Pre-grasp interaction as a manipulation strategy for movable objects

Lillian Y. Chang

CMU-RI-TR-10-28

Submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

The Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213
United States

December 2010

Thesis Committee:
Nancy S. Pollard, Chair
Christopher G. Atkeson
Jessica K. Hodgins
John M. Hollerbach

This research was supported, in part, by a National Science Foundation Graduate Research Fellowship and a NASA Harriett G. Jenkins Pre-doctoral Fellowship. Additional funding was provided by National Science Foundation through grants IIS-0326322, ECS-0325383, and CCF-0702443.

Abstract

Robotic systems have yet to match humans in skill for movement planning and tool manipulation. For example, humans can robustly grasp and manipulate objects even under task variation. However, successful grasping methods for robotic manipulators are often limited to structured environmental conditions. Our dual goals are to understand manipulation actions in humans and to add such skills to a robot manipulator's repertoire. In particular, we examine strategies for object acquisition, which is a common first component in manipulation actions.

Many approaches to automating robot motion for object acquisition have focused on reach-to-grasp tasks, where the arm motion and hand configuration are planned for grasping an object. With these solutions, the object placement often remains fixed in the environment until the object is carefully grasped from its presented configuration. In contrast, humans often take advantage of an object's movability to reorient and regrasp an object during the acquisition process.

This thesis investigates how such *pre-grasp interaction* can improve grasping through preparatory manipulation of the object's configuration. Specifically we studied the strategy of pre-grasp object rotation for grasp acquisition prior to a transport task. First, we examined human performance of the pre-grasp rotation strategy. A larger amount of pre-grasp object rotation correlated to a greater lifting capability, or maximum payload, of the grasping posture used at the time of object acquisition. In addition, when the task was more difficult due to increased object mass or increased upright orientation constraints, there was decreased variability in the object orientation selected for grasping. Second, we developed and evaluated a method for planning pre-grasp rotation for a robot manipulator. Our results show that the pre-grasp rotation strategy can improve a robot's manipulation capabilities by both extending the effective workspace for a transport task and improving the quality of the transport action.

To my family

Acknowledgments

It is my pleasure to thank the many people who have been part of my journey.

Above all, I must thank my advisor, Nancy Pollard, for her steadfast support during my graduate career. I am fortunate to have had Nancy's thoughtful guidance throughout the research process and her encouragement as I explored multidisciplinary interests. I would also like to thank my committee members Jessica Hodgins, Chris Atkeson, and John Hollerbach for their mentorship. Jessica generously shared her resources for graphics and motion capture support, and both Jessica and Chris welcomed me to their lab groups. John enthusiastically participated as my external reviewer, often making the effort to meet in person for discussion.

Sidd Srinivasa provided many suggestions that helped shape this work, as well as a computer in his lab and time on his robot. Not the least was Sidd's encouragement between the proposal and thesis completion. I am grateful to Garth Zeglin for his discussion during the proposal development, pep talks when I navigated difficult turns, and assistance with lab hardware and visitor demonstrations. Thanks also go to Sonya Allin and Bambi Brewer for sharing their experience and advice at various stages of the doctoral program.

I would also like to acknowledge Howard Seltman for his guidance on statistical testing and Roberta Klatzky for helping with the experimental design and analysis of the human studies. Dmitry Berenson and Rosen Diankov shared their motion planning expertise and code. Justin Macey assisted with the motion capture experiments, and Moshe Mahler modeled and rendered the robot simulation scenes. Jen Turken and Suzanne Muth gladly helped with logistics.

Many others contributed to the camaraderie in the Robotics Institute community, and several people moreover volunteered their manipulation skills for observation. I will fondly remember my fellow students and especially my first-year cubemates Peter Barnum and Kevin Yoon, whose mischief around and outside the office brought many smiles. I also thank the friends from my former communities who encouraged me from afar.

My family has my deep gratitude for their unwavering support. I owe my most heartfelt appreciation to my best friend for sharing this adventure together.

Contents

List of Figures

List of Tables

1 Introduction

1.1 Motivation

Dexterity is skillful interaction with the world. Dexterous manipulation may entail understanding physical properties of the environment, synthesizing robust policies for motor control, or achieving elegance and grace of movement. Much is still unknown about how to achieve dexterity and robust movement planning, both for human and robotic manipulation systems. Our dual goals are to understand manipulation skills in humans and to add such skills to the repertoire for robot manipulators.

Humans effortlessly manipulate objects as part of their daily interaction with the physical environment. Activities of daily living such as grooming, cooking, and eating require object acquisition by reaching and grasping, as well as in-hand manipulation by fine finger movements while using a tool. At any stage of a manipulation task, human performance requires selection of movement aspects such as arm posture, object contact, and joint compliance. Determining what metrics drive the selection of these components contributes to an understanding of human motor control.

For robotic systems, studying human manipulation strategies offers potential benefits for both functional and perceptual purposes. First, human manipulation provides examples of working action strategies that could be transferred to a robotic manipulator to increase its functionality. Applications that would benefit from increased functional dexterity include prosthetic devices and service robots. As another example, in space exploration, robotic manipulators may be required to operate autonomously in environments which are unfit for human operators and where teleoperation is precluded by long communication delays. Second, imitation of human strategies may provide perceptual benefits for robots which interact with humans. It is desirable for humanoid robots to perform manipulation tasks in a human-like way, either for social acceptability or predictability of motions in a cooperative task. Equipping a teleoperated manipulator with a library of automated sub-routines based on human strategies may also increase ease of control by a human operator.

This thesis investigates a specific type of manipulation skill for robust object acquisition: *pre-grasp interaction*. We observed and analyzed the pre-grasp interaction strategy in the context of human manipulation. We also developed a method for adopting this strategy in robotic manipulation and evaluated how it improves a robot's functional performance in object acquisition tasks.

1.2 Pre-grasp interaction

Dexterity includes a rich variety of manipulation skills. Examples include fixturing, transporting, tumbling, and throwing. The first phase of many tasks is object acquisition by grasping. Strategies for object acquisition in natural environments need to be robust to variations in task conditions such as object geometry, weight, and configuration in the environment.

Reaching and grasping for pick-and-place operations represents just one class of strategies for object acquisition. A trajectory for the manipulator arm motion and the hand configuration must be determined for reaching and grasping an object. In previous investigation of grasping, the object placement is often considered fixed in the environment and before the object is carefully grasped from its presented configuration. A novel initial object placement may necessitate a completely new plan for the reach-and-grasp action, including both the arm trajectory and the hand configuration for grasping.

However, the presented configuration of the object in the world could be suboptimal with respect to the desired grasp for a specified task. The object placement in the world could be suboptimal because a desired grasp contact point is unexposed, as in the case of a shelved book where only the spine is visible (Fig. 1.1(a)) or the case of thin coins resting flat on a table surface. In addition, obstacles in a cluttered environment may preclude a collision-free manipulator configuration for grasping the object. Even in the case where the desired contacts are reachable, the object placement may be suboptimal in terms of the manipulability or payload capacity of the manipulator pose required for grasping.

In several scenarios, the object has degrees of freedom which could be reconfigured prior to grasping without violating the task specifications. For example, a container filled with liquid may need to remain upright to prevent spilling, but it can still undergo planar displacements on the support surface. Thus it may be possible to adjust the object placement to a more desirable configuration which facilitates improved grasp quality. Observation of human manipulation suggests that directly grasping an object from its presented object configuration is not the only strategy for object acquisition. Instead, humans often take advantage of an object's movability in the environment to reorient and regrasp an object during the object-acquisition process. We hypothesize that object reconfiguration in the world prior to grasping should be considered as a manipulation strategy to improve dexterity. We refer to this strategy as *pre-grasp interaction*, because the reconfiguration prepares the object for the grasping action required for the subsequent primary manipulation task.

There are several reasons to examine this pre-grasp interaction strategy as an alternative to existing approaches for object acquisition. Taking advantage of object movability may increase performance by making good grasps possible for a greater number of initial object configurations. Pre-grasp interaction in the form of non-prehensile, or non-grasping, contact without lifting makes use of shared support

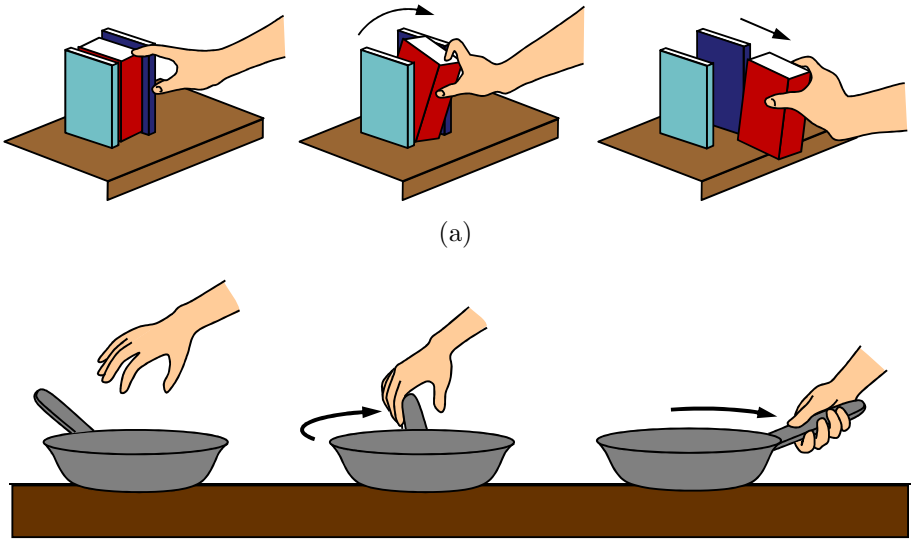

(a)

(b)

Figure 1.1: Pre-grasp interaction adjusts the object configuration in the workspace prior to grasping. (a) One example of pre-grasp manipulation is the pivoting of a shelved book in order to expose the book cover surfaces for grasping. (b) Pre-grasp object rotation of a cooking pan adjusts the object orientation such that the handle is reachable with the desired grasp and approach direction.

with the work surface, which can be more robust than direct grasps that must fully support the object load. The expense of tuning control parameters for complex manipulators can be reduced if first reconfiguring the object into a canonical pose enables a single well-tuned grasp to be applied to multiple initial poses. In addition, mimicking a pre-grasp manipulation strategy may lead to increased appearance of naturalness for an anthropomorphic manipulator interacting with a human. Finally, understanding the key challenges and limitations of the pre-grasp interaction strategy provides insight into the requirements for complex manipulation skills beyond pure grasping.

This thesis includes a survey of pre-grasp manipulation strategies used naturally by humans. The survey provides a framework for describing the variety of pre-grasp manipulation skills humans use in addition to direct grasping of objects. Similarities in the observed examples of pre-grasp manipulation provide insight into the common patterns of the object reconfiguration and the intent of pre-grasp manipulation.

A core part of the thesis examines pre-grasp object rotation as a specific example of pre-grasp manipulation. In pre-grasp object rotation, the object is pivoted in the plane of the support surface in order to reorient the object handle prior to grasping. Instead of completely re-planning a new direct grasp for novel object orientations, a pre-grasp rotation strategy decomposes the task into two components of object adjustment on the surface followed by execution of a canonical, well-tuned grasp.

One example of pre-grasp rotation is the reorientation of a cooking pan to achieve a more desirable handle direction for grasping (Fig. 1.1(b)).

We investigated whether humans consistently select a pre-grasp rotation strategy in response to changes in the presented object configuration. This pre-grasp interaction could be a mechanism for adapting novel task conditions to some intermediate state which affords improved grasp quality. We observed that the task factors of object weight and required angular precision change the human performance of pre-grasp rotation.

We also evaluated how pre-grasp rotation could improve robotic manipulation. A kinematic analysis of reachable object poses illustrates the extension of the workspace gained by object reorientation. The utility of pre-grasp rotation was demonstrated through the reuse of a manually-programmed grasping action as well as an increase in finding feasible grasping plans using our automated planning method. In some task scenarios, pre-grasp rotation also resulted in an increase in action quality in addition to action feasibility.

