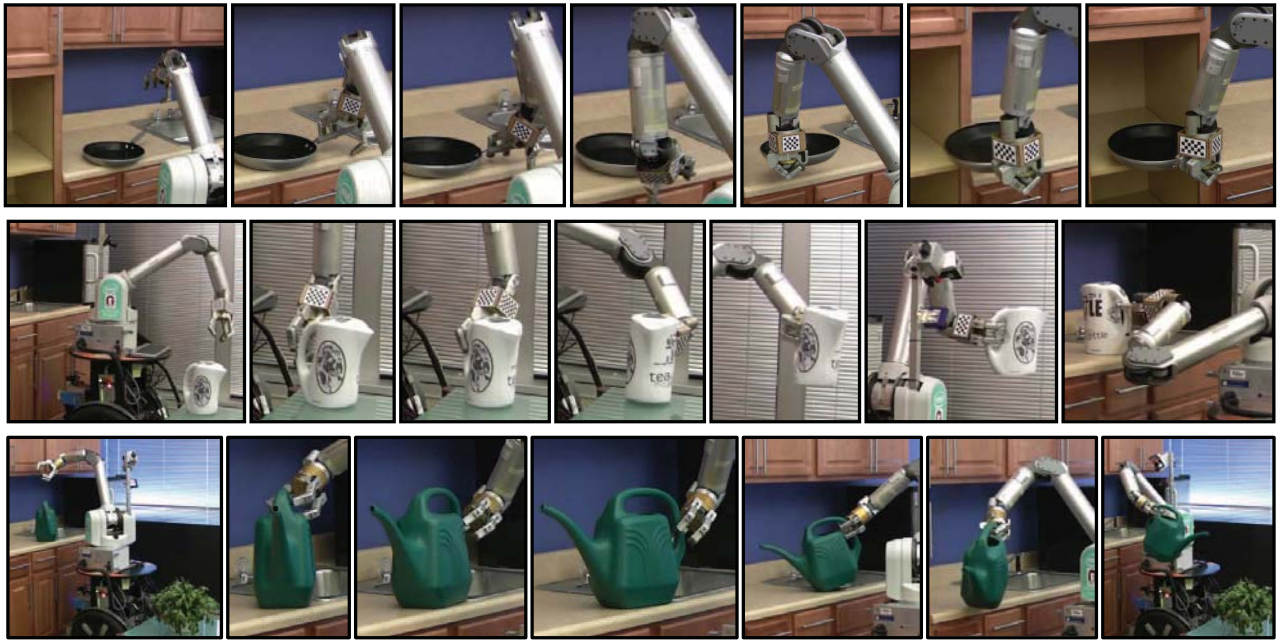
Fig. 7.   Example transport task plans demonstrated on the physical robot.

C. Physical demonstration

We demonstrated successful transport task plans on HERB [34], a manipulator platform consisting of the BarrettWAM 7-DoF arm and the BarrettHand end effector. Fig. 7 shows example sequences of the successful plans. Please also see this paper's accompanying video, as well as other videos available at:

http://personalrobotics.intel-research.net/
projects/pregrasp.php.

The motion plans were executed open-loop on the robot, after the object was placed in the demonstration environment according to the simulated task conditions. All plans were sensitive to the correct placement of the object to obtain successful contact conditions. However, successful physical demonstrations were possible for each example object, showing that our simplified model of object pre-rotation yields plausible manipulation plans.

VI. DISCUSSION

We have presented a method for incorporating pre-grasp rotation automatically in planning a complete manipulation sequence for transport tasks. Our approach extends the manipulator performance to find feasible plans for a wider range of task conditions compared to direct grasping. It also favors pre-grasp actions which result in low-cost manipulator configurations that have high payload safety margins. The pre-planning stage samples and evaluates only the transition state of the grasp acquisition configuration. The approach thus considers the key interaction between the component actions without needing to plan in the high-dimensional space of all possible manipulation path sequences.

An interesting result of our study is that the selection of a single point was often sufficient to improve the overall transport path, even though the planner for the transport component was completely agnostic to the payload cost metric. We believe this is because the limiting joints determining the maximum payload are often the distal joints at the manipulator "wrist." Due to the constraint to maintain nearly-upright object orientations during transport, the wrist configuration remains similar to that of the selected transition state determining the start of the transport path.