1.3 Contributions

This thesis makes the following contributions to the knowledge about human manipulation behavior and robot manipulation strategies:

- The identification of pre-grasp object interaction as a manipulation strategy, and a survey of human pre-grasp actions in natural activities and occupational tasks.
- An analysis of human performance of pre-grasp object rotation, including strategy selection criteria that determine task constraints.
- A method for automatically planning pre-grasp object rotation as part of an object transport task for a robotic manipulator.
- An evaluation of pre-grasp rotation as a strategy for increasing a robot manipulator's performance in terms of task feasibility and action quality.

The thesis also includes suggestions for future research on other pre-grasp interaction capabilities, besides pre-grasp rotation, for robot manipulators.

1.4 Organization

The remainder of this thesis is organized as follows. Chapter 2 provides background on related literature and describes where our approach fits within this context.

The first half of the thesis describes our observation of human pre-grasp manipulation. Chapter 3 presents a survey of human pre-grasp interaction examples in routine activities of daily living. Chapters 4 and 5 present two studies on how humans use pre-grasp object rotation for grasping objects by their handles.

The latter half of the thesis examines and evaluates pre-grasp rotation for a robot manipulator. Chapter 6 presents an analysis of the workspace limits of the alternative direct grasping strategy. Next, Chapter 7 presents an evaluation of pre-grasp rotation as a grasp reuse strategy . Chapter 8 presents a method for planning feasible pre-grasp rotation on a robotic manipulator completing a object transport task. Chapter 9 examines the quality of feasible transport plans and compares different optimization methods and cost metrics.

The thesis concludes with suggestions for future research in Chapter 10.

1.5 Publication note

Material in Chapter 3 previously appeared in [Chang and Pollard, 2009]. Parts of Chapter 4 appeared in [Chang *et al.*, 2008] and [Chang and Pollard, 2008]. Most of Chapter 5 was published as [Chang *et al.*, 2010a], co-authored with Roberta Klatzky. Chapter 7 was published as part of [Chang *et al.*, 2008], co-authored with Garth Zeglin. The material in Chapter 8 was published as [Chang *et al.*, 2010b], co-authored with Sidd Srinivasa.

2 Background

Our work examines how pre-grasp manipulation of a movable object could reconfigure the intermediate task conditions to improve the quality of the primary manipulation task. In particular, we focus on pre-grasp object rotation prior to grasping as an example of pre-grasp interaction, which uses the support surface to pivot an object to a more desirable handle orientation for grasping.

The ideas presented in this thesis build upon work from a number of related fields. This chapter discusses the thesis contributions in the context of previous studies of human manipulation (§2.1), human motor control (§2.2), robotic manipulation (§2.3), imitation learning (§2.4), and perceptual response (§2.5). Relevant work is also reviewed in more detail in later chapters.

2.1 Patterns of human manipulation

One contribution of this work is a framework for describing the types of pre-grasp interactions naturally used by humans in everyday tasks. The intent of the study is to provide insight and structure for understanding the complexity and variety of hand skills.

The dexterity of human manipulation has similarly motivated many researchers to understand the structure and motion of the human upper limb. The structural complexity of the human upper limb is both a source of its flexibility and an obstacle for understanding the biological mechanisms for control. Often the human arm is kinematically modeled with 7 degrees of freedom (DoFs) describing the shoulder (3-DoF), elbow (1-DoF), and wrist (3-DoF) joints. The redundancy allows multiple arm configuration solutions for a single configuration of the hand segment. In addition, the hand itself exhibits more than 20 DoFs for the palm, thumb, and finger joints. An overview of basic hand anatomy and function can be found in Napier [1993].

The high-dimensional kinematics allow for a wide variety of possible hand configurations. However, biomechanic and neuromuscular constraints suggest there are coordinated patterns underlying the apparent complexity. One approach to identifying such patterns has focused on categorizing hand poses as discrete grasp types. Napier [1993] proposed a basic taxonomy of power grasps and precision grasps, which are distinguished by contact with the finger or palm surfaces versus the finger tips.

More detailed taxonomies have been used to describe functional grasps for activities of daily living [Edwards *et al.*, 2002; Kamakura *et al.*, 1980] as well as for skilled machining tasks [Cutkosky, 1989]. In another approach, studies from the motor control community provide evidence that hand motion can be described as combinations of low-dimensional synergies or components. Work by Santello *et al.* [1998] and Mason *et al.* [2001] found that mimed reach-to-grasp hand shapes can be represented by just a few principal components in the joint angle space.

The survey of hand skills in §3 provides a framework for describing pre-grasp manipulation as a set of manipulation skills beyond direct grasping. Pre-grasp manipulation involves object reconfiguration in the environment and thus cannot be described only by the hand pose as classified in grasp taxonomies.

2.2 Human motor behavior and control

The first portion of the thesis (§4, 5) investigates the pre-grasp rotation strategy as a human behavior for manipulation tasks. The work contributes a study of pre-grasp manipulation to the field of human motor behavior and control.

Pre-grasp manipulation can be viewed as a coordinated pattern of movement for a certain class of manipulation tasks. Previous literature in the motor control community has already identified consistent patterns of human motion for the actions of reaching and grasping. Jeannerod [1981] has investigated the coordination between arm reaching motion and hand pre-shaping. The experiments analyzed timing correlation of the hand trajectory toward the object with the hand shape and also investigated the effects of sensory input during the reaching task. Lacquaniti and Soechting [1982] investigated the degree of coupling between the joints of the upper limb, and they found that shoulder and elbow but not wrist in particular were highly coupled during the tested reaching tasks. The research on grasp synergies by Santello *et al.* [1998] and Mason *et al.* [2001] demonstrated the low-dimensional variation in hand shape in response to different object geometries. Furthermore, Lukos *et al.* [2007, 2008] showed that hand grasp shape, as measured by fingertip contact points on the object, responded not just to object geometry but were also modulated in anticipation of asymmetries in the location of the object's center of mass.

Coordination of the hand for grasping also involves the regulation of force and compliance in addition to kinematic configuration. Johansson [1996] investigated how finger grip force magnitude responds to object surface friction. The experiments of Li *et al.* [1998] and Latash *et al.* [1998] suggest the existence of motor synergies in the force output of individual fingers, where force sharing patterns between digits in a multi-finger grasp address the motor redundancy of an overactuated system. The patterns of anticipatory contact point modulation in the hand grasp shape [Lukos *et al.*, 2007, 2008] may be chosen to better reject force or torque disturbances during the grasp [Ciocarlie and Allen, 2009]. Recent work by Balasubramanian and Mat-

suoka [2008] has also investigated force coordination in terms of muscle activation patterns for the index finger system.

A consistent pattern of coordination represents a particular subset of the full set of movements possible by a redundant and over-actuated system. One approach to understanding the selection of particular strategies is to describe manipulation actions by the optimization of some cost criteria. Previous studies of possible optimization criteria of arm control have investigated minimum jerk [Flash and Hogan, 1985] and minimum torque change [Uno *et al.*, 1989] for arm motion trajectories. In addition, work in the biomechanics community has examined how static postures for lifting tasks may be predicted by energy or effort costs [Dysart and Woldstad, 1996; Chang *et al.*, 2001]. For certain tasks such as one-handed lifting motions [Park *et al.*, 2005], humans may select between multiple movement strategies according to individual preference or capabilities. In two-handed lifting tasks, the selection between either a stooping or squatting whole-body lifting posture may depend on factors such as individual height or strength [Burgess-Limerick and Abernethy, 1997].

The work discussed in §4 and §5 contributes to the literature by demonstrating the consistent pattern of pre-grasp rotation in humans and examining factors which affect the strategy selection. In particular, our work extends the understanding of human behavior to task domains where more complicated actions are performed. For tasks where the object is easy to grasp, a direct reach-to-grasp action may be sufficient. However, more difficult tasks may require or encourage the use of pre-grasp interaction as a more successful manipulation strategy. We show that humans use pre-grasp object rotation for demanding tasks involving heavy objects and strict angular precision requirements when the object is presented with the handle in a non-canonical orientation relative to the person's body.

Pre-grasp manipulation involves changing the object configuration in the world prior to grasping. To understand how an object can be reconfigured in the world, one must have or acquire a model of the object affordances. Lederman and Klatzky [1993] describes manual exploratory procedures used by humans to identify object properties. Stoytchev [2005] provides an overview of literature on object affordances and tool usage. This thesis does not address the acquisition of object models and assumes a priori knowledge of an object's degrees of freedom.

2.3 Robotic manipulation

The latter part of the thesis (§6, §7, §8, and §9) focuses on pre-grasp rotation on a robot manipulator. A major component is the implementation and evaluation of a planning framework which incorporates pre-grasp rotation with reaching and grasping actions. The contribution of this portion of the work is to provide a method for synthesizing plausible actions and to demonstrate the functional advantages of pre-grasp manipulation compared to direct grasping approaches.

Both pre-grasp manipulation and direct grasping involve grasping and reaching components. Previous robotics literature has investigated in depth the problems of grasp planning and motion planning. This thesis does not address grasp planning, and our implementation of pre-grasp manipulation uses predefined grasps. The predefined grasps can be manually selected or automatically determined by the methods discussed in §2.3.1. The implemented framework in §8 and §9 primarily builds upon the existing sampling-based methods for motion planning of a robotic manipulator, which are discussed in §2.3.2.

In addition, motion planning also includes actions which could be used for the pre-grasp manipulation component, such as planar pushing or tumbling. Section 2.3.2 also reviews relevant work in this area. In many cases, these actions have been analyzed as the primary manipulation task rather than as a pre-grasp interaction. For example, for a primary pushing task, the end state of a pushing action is already specified as the task goal. Our investigation in §8 examines how the choice of the intermediate object configuration can be selected to specify the goal of pre-grasp action and the target for the reaching action.

The presented implementation of pre-grasp rotation in §7–§9 focuses on the example of pre-grasp rotation where the pre-grasp action and the subsequent primary grasping action can be considered as two distinct segments of the motion. However, in-hand manipulation is an example of pre-grasp manipulation where the motion may occur as a single continuous action leading to the final whole-hand grasp for the specified task. Previous research on in-hand manipulation of objects is discussed in §2.3.3.

2.3.1 Grasp planning and analysis

Several efforts in robotics have studied grasping from the approach of finding appropriate contact points for finger placements. Analysis of a candidate grasp, as defined by the set of contacts placed on an object, involves evaluation of properties such as form closure and force closure in order to determine whether a grasp is able to sufficiently constrain the object configuration. Mason [2001] provides an overview of grasp analysis.

One research focus is the development of algorithms for synthesizing the contact placements for grasping a given object [Mishra *et al.*, 1987; Li and Sastry, 1988; Ponce *et al.*, 1997; Pollard, 2004]. Quality metrics used for evaluation and optimization of grasp force and stiffness properties have been explored by several studies, such as Zhu *et al.* [2001]; Lin *et al.* [2000]. Often the metric describes the interaction of the object and the contact points of the manipulator end effector. Grasp quality measures by Li *et al.* [1989]; Fu and Pollard [2006] consider both the task constraints as well as the mechanical constraints of the manipulator.

Another approach views grasps as prototypical hand shapes, as suggested by the taxonomies of human grasps such as that proposed by Cutkosky [1989] (§2.1). In

imitation learning applications, the grasp type for a particular object is determined by the classification of the hand shape demonstrated by a human, as used in the approaches by Ikeuchi and Suehiro [1994]; Kang and Ikeuchi [1997]; Ekvall and Kragic [2005]; Moussa and Kamel [1995]; Bernardin *et al.* [2005]. The approach by Li and Pollard [2005] frames grasp synthesis as a shape-matching problem, where candidate hand postures for grasping a new object model is chosen from a database of several examples of static grasp poses.

2.3.2 Motion planning

A collision-free arm configuration and trajectory is also required for manipulation in addition to the hand grasp. Several works have investigated the use of sampling-based planning techniques, such as Rapidly-exploring Random Trees (RRTs) [LaValle and Kuffner, 2000], for path-planning of high-dimensional manipulator arms. Bertram *et al.* [2006] and Weghe *et al.* [2007] present modifications to the standard RRT algorithm in order to avoid using inverse kinematics for solving the goal configuration of the arm. Berenson *et al.* [2007] and Diankov *et al.* [2008a] precompute possible grasps for an object and then plan the manipulator path to reach a feasible grasp. Similarly, Gienger *et al.* [2008] first build a task map of the grasp manifold for an object and then plan arm motion toward the manifold represented as attractor points. Although some of these approaches consider how the object grasp and arm motion can be planned together, the object placement in the environment is considered stationary.

Another body of work in the motion planning literature focuses on planning manipulation to reconfigure objects in the workspace. Methods for planning non-prehensile pushing of objects have been investigated by Lozano-Perez *et al.* [1984]; Mason [1986]; Lynch and Mason [1995]. Push-planning methods have also been demonstrated on other humanoid platforms in recent work [Stilman and Kuffner, 2005; Hauser *et al.*, 2007; Ng-Thow-Hing *et al.*, 2007]. Other examples of manipulation maneuvers which could be used for pre-grasp object reconfiguration are toppling and tumbling [Lynch, 1999]. In these scenarios, the reconfiguration of the object on the support surface is specified as the primary task goal, rather than as a subtask preceding a grasping task. Previous literature which has examined the pushing task as a pre-grasp action for a prehensile grasp includes the work by Simeon *et al.* [2002] and Sahbani *et al.* [2002], where the manipulation consists of pushing and regrasping a movable object in a sequence of alternating transit and transfer steps. Recent work by Berenson *et al.* [2009b] presents a method for planning dragging motions of grasped objects prior to lifting in order to meet the task constraints. In these examples, the object is already grasped at the start of the dragging motion. This thesis builds upon the previous approaches by focusing on pre-grasp interactions which occur prior to the object acquisition.

2.3.3 In-hand manipulation

Besides transporting and reconfiguring objects in the environment, pre-grasp manipulation also includes reorienting objects within the hand. Work by Cole *et al.* [1989, 1992] has developed controllers for using rolling and sliding finger contacts to achieve a desired object trajectory. Brock [1988] presents an analysis of in-hand pivoting which uses gravity to reorient an object by controlled slipping.

Pre-grasp manipulation can be especially important for the task of lifting objects from a work surface into the hand for tool use where the placement of the object on the surface prevents a direct approach of the final desired tool grasp. Kleinmann *et al.* [1996] define six strategies for the transition between a precision fingertip grasp to an enveloping power grasp. The six strategies proposed are rolling, snapping, stepping, gravity, assistance, and replacing, and some of the strategies include the use of an additional object or surface in the environment to manipulate the grasped object. In a similar approach, Kaneko *et al.* [2000] develops a flow chart structure dictating the procedure of how a robotic manipulator would explore five strategies for lifting cylinders depending on the object size and surface properties. Large objects may be grasped directly, but smaller objects require pre-grasp manipulation such as sliding, rolling, regrasping, or regrasping with rolling contact to detach the object from the surface.

2.4 Example-based learning

This thesis investigates a strategy observed in human manipulation and the potential benefits of such a strategy for a robot manipulator. Our approach is to transfer the human manipulation skill by a high level, manual extraction of the strategy. As such, it is not an example of imitation learning of the manipulation actions. Imitation learning may be appropriate for learning parameters of a given pre-grasp action, such as the amount of rotation or good contact placements. Such a framework for imitation learning would necessarily include the assumption of the pre-grasp strategy for object adjustment, rather than learning the strategy directly from observation. Following this approach, we also present a taxonomy of pre-grasp manipulation strategies which is the result of manual classification of observed hand activity.

The design and control of robotic manipulators is often inspired by the amazing dexterity of biological manipulation systems. Example-based learning and imitation learning are approaches for leveraging human demonstration to improve robot performance on the same task. Imitation learning for grasping has been explored for recognizing the demonstrated grasp type in order to execute corresponding pre-programmed motor primitives on a robotic manipulator [Kang, 1994; Ikeuchi and Suehiro, 1994; Kang and Ikeuchi, 1997]. More recent work has investigated methods for learning from kinesthetic demonstrations to reproduce arm trajectory motions for a particular task [Billard and Siegwart, 2004; Calinon and Billard, 2008]. Ratliff

et al. [2007] present a method for learning an objective cost function from examples when there are many features which could describe a complex task.

Several works also present methods for choosing grasp contact points on an object based on similar grasp examples. Pollard [2004] uses a single grasp example to efficiently constructs families of grasps for tasks with a large number of contacts. Saxena *et al.* [2008] learn contact grasping points from synthetic 2-D images without constructing 3-D object models. Hsiao and Lozano-Perez [2006] present a method for modifying whole-body grasps to new objects by transforming sub-parts of the object shape.

The use of human examples for generating natural-looking motion has also been explored in computer animation of virtual characters [Yamane *et al.*, 2004; Safonova *et al.*, 2004, e.g.]. Adapting an example motion to a novel situation is particularly challenging for actions which involve making and breaking contact with the environment. Relevant work in the synthesis of hand motion for animated virtual environments has investigated grasping motions [Pollard and Zordan, 2005] and finger-gaiting manipulations Kry [2005]. Pollard and Zordan [2005] generate physically-based hand motion using a controller that models both active and passive components properties of a single grasping example. The interaction capture approach presented by Kry [2005] measures the motion and forces of manipulation tasks, from which joint compliances and target trajectories are estimated for application to new object geometries.

2.5 Perceptual response to manipulation actions

Although a major challenge for robotics research is to develop methods for automating functional manipulation tasks, another open question is how to design motion which appears natural or pleasing to a human observer. For robotic manipulators to be adopted as assistive devices in human environments such as the home or workplace, it may be preferable for the manipulator actions to avoid awkward or unpredictable motion even if the manipulation task is successfully completed.

Because pre-grasp rotation is observed to be as a consistent strategy used by humans (§4, §5), adopting such a strategy in a robot may improve the perception of naturalness. In an informal study, not presented in this thesis, we have begun a preliminary investigation of whether the difference between a grasping strategy using pre-grasp rotation and a direct grasping strategy influences the impression of a manipulator's apparent dexterity.

In related work, Huber *et al.* [2008] and Shibata and Inooka [1998] have investigated how the timing of a robot's reaching motions affect a human's perception of its naturalness. In the neuroscience community, Stanley *et al.* [2007] and Kilner *et al.* [2007] have also investigated movement timing and the degree of anthropomorphism in the agent as factors which influence a human's interference response to movement

stimuli. These works focused on the motion of the agent. Future work beyond this thesis is needed to determine how factors such as timing or robot anthropomorphism interact with the pre-grasp interaction strategy to change a human's perceptual response to the manipulation action.

3 Pre-grasp interaction in natural human manipulation activities

Humans typically use a few prototypical reaching and grasping actions to pick up objects during manipulation tasks. In daily life, humans must grasp objects from a variety of initial configurations, including many of which may not be well-matched to canonical grasps and the arm approach directions. In our observation of human grasping, we notice that humans do not always grasp an object directly from its presented configuration. Instead, humans often manipulate the object to adjust its configuration prior to grasping. For example, a person might slide a mug on a table closer to the body by pulling on the handle with unidirectional or non-grasping contact. When grasping a pen from a table surface, the fingers may quickly pivot the pen to orient the tip for the subsequent writing task.

These are examples of what we refer to as *pre-grasp interaction*. Pre-grasp interaction occurs whenever the manipulation first adjusts the object configuration prior to the final grasp. In several cases, the adjustment may occur while the object is partially supported by the environment, such as a table surface. This approach takes advantage of the object's movability on the supporting structure to effectively change the intermediate task parameters. Thus the anticipation of the grasping task includes object reconfiguration as a motion in addition to the arm reaching and hand pre-shaping movements.

This chapter includes a description of typical types of pre-grasp interaction from our survey of human hand activity in natural settings. In the observational survey, we recorded video of people performing manipulation tasks in natural settings such as the home or place of occupation. The activities included tasks such as sorting office supplies, washing dishes, and moving furniture. We organized the survey examples into a framework that describes a pre-grasp interaction by the object reconfiguration and the task constraint that is modified. This framework provides the broad context of pre-grasp interaction strategies in which the specific example of pre-grasp object rotation is situated.

The survey results presented in this chapter have previously appeared in workshop proceedings [Chang and Pollard, 2009].

3.1 Pre-grasp interaction characteristics

Pre-grasp interaction is any manipulation that reconfigures the object for a subsequent interaction when the task which does not specify the reconfiguration (Fig. 3.1). This definition allows for a broad concept of pre-grasp interaction as any "preparatory action" on the object prior to the final task phase. In this sense, many multi-step operations or complex manipulation processes inherently include pre-grasp interaction between the initial presented task conditions and the final primary interaction required for task completion. All actions that do not include pre-grasp interaction are referred to as "direct" manipulation strategies. Direct interactions consist of only the primary interaction that achieves the task goal.

We characterize a pre-grasp interaction by two aspects: (1) the object reconfiguration and (2) the pre-grasp intent relative to the primary task goal. Without either component, the interaction is instead a direct manipulation where there is only manipulator reconfiguration prior to the primary interaction. Note that the primary interaction may include object reconfiguration (such as in lifting or pushing), which is distinct from pre-grasp reconfiguration if the task explicitly specifies the motion.

The pre-grasp object reconfiguration can be described according to the degrees of freedom which are reconfigured by the pre-grasp interaction. Examples of object reconfiguration classes include the following:

- Planar displacement – Rigid planar displacement consisting of 3-DoFs of 1-DoF in-plane rotation with 2-DoF in-plane translation is commonly observed for interactions where the object is supported on a flat surface such as a table. Pre-grasp planar displacement can also occur in the manipulation of stacked objects, such as removing the top stacked book by sliding (Fig. 3.2(a)).
- Rigid displacement – General 6-DoF rigid displacements include pivoting or tumbling out of the horizontal support plane. For example, a piece of furniture may be tumbled in order to reach a particular handhold before transporting (Fig. 3.2(b)).
- Morphological reconfiguration – Pre-grasp interaction can also change the shape of an articulated or deformable object. The change in shape may create a better hand-hold for a desired grasp, as in the case of reorienting the handle of a bucket before lifting (Fig. 3.2(c)). Morphological reconfiguration may also include the manipulation of a set of objects which are separate but conceptually linked to each other, as in a pile of homogeneous objects such as game tokens or an assembly of multiple objects such as a container with a matching lid.

It is assumed that the reconfiguration does not violate any specified task constraints. In fact, the motion should improve a presented object configuration that is suboptimal. This assumed improvement in the object configuration is related to the intent of the pre-grasp manipulation.

The intent of a pre-grasp interaction is described as the benefit of the object re-

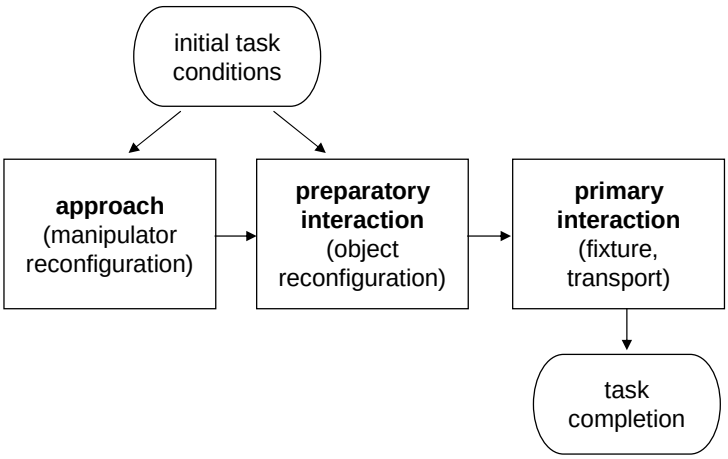

Figure 3.1: Pre-grasp manipulation refers to any preparatory interaction which reconfigures the object prior to the primary interaction specified by the task goal.

configuration relative to the primary interaction. The presented object configuration in the initial task conditions may be suboptimal with respect to the environment, manipulator, or object. These constraints are described in terms of how they limit the final grasp if only direct grasping without pre-grasp interaction is used.