The limitations of this initial approach for finding complete manipulation plans suggest several directions for future extensions. First, this work restricted the pre-grasp interaction to 1-DoF rotation about a fixed pivot axis. This simple kinematic model limited the number of candidate transition states in the search, and it resulted in plausible plans which were successful on the physical platform. A limited amount of manual experimentation was sufficient to locate an approximate fixed axis for modeling our example objects with pure rotation and negligible translation. However, other objects may require modeling the pushing dynamics or sensing the resultant displacement. In addition, increasing the degrees of freedom for the pre-grasp interaction will require efficient sampling of the pre-grasp transition states. While we selected the globally optimal pre-grasp transition state, it may be necessary to instead limit the evaluation to local states near the initial presented object pose.

Further exploration of different path cost metrics is warranted for a broader set of applications. We evaluated our plans according to the payload cost metric because of its relevance to the load-bearing component of the transport task. For other tasks, it may be desirable to use pre-grasp interaction to improve the quality of the overall plan with respect to other posture-dependent metrics.

Future refinements could also investigate heuristics for

deciding when to optimize the transition state. Our results illustrated cases where the configuration costs have small variation across the grasp set. The distribution of costs at the pre-planning evaluation stage could be used to selectively plan for pre-grasp rotation only in cases with potential for large improvements in path cost. Furthermore, in applications where path cost is of particular priority, direct modification of the individual component planners is likely to be more effective than only indirect influence through transition selection.

Finally, robust execution in the presence of modeling uncertainty is a relevant issue but is beyond the scope of this work. We demonstrated the example transport plans in open-loop execution and used a high-fidelity simulation model of the environment and object geometry. The sensitivity to the placement of the object and robot illustrate the need for accurate localization methods for manipulation. Sensing feedback and local re-planning are ways to address the expected challenges that arise with approximate object models, localization uncertainty, and errors in execution of motor actions.

References

[1] M. Gienger, M. Toussaint, and C. Goerick, "Task maps in humanoid robot manipulation," in *Proc. IEEE Conf. Intelligent Robots and Systems (IROS)*, 2008, pp. 2758–2764.

[2] D. Berenson, S. S. Srinivasa, D. Ferguson, A. C. Romea, and J. Kuffner, "Manipulation planning with workspace goal regions," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2009, pp. 618–624.

[3] L. Y. Chang, G. J. Zeglin, and N. S. Pollard, "Preparatory object rotation as a human-inspired grasping strategy," in *Proc. IEEE Int. Conf. Humanoid Robots (Humanoids)*, 2008, pp. 527–534.

[4] L. Y. Chang, R. L. Klatzky, and N. S. Pollard, "Selection criteria for preparatory object rotation in manual lifting actions," *J. Motor Behavior*, vol. 42, no. 1, pp. 11–27, Jan 2010.

[5] J.-C. Latombe, *Robot Motion Planning*. Norwell, MA, USA: Kluwer Academic Publishers, 1991.

[6] T. Lozano-Perez, "Spatial planning: A configuration space approach," *IEEE Trans. Computers*, no. 2, pp. 108–120, Feb 1983.

[7] ——, "A simple motion-planning algorithm for general robot manipulators," *IEEE J. Robot. Automat.*, vol. 3, no. 3, pp. 224–238, Jun 1987.

[8] J. M. Ahuactzin, K. Gupta, and E. Mazer, "Manipulation planning for redundant robots: A practical approach," *Int. J. Robotics Research*, vol. 17, no. 7, pp. 731–747, Jul 1998.

[9] L. E. Kavraki and J.-C. Latombe, "Randomized preprocessing of configuration space for path planning: articulated robots," in *Proc. IEEE Conf. Intelligent Robots and Systems (IROS)*, vol. 3, 1994, pp. 1764–1771.

[10] L. E. Kavraki, P. Svestka, J.-C. Latombe, and M. H. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE J. Robot. Automat.*, vol. 12, no. 4, pp. 566–580, 1996.

[11] S. M. LaValle and J. J. Kuffner, *Robotics: The Algorithmic Perspective. Int. Workshop on the Algorithmic Foundations of Robotics*, Hanover, NH, 2000, ch. Rapidly-exploring random trees: Progress and prospects.

[12] D. Bertram, J. Kuffner, R. Dillmann, and T. Asfour, "An integrated approach to inverse kinematics and path planning for redundant manipulators," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2006, pp. 1874–1879.

[13] M. VandeWeghe, D. Ferguson, and S. Srinivasa, "Randomized path planning for redundant manipulators without inverse kinematics," in *Proc. IEEE Int. Conf. Humanoid Robots (Humanoids)*, 2007, pp. 477–482.