- Environment constraints – In the initial configuration, the desired contact surface on the object could be unexposed or blocked by an obstacle in the environment. The obstacle in the environment could be the supporting structure if the desired grasp requires contact with the object's bottom surface. This case occurs when a sheet of paper resting on a surface is to be grasped on both faces (Fig. 3.3(a)). Clutter in the environment may also block the approach to a desired contact area.
- Manipulator constraints – Given the initial configuration of the manipulator relative to the object, a desired grasp of the object could be limited by a suboptimal manipulator posture using direct reaching. For example, a direct grasp may require an arm posture associated with low kinematic manipulability. A person may prefer to reorient an object using bimanual regrasping while maintaining manipulability in both arms rather than reaching across the body with one hand if it requires a near-singular arm posture (Fig. 3.3(b)). Other factors which might affect the preference for a particular manipulator posture include torque capability constraints and grasp stability.
- Object constraints – In some situations, the presented object morphology may not afford the desired grasp contacts or may have less desirable physical properties. The reshaping of an articulated object for grasping, such as the previous example of the bucket handle (Fig. 3.2(c)), is similar in intent to overcoming environmental constraints because the ungrasped object subparts are obstacles to the grasped subpart. In the bucket example, the reconfiguration also

(a)

(b)

(c)

Figure 3.2: Examples of object reconfiguration classes for pre-grasp manipulation. (a) Planar displacement: Books are slid off the top of a stack to grasp the bottom surface. (b) Rigid discplacement: A cabinet is tumbled to reach a handhold on the bottom in order to carry it sideways. (c) Articulated motion: The bucket handle (highlighted in white) is lifted from the side to achieve a whole-hand grasp of the handle.

changes the grasp location relative to the object's center of mass. Pre-grasp manipulation may also gather deformable materials or a set of multiple objects into a compact shape more suitable for grasping. An example of this case is the scooping of food scraps into a cupped hand (Fig. 3.3(c)).

The different constraints addressed by the pre-grasp intent are not necessarily mutually exclusive. In the example of bimanual regrasp of a pan by its handle (Fig. 3.3(b)), the intent of the pre-grasp action could be a combination of avoiding environment clutter and using an arm posture with improved manipulability.

In addition to the variety in reconfiguration and intent, there is also a range of scales of pre-grasp interaction with respect to the manipulator. While this range is

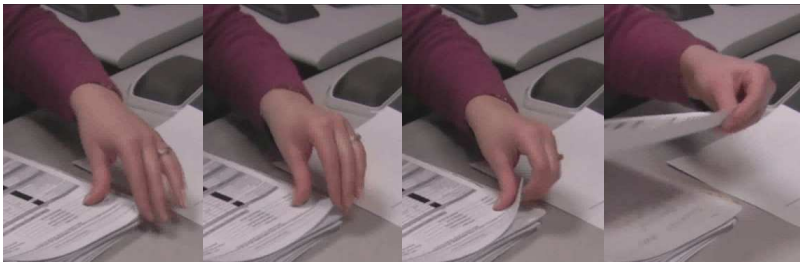

(a)

(b)

(c)

Figure 3.3: (a) Environment constraints: The top sheet of paper is lifted from the stack to expose the underside for grasp contact. (b) Manipulator constraints: The pan is tumbled by the right hand so that the left hand can grasp the handle. (c) Object constraints: The pile of peelings is reshaped between two hands before lifting from the cutting board.

not a fundamental aspect defining the pre-grasp interaction, it highlights the richness of scenarios that are relevant to this strategy. Examples of the different interaction scales are related to the final grasp contacts:

- unimanual grasp of a writing utensil,
- bimanual grasp of a large pot,
- whole-body grasp of large box making contact with the upper arm and torso, and
- cooperative manipulation between multiple agents for lifting a piece of furniture.

We hypothesize that the scale of manipulator interaction is partially determined by the object size and mass, because larger and heavier objects are likely to require wider grasp apertures and contact forces for manipulation. For example, objects

manipulated purely by unimanual contact are expected to be smaller or lighter than objects handled in cooperative manipulation between two people. Another factor in determining the scale of interaction is contact with the environment, which shares support of the object load. The environmental structure supporting the object effectively becomes an extension of the manipulator. This scenario often occurs when there is non-grasping, or non-prehensile, contact with the object, as in the case of pushing of an object which is supported against gravity by a table.

3.2 Video survey of natural pre-grasp interactions

We surveyed examples of natural human pre-grasp interaction strategies in order to provide context for the specific example of pre-grasp rotation. Our goal was to develop a taxonomy for classifying the variety of pre-grasp action primitives which are integrated into complex reach-to-grasp tasks. We were specifically interested in surveying human hand activity in natural settings in contrast to instructed tasks within a laboratory environment. In this way, we could capture the richness of pre-grasp interactions beyond the direct reach-to-grasp actions studied previously in the literature.

In the video survey of human hand activity, we filmed people performing manipulation tasks in natural settings such as the home or place of occupation. All participants provided informed consent. In all observations, the participants performed manipulation skills which had been practiced previously as part of their regular occupation. There were a total of 10 sessions of both individual and group manipulation activities, such that overall 38 people were filmed (Table 3.1). The sessions covered activities for housekeeping, food preparation, office work, and mechanical repair. Specific tasks include sorting office supplies, washing dishes, and moving furniture.

We found that there is indeed a broad class of pre-grasp interactions where the object is not grasped directly from its presented placement in the environment. Our framework describes the survey examples according to two main aspects of the pre-grasp interaction described earlier in §3.1. The first aspect is the type of object re-configuration resulting from the interaction. The second aspect is the underlying intent of the interaction to improve the posture quality or grasp quality of the manipulation action.

First, the object reconfiguration is classified by a taxonomy based on the degrees of freedom which were adjusted by the pre-grasp interaction (Fig. 3.4). The object motion may be completely comprised by planar displacement. This case is common in examples of non-prehensile pre-grasp interaction where the object is primarily supported on a horizontal surface. Alternatively, for a bulky piece of furniture, the pre-grasp tumbling interaction may result in general 6-degree-of-freedom rigid displacement. In more complex cases, the pre-grasp interaction may cause a mor-

Table 3.1: Daily and occupational tasks observed in video survey of natural manipulation activities.

Category	Session (number of participants, total footage recorded)
Food preparation	Potluck lunch set up (5 people, 24 min.)
	Cooking dinner (1 person, 21 min.)
	Cooking at a restaurant (4 people, 30 min.)
Housekeeping / cleaning	Party set up (5 people, 11 min.)
	Lab clean up (14 people, 13 min.)
	Fish tank maintenance (1 person, 5 min.)
Office work	Desk work (1 person, 5 min.)
	Library work (1 person, 10 min.)
Other	Bicycle repair (3 people, 34 min.)
	Moving furniture (3 people, 12 min.)

phological reconfiguration of a deformable or articulated object, such as a bucket with a hinged handle.

Second, the pre-grasp interaction is described by the intent of the object adjustment in relation to the final grasp. The presented configuration of the object in the environment could be suboptimal for direct reach-to-grasp object acquisition due to preferences for a particular body posture and/or grasp. When the handle on a cooking pan is oriented away from the person, a direct grasp of the handle may be feasible but could require lifting the heavy pan from an uncomfortable body posture with limited lifting capability. In other scenarios, the intent of the pre-grasp interaction may be to improve the grasp quality rather the posture quality. This grasp improvement is especially relevant to situations where environmental clutter occludes the desired grasp contact surfaces, as in the case of a shelved book where only the spine is exposed in the initial task condition. The observed examples suggest that the intent of pre-grasp interaction was often a combination of preferences for both posture quality and grasp quality, and potentially other optimization metrics that influence the constraints described in §3.1.

The presented pre-grasp interaction framework suggests several approaches for improving the dexterity of robotic manipulators. Taking advantage of object movability may extend the effective workspace by changing the environmental constraints when direct reach-to-grasp actions are of insufficient posture and/or grasp quality. Non-prehensile pre-grasp interaction could reduce the load on the manipulator by using shared support with the work surface during the initial interaction with the object. Moreover, the expense of tuning control parameters for complex manipula-

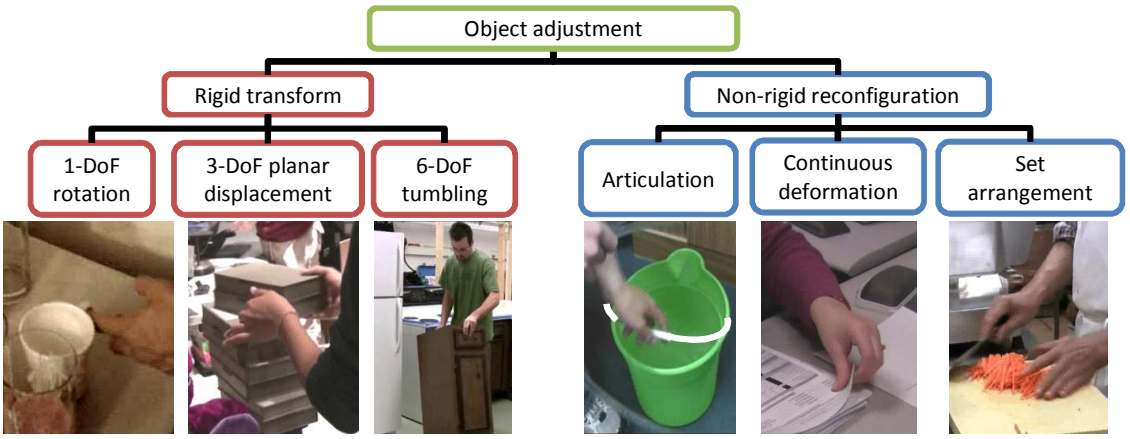

Objects in observed examples

mug	books	cabinet	bucket	paper	peelings
stapler	CD cases	pan	ladder	rug	bolts
pan	pitcher	wrench		foil	paper stack
	large box	paperclip		plastic bag	set of jars
	plates	spatula			chopped
	food carton	chopsticks			produce
		books			

Figure 3.4: Taxonomy for the object reconfiguration aspect of pre-grasp interaction. Examples of rigid planar transformations included rotation of a cup by its handle and sliding books off the top of a stack. General rigid tumbling was used to achieve a whole body grasp of a bulky piece of furniture. Pre-grasp interaction was also observed for non-rigid objects. A hinged bucket handle was rotated to achieve a hook grasp, and a piece of paper was curled to achieve a pinch grasp. Multiple objects were also rearranged as a set, such as in the scooping interaction with a pile of carrot peelings.

tors can be reduced if pre-grasp object reconfiguration enables the reuse of a single well-tuned grasp action for multiple initial placements. Finally, because pre-grasp strategies are part of natural human manipulation, incorporating them in the repertoire of assistive or teleoperated manipulators could facilitate more intuitive control for human operators.

3.3 Pre-grasp rotation as an example interaction

The remainder of this thesis examines pre-grasp object rotation as a specific example of a pre-grasp interaction strategy. The object reconfiguration for pre-grasp rotation consists of 1-DoF angular rotation in the plane of the support surface. One example of pre-grasp rotation is the reorientation of a cooking pan on a horizontal surface

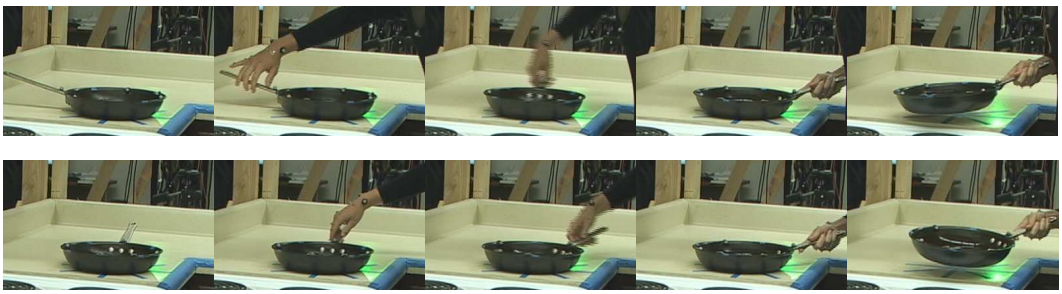

Figure 3.5: Pre-grasp manipulation changes the object configuration in the workspace prior to grasping. One example of pre-grasp manipulation is pre-grasp object rotation, shown here for a cooking pan with a handle, where the object orientation is adjusted before grasping.

to achieve a more desirable handle direction for grasping (Fig. 3.5). Non-prehensile manipulation can be used to take advantage of the supporting surface in order to pivot the object without grasping.

In particular, this thesis studies examples of pre-grasp rotation achieved by unimanual manipulation. By limiting the manipulation to unimanual contact, we examined the use of pre-grasp manipulation in more demanding tasks with constrained resources. Pre-grasp rotation can be also performed using bimanual manipulation, whole body grasps, or cooperative manipulation between multiple agents. In these cases, the task may be more difficult and involve larger or heavier objects that require additional grasping contacts.

In addition, this work focused on objects with handles in order to address the use of pre-grasp rotation independently from the selection of grasp contact points. We assumed that the handle represents the desired contact set for the specified task.

In the following chapters, pre-grasp rotation is examined in detail from a variety of aspects. The earlier chapters investigate the human performance of pre-grasp rotation for object lifting tasks. The later chapters present a method for and evaluate the functionality of pre-grasp rotation for robotic manipulators.

4 Pre-grasp object rotation in humans

The previous chapter presented examples of human manipulation activities in natural settings. In particular we reviewed examples of different types of pre-grasp manipulation observed in natural hand activities. A representative class of pre-grasp motion is planar displacement on a horizontal support surface to reconfigure the object prior to grasping. An even more constrained case within this class is pure rotation of the object in the plane.

This chapter presents a study of pre-grasp object rotation as strategy used by people for manipulating familiar household objects. The first goal of the study is to quantify the consistency of object rotation as a pre-grasp strategy. When there is pre-grasp manipulation, the action in anticipation of a grasping task not only consists of changes in the manipulator's posture such as arm reaching movement and hand pre-shaping, but it also includes changes in the object configuration in the environment prior to grasping. Our experimental results found a consistent trend that the amount of object rotation changes quadratically with respect to the initial orientation angle.

The second goal of the study is to examine possible criteria which make the pre-grasp rotation strategy advantageous compared to the alternative of a direct grasping strategy. In the case of lifting heavy objects, rotating the object prior to lifting may move the handle to a familiar canonical configuration. The canonical orientation may be preferable to the orientation initially presented because it requires less exertion in the lifting posture, allows more stable grasps of the handle, or is perceived to be more comfortable. Our results show that for increased rotation amount there is a corresponding decrease in upper body joint torques for the lifting posture. In addition, initial handle orientations further from the canonical orientation resulted in greater differences in the selected handle grasps.

In the following chapter, we present another study which examines how the specific task factors of object weight and required angular precision affect how humans use pre-grasp object rotation for grasping.

Work in this chapter has been previously published as part of Chang *et al.* [2008]. Additional details on the experimental procedure and analysis can also be found in Chang and Pollard [2008].

4.1 Related literature

Humans manipulate objects daily in response to a wide variety of task conditions such as object shape, weight, and placement in the environment. Some underlying characteristics of human reaching and grasping have been found to be task invariant and subject invariant, such as the timing coordination of hand transport with hand pre-shaping during reaching [Jeannerod, 1981] or motor synergies in hand force regulation [Li *et al.*, 1998; Latash *et al.*, 1998]. However, other aspects of human manipulation may change with the task condition, such as finger grip force magnitude in response to object surface friction [Johansson, 1996].

In particular, postural components of a manipulation action can vary according to the specific task parameters. Bongers *et al.* [2004] reported that the stopping distance and tool-use posture in a reaching manipulation task was dependent on the tool length and object size. Work by Stelmach *et al.* [1994] and Desmurget *et al.* [1996] showed that wrist angle and wrist trajectory changed according to the orientation of the cylindrical object placed in the same position. Similarly, Rand and Stelmach [2005] studied how grasp aperture and arm transport in reach-to-grasp movements for actions with or without forearm pronation depending on the starting object orientation.

For some tasks, there may be discrete choices for the postural response. Rosenbaum and colleagues have investigated the selection between overhand and underhand grips for a variety of handle transport and handle rotation tasks where a cylinder is grasped and placed at several different goal configurations [Rosenbaum *et al.*, 1990, 1992, 1993b]. The findings suggest that the choice of initial hand grasps is due the perceived end-state comfort of a task [Rosenbaum *et al.*, 1992; Short and Cauraugh, 1997; Zhang and Rosenbaum, 2008]. Similarly, in bimanual sagittal lifting tasks, there are discrete alternative movement techniques of stooping and squatting [Park *et al.*, 2005; Burgess-Limerick and Abernethy, 1997]. The selection between the two postural techniques differs between individuals and may depend on factors such as individual height or strength [Burgess-Limerick and Abernethy, 1997].

Previous work has largely focused on postural responses to task conditions, and usually the manipulated object is grasped directly from its presented configuration in the workspace. Less well-studied are movement strategies where the object configuration may be manipulated prior to grasping. This pre-grasp manipulation strategy takes advantage of the object's movability on the support surface to effectively change the intermediate task parameters, such as making a handle on an object more reachable. In such cases, the action in anticipation of a grasping task not only consists of postural changes such as arm reaching movement and hand pre-shaping, but it also includes changes in the object configuration in the environment prior to grasping. In recent work, Rosenbaum and Gaydos [2008] presented a method for evaluating the relative costs of alternative choices for object-positioning and object-rotation tasks, but the object reconfiguration was the task itself rather than a pre-grasp motion for

a subsequent action such as lifting.

There have also been previous studies examining possible optimization criteria for motor control. Exertion or energy as measured by the magnitude of joint torques has been used for simulating walking behavior [Chow and Jacobson, 1971] and predicting postures for sagittal plane lifting [Dysart and Woldstad, 1996; Chang *et al.*, 2001]. Several studies [Fischman, 1998; Short and Cauraugh, 1997; Rosenbaum *et al.*, 1992] have investigated end-state comfort for hand grasps of elongated objects. It is possible the perceived grasp comfort may be related to the grasp location on the object, as work by Turvey *et al.* [1999] found that object inertia asymmetry was a factor in weight perception.

4.2 Experiment overview

This study specifically examines object rotation as a pre-grasp manipulation strategy for lifting heavy objects from different presented object orientations. The experiments recorded how people perform an object transport task in response to changes in the initial object orientation. Participants approach an object on a tabletop and transport the object laterally to the goal area on the tabletop. We aim to capture the participants' natural grasping strategy behavior by allowing the participants to move at a self-selected speed, select the standing position for the transport task, and choose the grasp points on the object. These movement choices are recorded by tracking participant's body posture with marker-based camera system. In addition, we track the object to record how and whether the participant takes advantage of the object's movability on the surface to effectively change the intermediate task parameters for grasping.

In the analysis, we investigate how the amount of pre-grasp object rotation changed in response to eight different object orientations on a horizontal surface. Other metrics computed from the body posture at time of object lift-off included the joint torques at the torso, shoulder, elbow, and wrist, and the grasp configuration with respect to the object handle. All of the response variables were compared between different task constraints where pre-grasp object rotation was permitted or not permitted in the verbal task instructions.

4.3 Transport task for two objects with handles

4.3.1 Participants

Ten right-handed adults (5 male, 5 female) volunteered for the study (age = 26.7 $\pm$ 3.5 years [mean $\pm$ standard deviation], height = 167.1 $\pm$ 9.1 cm, weight = 58.7 $\pm$ 10.9 kg). All participants signed informed consent forms approved by the Institutional Review Board.

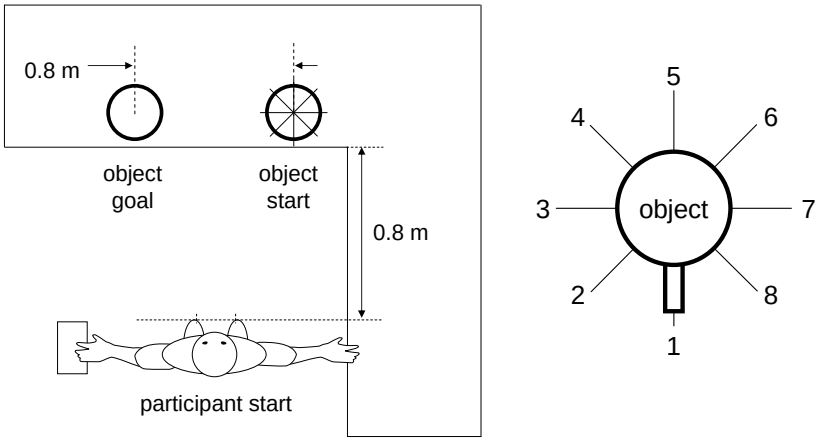

(a) (b)

Figure 4.1: General layout of the experimental setting. (a) Participants started in a standing position facing the countertop setting. Participants transported the objects from the start position to the goal position with their right hand. (b) In each trial, the object started in one of eight orientations defined by the handle direction. In the figure, the object is in orientation 1, the baseline orientation.

4.3.2 Transport task conditions

Participants performed the object lifting tasks in a kitchen counter top setting (Fig. 4.1). The object start position was located on the right side counter area. The object goal area was located 77 cm to the left of the start position and was marked by a circular cover over the bottom left stove burner. At the start position, the object was presented in one of eight possible orientations, indicated by the direction of the object handle. In the baseline orientation, the handle directly faces the participant, and the remaining 7 orientations spanned a full 360 degrees in 45 degree increments. The two handled objects tested for all ten participants were a large plastic water jug and a cast iron frying pan, both without lids (Fig. 4.2). The objects were filled with water for a total mass of 3.4 kg for the jug and 1.5 kg for the pan.

4.3.3 Motion capture of task performance

Kinematic data for the participant and objects were recorded at 120 Hz using a Vicon camera system (Vicon Motion Systems, Los Angeles, California, USA). Motion of the full body, including hands and fingers, was tracked by 80 reflective markers attached to the participant. The marker set included 4 markers on the hips, 8 markers on the torso, 7 markers on the arm, 6 markers on the hand dorsum, and 10 markers on the fingers. Prior to the lifting tasks, a calibration trial was recorded where the participant exercised the range of motion for all joints. In the following experiments, each object was tracked by 5 attached markers.

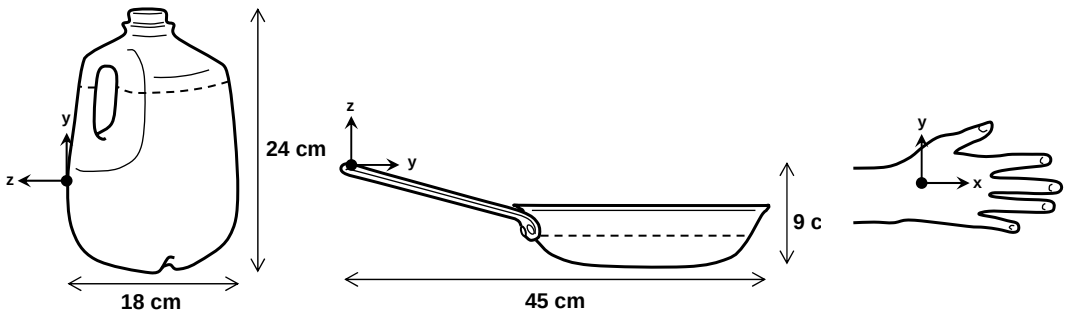

(a) (b) (c)

Figure 4.2: Participants transported two different handled objects filled with water, (a) an uncapped water jug and (b) a cast iron cooking pan. The dashed line indicates the approximate water level. The coordinate systems located at the object handles provide a reference frame for measuring the (c) hand dorsum configuration during the object grasp.