[14] D. Berenson, R. Diankov, K. Nishiwaki, S. Kagami, and J. Kuffner, "Grasp planning in complex scenes," in *Proc. IEEE Int. Conf. Humanoid Robots (Humanoids)*, 2007, pp. 42–48.

[15] R. Diankov, S. S. Srinivasa, D. Ferguson, and J. Kuffner, "Manipulation planning with caging grasps," in *Proc. IEEE Int. Conf. Humanoid Robots (Humanoids)*, 2008, pp. 285–292.

[16] K. Yamane, J. Kuffner, and J. K. Hodgins, "Synthesizing animations of human manipulation tasks," *ACM Trans. on Graphics (Proc. SIGGRAPH)*, vol. 23, no. 3, pp. 532–539, Aug 2004.

[17] D. Berenson, S. S. Srinivasa, D. Ferguson, and J. J. Kuffner, "Manipulation planning on constraint manifolds," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2009, pp. 625–632.

[18] T. Simeon, J. Cortes, A. Sahbani, and J. Laumond, "A manipulation planner for pick and place operations under continuous grasps and placements," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, vol. 2, 2002, pp. 2022–2027.

[19] A. Sahbani, J. Cortes, and T. Simeon, "A probabilistic algorithm for manipulation planning under continuous grasps and placements," in *Proc. IEEE Conf. Intelligent Robots and Systems (IROS)*, vol. 2, 2002, pp. 1560–1565.

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

[21] T. Lozano-Perez, M. Mason, and R. H. Taylor, "Automatic synthesis of fine-motion strategies for robots," *Int. J. Robotics Research*, vol. 3, no. 1, pp. 3–24, 1984.

[22] M. T. Mason, "Mechanics and planning of manipulator pushing operations," *Int. J. Robotics Research*, vol. 5, no. 3, pp. 53–71, 1986.

[23] K. M. Lynch and M. T. Mason, "Controllability of pushing," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 1995, pp. 112–119.

[24] M. Stilman and J. J. Kuffner, "Navigation among movable obstacles: Real-time reasoning in complex environments," *Int. J. Humanoid Robotics*, vol. 2, no. 4, pp. 479–503, 2005.

[25] K. Hauser, V. Ng-Thow-Hing, and H. Gonzalez-Banos, "Multi-modal motion planning for a humanoid manipulation task," in *Proc. Int. Symp. Robotics Research (ISRR)*, 2007.

[26] V. Ng-Thow-Hing, E. Drumwright, K. Hauser, Q. Wu, and J. Wormer, "Expanding task functionality in established humanoid robots," in *Proc. IEEE Int. Conf. Humanoid Robots (Humanoids)*, 2007, pp. 136–142.

[27] K. Lynch, "Toppling manipulation," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, vol. 4, 1999, pp. 2551–2557.

[28] P. Tournassoud, T. Lozano-Perez, and E. Mazer, "Regrasping," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, vol. 4, 1987, pp. 1924–1928.

[29] E. Yoshida, M. Poirier, J.-P. Laumond, R. Alami, and K. Yokoi, "Pivoting based manipulation by humanoids: a controllability analysis," in *Proc. IEEE Conf. Intelligent Robots and Systems (IROS)*, 2007, pp. 1130–1135.

[30] E. Yoshida, M. Poirier, J. P. Laumond, O. Kanoun, F. Lamiraux, R. Alami, and K. Yokoi, "Whole-body motion planning for pivoting based manipulation by humanoids," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2008, pp. 3181–3186.

[31] E. Yoshida, M. Poirier, J.-P. Laumond, O. Kanoun, F. Lamiraux, R. Alami, and K. Yokoi, "Pivoting based manipulation by a humanoid robot," *Autonomous Robots*, vol. 28, no. 1, pp. 77–88, Jan. 2010.

[32] D. Berenson, J. Chestnutt, S. S. Srinivasa, J. J. Kuffner, and S. Kagami, "Pose-constrained whole-body planning using task space region chains," in *Proc. IEEE Int. Conf. Humanoid Robots (Humanoids)*, 2009, pp. 181–187.

[33] R. Diankov and J. Kuffner, "Openrave: A planning architecture for autonomous robotics," Robotics Institute, Pittsburgh, PA, Tech. Rep. CMU-RI-TR-08-34, 2008.

[34] S. S. Srinivasa, D. Ferguson, C. Helfrich, D. Berenson, A. C. Romea, R. Diankov, G. Gallagher, G. Hollinger, J. J. Kuffner, and M. VandeWeghe, "HERB: a home exploring robotic butler," *Autonomous Robots*, vol. 28, no. 1, pp. 5–20, Jan. 2010.