In each lifting task trial, the participant started facing the counter with toes at a distance of 78 cm from the counter edge, such that the object was outside of arm's reach. The participants arms were held out to the side with the hands resting prone at table height, which improved the tracking of the hand markers by the camera system. For all trials, the task goal was to move the object to the goal position without spilling any water. No specific object orientation was required at the goal. Participants were instructed to perform the transport task at a self-selected speed with no time constraints. We investigated the performance of the transport task for the two objects over the 8 possible object starting orientations. Each set of 8 trials started with the object handle in the baseline handle orientation, and subsequent 7 trials progressed clockwise in the handle orientation.

4.3.4 Three levels of task constraints

The experiment consisted of three phases to examine three types of task scenarios (Fig. 4.3). The first phase served as practice trials to familiarize the participants with the task setting. Participants were instructed to complete the transport task with no restrictions, using either or both hands as desired. In the first set of 8 trials, participants transported the filled water jug from the 8 initial handle orientations. Then, in the second set of 8 trials, the filled frying pan was transported from the 8 initial handle orientations.

The second phase investigated unimanual lifting performance in response to the different object handle orientations. The purpose of this phase, as the main portion of the experiment, was to observe to what extent object pre-grasp adjustment would be used as a strategy to compensate for changes in object orientation. Participants were instructed to complete the transport task using only their right hand to contact the object. Besides this unimanual constraint, there were no restrictions on the task

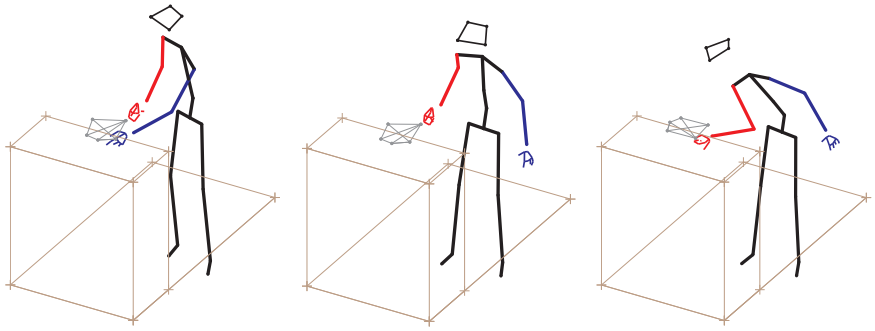

(a) (b) (c)

Figure 4.3: Participants lifted objects in three task scenarios. (a) Unconstrained practice trials where bimanual manipulation was permitted. (b) Unimanual constraint trials where only the right hand was permitted to contact the object. (c) Object motion constraint trials where lateral motion of the object on the surface was not permitted.

performance. The verbal instructions did not suggest pre-grasp object motion as a strategy, as it was our intent to observe what strategies the participants would naturally select. Two sets of 8 trials were completed for the jug to obtain a measure of performance repeatability. Next, two sets of 8 trials were completed for the pan.

The final phase tested how participants would respond to different object handle orientations in the absence of pre-grasp object adjustment on the counter surface. The task performances from this third phase provide a reference measurement for analyzing the object adjustment motion in the second phase. Participants were instructed to transport the object using only right hand contact and without any lateral sliding on the surface prior to lifting the object from the start position. Participants were given the option of aborting a trial at any time during the task if they could not complete the task under the given constraints. The first set of 8 trials was completed for the jug, and the second set of 8 trials was completed for the pan.

The experiment required approximately 90 minutes for a single participant.

4.4 Analysis of task response

Kinematic skeletal models for each participant were fit to the recorded marker trajectories as follows. The generic skeletal template modeled the lower back, shoulder, elbow, and wrist as spherical joints with three rotational degrees of freedom. An individual skeleton was automatically fit to the manually-labeled range of motion data using the Vicon IQ software's subject calibration function (Vicon Motion Systems, Los Angeles, California, USA). For the lifting task trials, the joint center locations of the individual skeletal model were fitted to the labeled marker trajectories automatically using the Vicon IQ software's kinematic fit function. Gaps in the marker

trajectories due to occlusions, which occurred frequently for the lower body markers and hand dorsum markers, were not manually reconstructed. Due to the frequent occlusions, reliable joint center locations were obtained only for the torso, shoulder, elbow, wrist, and finger knuckle joints.

4.4.1 Representative event time points

In each trial, four time points represent key events in the task performance. The key time points were determined by the following events: (1) initial body movement, (2) initial object contact, (3) object lift-off from the surface, and (4) object arrival at the goal. These mark the start and end points of three segments of the motion (Table 4.1). First, reaching and approach to the object start with the initial body movement from the participant's start position and ends when the object is first contacted ($T1$). Then, object manipulation and grasping occur between the initial object contact and the object lift-off from the surface ($T2$). Finally, object transport starts after the object lift-off and ends on object arrival at the goal position ($T3$). The four key time points were estimated automatically from the trial marker trajectories based on manually-selected thresholds, which are described in detail in Chang and Pollard [2008]. The time durations for the reaching, object manipulation, and object transport segments of the task are determined from the differences between the four key time points. Total task duration is the sum of the three segment durations.

The third key time point, object lift-off from the surface, is the focus of our data analysis. Four metrics were computed from the participant's body pose at the lift-off time frame: object rotation, joint torque load, grasp orientation, and grasp location (Table 4.1). Object rotation was measured as the change in angular orientation with respect to the horizontal plane between the object's starting orientation and the object orientation in the lift-off frame. We computed the absolute amount of rotation so that there was no distinction between clockwise or counterclockwise rotation. Upper body joint torques were estimated from the lift-off body pose as follows. The segment mass of the torso, right upper arm, right forearm, and right hand were estimated as a fraction of total body mass according to the anthropomorphic data reported in Clauser *et al.* [1969]. Locations of the segment center of mass were also estimated as a fraction of segment length according to the results of Clauser *et al.* [1969]. Given the fitted joint center locations for the lower back, shoulder, elbow, and wrist, joint torques were calculated from the loads due to distal limb segment weight and the object weight. The four joint torques were combined into a single metric of the sum of squared joint torques. The configuration of the hand dorsum coordinate system was then computed with respect to the reference coordinate system attached to the object handle (Fig. 4.2). The grasp orientation is measured as the angle magnitude of the single axis-angle rotation which would align the hand dorsum coordinate system to the object handle coordinate system. The grasp location is measured as the distance between the hand coordinate system origin at the proximal

Table 4.1: Response metrics observed for each trial where the participant lifted and transported a handled object from the presented object start orientation. The linear mixed-effects models tested the differences in each of the metrics between the unimanual constraint trials and object motion constraint trials.

Response metric	Computation notes
Timing	
Approach and reaching, $T1$	Initiation of body movement to initiation of object movement
Manipulation and grasping, $T2$	Initiation of object movement to object lift off from surface
Object transport, $T3$	Object lift off surface to object arrival at goal position
Total time, $T1 + T2 + T3$	Initiation of body movement to object arrival at goal position
Object lift-off posture	
Object rotation	Angle change at lift-off from initial orientation, on surface plane
Joint torque load	Sum of squared torques over torso, shoulder, elbow, and wrist
Grasp orientation	Angle of single rotation between hand frame and object frame
Grasp location	Distance between origins of hand frame and object frame

end of the third metacarpal and the object coordinate system origin at the base of the handle (Fig. 4.2).

4.4.2 Candidate criteria for strategy selection

The overall set of dependent variables examined in this study were difference in metrics (Table 4.1) between the tasks with the unimanual constraint (second phase) and the tasks with the additional object motion constraint (third phase). The difference in metrics were computed between matched pairs of trials performed by the same participant on the same object for the same initial handle configuration, with only a difference in the task constraint. The differences were computed for the two sets (repetitions) of 8 trials per object in the unimanual constraint phase with respect to the single set of 8 trials per object in the motion constraint phase.

4.4.3 Statistical analysis using linear mixed-effects models

We analyzed the difference metrics with linear mixed-effect (LME) models [Verbeke and Molenberghs, 2000] using the NLME package [Pinheiro and Bates, 2000] for R 2.6.2 [R Development Core Team, 2008]. Our study contains repeated measures because each participant performs the lifting task for both objects and all 8 orientations. Thus the observations are not independent, and we cannot use a standard linear analysis of variance (ANOVA) to correctly test the significance of the explanatory variables. LME models address the correlated errors from dependent observations by including random effects in addition to the mean fixed effects. In our study, the data is grouped by participant such that the LME model accounts for the correlation between the repeated observations for each participant. LME models can also handle missing observations without discarding all observations for one participant. This feature allowed us to include data from participants even if the lifting task was not completed for all handle orientations.

The explanatory variables available as fixed effects for the LME models included the object, the initial handle orientation, the square of the initial handle orientation, and the task repetition (either the first or second observation) of the unimanual constraint trials. Model selection was used to determine an appropriate combination of available explanatory variables for modeling each dependent variable. In the model selection process for a given dependent variable, multiple LME models were fitted for different combinations of possible explanatory variables. Using a standard statistical method to favor better data fits without using too many model parameters, we selected the model with the lowest Bayesian Information Criterion (BIC) score [Pinheiro and Bates, 2000] as the final LME model for that dependent variable. The t-test results for the final LME model indicate which explanatory variables were statistically significant as fixed effects. In addition, the significance of each random effect in the final LME model was checked using a likelihood L ratio test comparing the selected LME to a linear model without the random effect [Pinheiro and Bates, 2000].

Separate LME models were selected for each of the four dependent variables computed from the lift-off pose: difference in object rotation, difference in sum of squared joint torques, difference in hand grasp orientation, and difference in hand grasp position.

4.5 Natural pre-grasp rotation under unimanual constraint

When participants were only restricted by the unimanual constraint, they often rotated the object on the countertop surface to a new orientation before lifting and transporting the object to the goal. The resulting body poses at object lift-off (Fig.

4.4) were similar in terms of the upright torso orientation and object handle directed toward the participant. In contrast, when the object rotation strategy was precluded by the object motion constraint, the resulting body poses at object lift-off were more varied in the torso orientation and arm configuration. For trials where the object handle faced away from the participant, the torso was often tilted over the countertop surface with the elbow extended away from body to achieve the grasp of the object handle. One participant chose to abort one lifting trial in the object motion constraint phase after grasping and attempting to lift the pan from handle orientation 5 without object motion along the surface.

4.6 Comparison to direct grasping strategy with no object rotation

4.6.1 Timing for task completion

Differences in the timing between the unimanual constraint and object motion constraint trials suggests that rotation adjustment provides a trade-off between different segments of the movement (Table 4.2). Out of the 320 possible timing difference observations for each motion segment, only 318 could be computed because of the missing reference measurement due to the single aborted trial. Object manipulation duration $T2$ between initial object movement and object lift-off was longer in the unimanual constraint trials than in the reference object motion constraint trials for 93% of the 318 time differences, and the mean duration difference was 139 msec. This time difference is not surprising since we expect object adjustment to require additional time to only lifting the object. However, the reaching time $T1$ and object transport time $T3$ were shorter for the unimanual constraint trials than for the object motion constraint trials for 90% and 81% of the 318 trials, respectively. The mean time differences were -199 msec for reaching and -105 msec for object transport. In addition, the total task time, $T1 + T2 + T3$, from initial body movement to object arrival at the goal position was shorter for the unimanual constraint trials for 75% of the 318 trials, with a total mean duration difference of -165 msec.

4.6.2 Object lift-off posture

Under the unimanual constraint, the amount and direction of object rotation varied depending on the initial handle orientation, as did the object orientation at the lift-off time frame (Fig. 4.5). In general, the selected handle orientation for lifting were clustered in the region between handle orientation 1 and 8 on the participants' right side. Thus, the scale for the handle orientation variable is centered around the orientation halfway between orientation 1 and 8 for all LME models. The centering is achieved by re-coding orientations 1 through 8 as values 0.5, 1.5, 2.5, 3.5, -3.5,

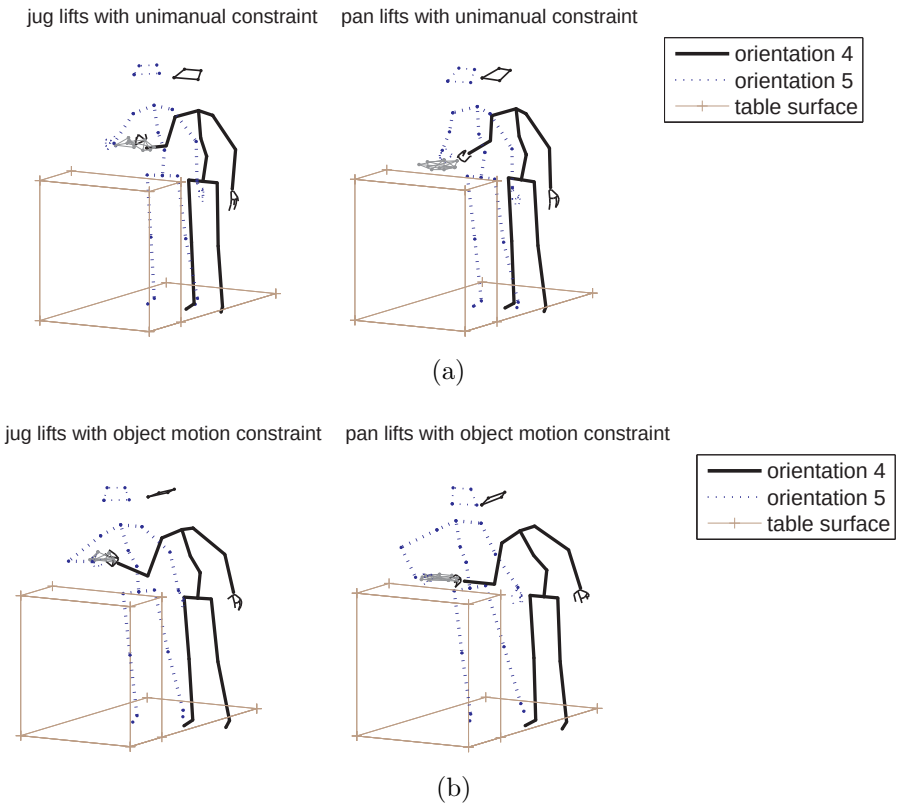

Figure 4.4: Visualization of the body postures at object lift-off for a sample subject. Poses are shown for the trials with initial object orientation 4 and 5, where the handle faced away from the participant. (a) Poses for the trials with the unimanual constraint. (b) Poses for the trials with the additional object motion constraint. The lift-off poses for the unimanual constraint trials were more similar to each other because the allowed object motion adjusted the handle direction toward the participant. When object motion was not permitted, the lift-off poses are more varied. In the most extreme cases for initial object orientation 4 and 5, the torso is tilted over the countertop and the elbow is held away from the side of the body.

-2.5, -1.5, and -0.5.

The differences in object rotation (Fig. 4.6) between the unimanual constraint trials and object motion constraint trials are largest for the initial handle orientation opposite the central handle direction facing toward the participant. The differences are small for the region near the central direction and then increase more quickly once the initial handle orientation is outside the region of preferred lifting orientations. The parabolic shape for rotation differences suggested the inclusion of the squared orientation term as a factor in the LME models. Here we highlight the key findings from the final LME models. Detailed LME results are provided in Chang and Pollard [2008].

Table 4.2: Timing duration differences for the three segments of the object transport task (Table 4.1). Each timing difference is the duration in one of the two unimanual constraint trials minus the duration in the corresponding object motion constraint trial. For each timing metric, there were a total of 318 timing differences.

Timing segment	Timing difference T(motion constraint)$-T$(unimanual constraint)	
	Mean (ms)	Count of negative differences (out of 318)
Approach and reaching, $T1$	-199	286
Manipulation and grasping, $T2$	139	21
Object transport, $T3$	-105	257
Total time, $T1 + T2 + T3$	-165	238

In the final selected LME model for object rotation with the lowest BIC score, the fixed effects of linear orientation, squared orientation, and orientation^2(object) interaction were all significant ($\alpha = 0.05$). There is slightly more object rotation for the pan than the jug for handle orientations far from the center, as modeled by the orientation^2(object) interaction effect (Fig. 4.6(b)). In addition, the result of the likelihood ratio test confirms that the random effect of squared orientation should be included in the LME model ($p < 0.0001$) in addition to the fixed average effect to account for inter-participant differences in the curvature of the object rotation trend.

The difference in the sum of squared joint torques also exhibited a quadratic trend with initial handle orientation (Fig. 4.7). The joint torque metric for the object motion constraint trials was greater than those for the unimanual constraint trials, as seen from the primarily non-negative differences (Fig. 4.7). This result suggests that there is a trade-off between pre-rotating the object prior to the lift-off and exertion of larger joint torques at lift-off. As with object rotation, the fixed squared orientation, fixed orientation^2(object) interaction, and random squared orientation were all significant effects ($p < 0.0001$). There is an increase in torque metric curvature for the pan over the jug (Fig. 4.7(b)). This increase is consistent with the increase in object rotation curvature for the pan over the jug, suggesting that the pre-grasp adjustment has a greater role in the pan lifting trials. Another result to note is the asymmetry of the regression curve fit for the jug lifting trials. The smallest difference in the torque metric occurred closer to initial handle orientation 8 rather than between orientation 1 and 8.

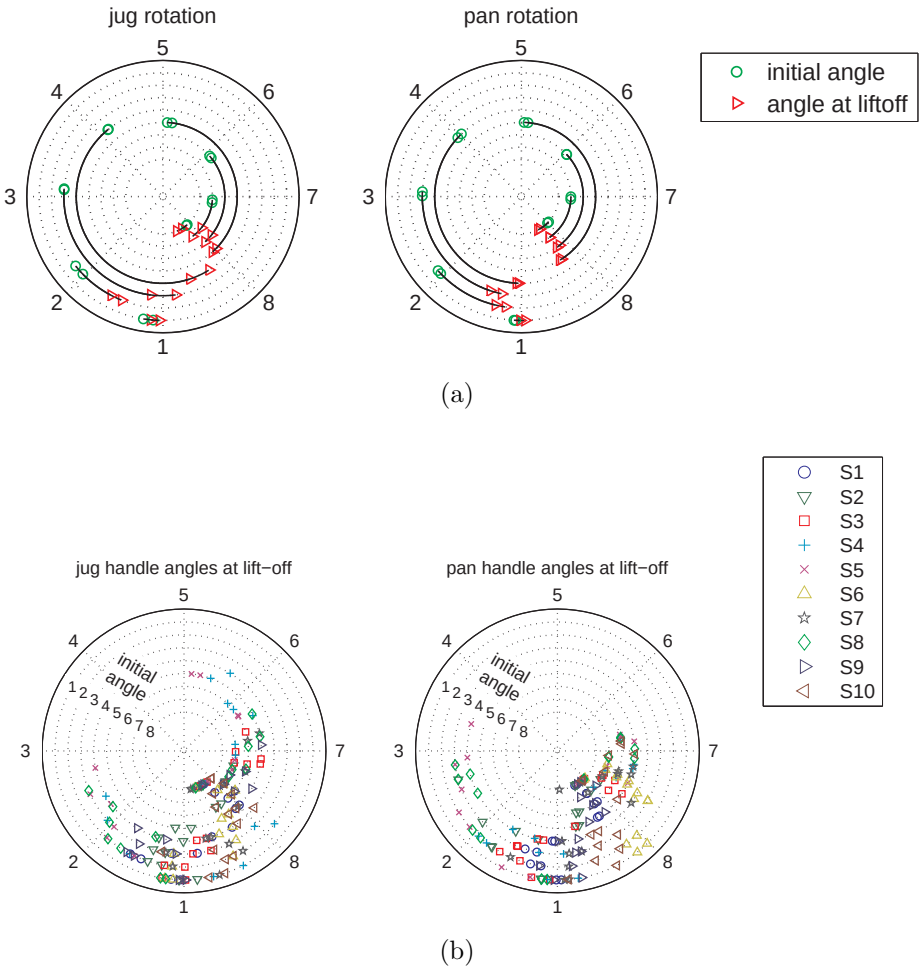

Figure 4.5: Visualization of the object rotation prior to lift-off for different initial handle orientations. (a) Sample initial and lift-off orientations for the task trials with unimanual constraint for an example participant. (b) Object lift-off angles for the unimanual constraint phase for all 16 trials for each subject.

4.6.3 Changes in hand grasp of object

Hand grasp configuration differences (Fig. 4.8) also increased when the initial handle directions faced away from the participant. The significant linear orientation effect for the grasp orientation suggests there is an asymmetry in the grasps used to the lift the object. The grasp differences are smaller for handle orientations 5 to 8, which were generally on the right hand side. Handle orientations 1 to 4 required the right hand to cross the body to reach the left side for the object motion constraint trials, which led to large differences in grasp orientation. In contrast, the response of grasp location was much more symmetric, as indicated by the absence of any significant linear orientation effects in the LME model. As with the torque metric,

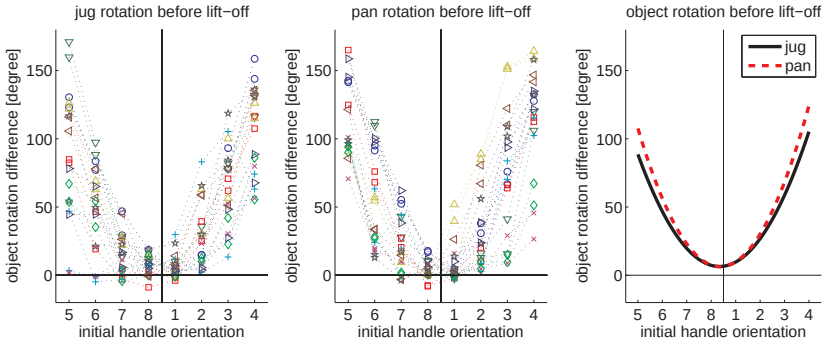

Figure 4.6: Difference in object rotation versus initial handle orientation. (a) Individual participant results. (b) Mean regression curve determined from the LME. The differences are the object rotation amounts in the two unimanual constraint trials minus the object rotation in the vertical constraint trial. The amount of object rotation prior to lift-off increases as the handle orientation moves further from the baseline orientation naturally preferred for the lifting task.

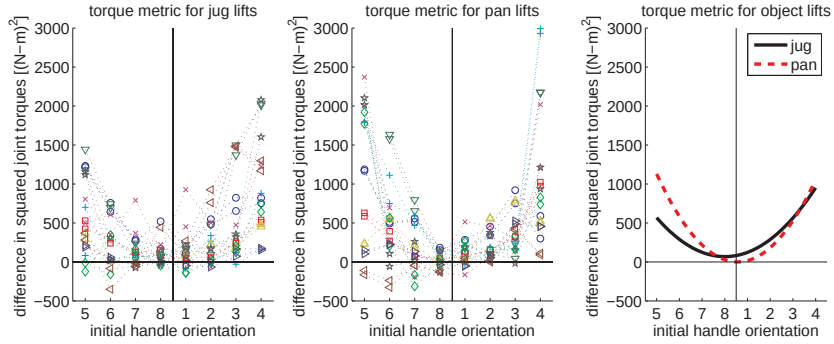

Figure 4.7: Difference in the sum of squared joint torques versus initial handle orientation. (a) Individual participant results. (b) Mean regression curve determined from the LME. The differences are the torque metrics in the two unimanual constraint trials subtracted from the torque metric in the vertical constraint trial. The trend in the torque metric is similar to that for the object rotation metric.

the difference in the trend curvature between the two objects suggests that the trade-off of pre-grasp rotation is more pronounced for the pan than for the jug. For some participants, the grasp location for lifting the pan in the object motion constraint phase changed dramatically. Instead of grasping close to the handle end as they did in the unimanual constraint phase, they lifted the pan with a grasp closer to the center of the pan when the handle was further from reach. The use of pre-grasp rotation strategy when it was permitted in the unimanual constraint phase might be

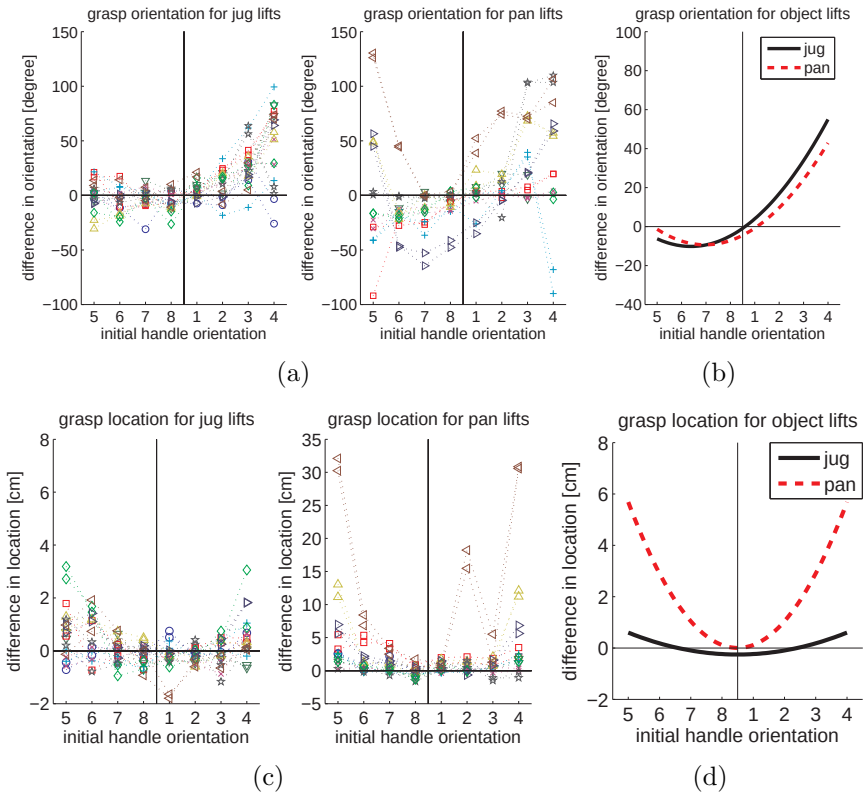

Figure 4.8: Difference in grasp as represented by the hand dorsum orientation and location with respect to the object frame. The differences are the grasp metrics in the two unimanual constraint trials subtracted from the grasp metric in the unimanual and object motion constraint trial. (a) Individual participant results for grasp orientation difference. (b) Mean regression curve determined from the LME for the grasp orientation difference. (c) Individual participant results for the grasp location difference. (d) Mean regression curve determined from the LME for grasp location difference.

due to the preference to grasp the object near the end of the handle, even though other grasps were feasible when object motion on the surface was not permitted.

For each of the four dependent variables computed from the lift-off pose, the final LME model did not include the effect of task repetition. The effect of task repetition was not statistically significant for significance level $\alpha = 0.05$.

4.7 Discussion and summary

Overall, we have found that the pre-grasp rotation of heavy objects increases with the change in handle orientation away from the central preferred direction. When participants are instructed not to pre-rotate the object prior to lift-off, they are still able to successfully complete the lifting and transport task. However, without

adjusting the object orientation prior to lifting, participants performed the lifting task with different body poses with tilted torsos and extended elbow positions in order to reach the object handle. In a comparison of the unimanual constraint trials to the object motion constraint, the differences in the joint torque metric and grasp configuration in the object frame computed from the body pose at lift-off follow a quadratic trend. This result suggests that the pre-grasp object adjustment may be desirable because it allows the object lift to be performed with lower joint torque load in the upper body and/or with a preferred grasp of the object handle.

Our experiments investigated the pre-grasp object adjustment in the specific context of right-handed lifting and lateral transport across the body. We focused on the effect of the initial object orientation on the selected body posture at object lift-off, but several other factors may affect the pre-grasp manipulation. We would expect similar adjustment strategies in other tasks with different constraints. For example, changing the location of the goal may result in a shift in the region of handle orientations at lift-off. Other task constraints to consider include whether the object manipulation is performed with the right or left hand, timing restrictions for the completion of the task, object weight, and object handle geometry.

Our analysis focused on the difference in performance in terms of metrics computed from the single time frame defined by the object lift-off. The body pose at lift-off was chosen as a representative snapshot of the performance. It may be of interest to analyze performance metrics over the entire trial for a dynamic analysis of the motor task. In addition, an altered methodology designed for specifically measuring the time course of the different movement segments would allow for extended inquiry of the cognitive models for planning and evaluating alternative manipulation strategies.

The pre-grasp rotation strategy in response to lifting a heavy object from different handle orientations is one example of how object configuration is adjusted in addition to the postural response. Other similar pre-grasp manipulation strategies include sliding, rolling, or tumbling maneuvers which re-configure the object prior to grasping. Directions for future work in robotics include examining which features of these human manipulation strategies should be imitated in humanoid robots. Some features, such as the arm configuration, might be essential for a robot to appear human-like while other features, such as optimization of the joint torques or pre-grasp object rotation to achieve a higher quality grasp, might be important heuristics for achieving robust performance in difficult task conditions. Further studies of human motor control may also uncover new concepts that may improve imitation learning for robotic manipulation.

5 Selection criteria for human performance of pre-grasp rotation

The previous chapter presented the initial experiment examining pre-grasp rotation as a consistent movement strategy used by humans for lifting heavy objects with handles (§4). Compared to a direct grasping strategy where no object adjustment is permitted, pre-grasp rotation of the object allowed for manual lifting postures with decreased upper body joint torque and more convergent grasp choices. The objects manipulated in the previous study were two everyday kitchen objects, a water jug and a cooking pan, with different geometry and different mass. The participants were also relatively unconstrained in the task precision, grasp choice, and standing position.

In this chapter we investigate how people adapt the pre-grasp object rotation strategy in response to controlled task difficulty factors and whether postural lifting strength may predict the choice of lift-off angles. In the first experiment, we examined the effect of object mass and required task precision on a lifting comfort zone, measured by the variability in object orientations at lift from the surface. Increased task difficulty, i.e. heavier object or higher precision requirements, resulted in decreased variability of lift orientation. In addition, the amount of pre-lift rotation correlated with the resulting change in the postural lifting capability, as measured for different object orientations.

In the second experiment, we compare the task completion time and success rate of the natural rotation strategy to direct grasping without pre-grasp rotation. Task completion time and success rate decreased, and initial object orientation affected pre-lift time. Our results suggest that lifting from the comfort zone produces more robust performance at a cost of slower completion; moreover, physical rotation could be replaced by mental planning when direct grasping is enforced.

The material in this chapter has been previously published in the *Journal of Motor Behavior* [Chang *et al.*, 2010a].

5.1 Related literature on movement strategies

In many scenarios, there are multiple actions which can achieve the specified task goal. The redundancy of the human motor system allows for alternative movement

strategies for completing the same task, and a key question is to understand what criteria determine the selected motor plan. This study examines criteria which could explain how people choose between alternative movements of either a single direct action or a sequence of sub-actions to complete a manual lifting task.

The experiments in the previous chapter §4 [Chang *et al.*, 2008] has shown that people sometimes choose to use an extra sub-action of pre-grasp object rotation to complete a lifting task even though a direct reach-to-grasp action would succeed. The pre-grasp object rotation was used to first physically reorient the object handle without supporting the entire gravitational load of the object. Then, the final grasp was formed to lift the object for the instructed object transport task. The strategy using the two sub-actions of a physical rotation and a lift resulted in more convergent body postures and grasp choices at the object lift-off point compared to the strategy of using a single direct lift. An open question not answered by the previous study in §4 [Chang *et al.*, 2008] is what criteria determine when and how much object rotation is incorporated into the motor plan as an alternative to a direct lifting plan.

A principal question investigated by the present study is whether posture-dependent strength capabilities can explain the preference for the pre-grasp object rotation strategy. According to the posture-based planning model [Rosenbaum *et al.*, 1993a,b, 1995], motor plans are based on goal postures rather than the movements between postures. The choice of using pre-grasp object rotation instead of a single direct action may be driven by the preference for a particular posture for the primary lifting action that achieves the task goal. The body of work on the end-state comfort effect [Rosenbaum *et al.*, 1992; Short and Cauraugh, 1997; Fischman, 1998; Zhang and Rosenbaum, 2008, e.g.,], has provided evidence that the comfort of a final goal posture affects the manual action, even if it requires awkward postures at the beginning of an action. In the previous work, end-state comfort has been evaluated by awkwardness ratings [Rosenbaum *et al.*, 1990, 1993b; Short and Cauraugh, 1997] or distance from neutral joint positions [Zhang and Rosenbaum, 2008]. In this work, we examine whether the choice of the object lift-off posture in a manual task is determined by the lifting capability of the posture. This relationship is theoretically important because it contributes further understanding of the biomechanical aspects that may be associated with psychophysical assessment of postural comfort. Individual strength has been investigated previously for predicting choices between alternative movement strategies of stooping versus squatting bimanual lifting tasks [Burgess-Limerick and Abernethy, 1997], but the focus was on strength capabilities between individuals rather than the differences between postures for one individual.

A second question investigated in this study is whether pre-grasp object rotation is used to achieve an object lift-off posture which is optimal versus satisfactory for the task. If individuals maximize the efficiency of the lift, they would tend to select postures with the maximum lifting capability even for different levels of task precision or object mass. However, if only a minimum efficiency threshold has to be satisfied, then there may be a larger range of selected lift-off postures when the task

is less difficult. In this investigation, we examine whether the preference for selected lifting postures is sensitive to two task difficulty factors: object mass and task precision. Previous work by Latash and Jaric [2002] studied how postural movements in drinking tasks were strongly correlated with the task precision, as defined by a geometric task difficulty index related to the critical angle for spilling. Results from Rosenbaum *et al.* [1996] also showed that the end-state comfort effect is diminished for tasks requiring less precision. We believe that in addition to task precision, object mass is also a factor which affects posture selection, especially for more demanding manual tasks. Manipulation requires not only meeting specific kinematic constraints specified by the task, but also applying or resisting forces due to the interaction with the object. The ability to meet the load requirements for a task are posture dependent, since muscle strength varies for different operating postures. Thus tasks involving heavier objects may result in the selection of a posture with greater lifting capability in order to satisfy a minimum efficiency or strength margin constraint. Our study expands on the previous literature by examining how posture selection may be sensitive to both task precision and object load.

5.2 Study objectives

Experiment 1 in this study investigates the questions of (1) whether the pre-grasp rotation is correlated with the change in lifting capability, and (2) whether the selected lifting postures are relatively constant or vary with task difficulty. First, we hypothesized that people would choose to use the pre-grasp rotation to avoid lift-off angles with lower lifting capability in favor of selected lift-off angles with higher lifting capability. This hypothesis predicts that the amount of object rotation would be correlated with the difference in lifting capability between the initial orientation and the lift-off orientation. Second, we examine the effects of task precision and object mass on the task performance. The contribution of these two factors to the task difficulty factors was evaluated by the effect on total time for task completion. It was expected that the selection of lift-off postures would be more constrained for increased task difficulty due to more stringent satisfaction requirements. This effect could be observed in the lift-off object configuration in two ways. The amount of object rotation prior to lifting could increase with increased object mass or increased angular precision. In addition, the variability in the object lift-off angles could decrease with increased object mass or increased angular precision.

It is possible that lifting capability represents a "hard", or immutable, constraint on the movement strategy selection rather than a "soft" constraint or preference. Pre-grasp object rotation may be selected because the task is not feasible using the alternative strategy of direct reach-to-grasp lifting, where the object is not rotated before lift-off. In Experiment 2, we examined whether the manual task was feasible using only direct grasping without object rotation. We found that for most of the

task conditions, participants could successfully complete the task using only direct grasping even though pre-grasp rotation was spontaneously used in Experiment 1 for the same conditions. Moreover, the dependence of physical rotation on the initial object orientation, as found in Experiment 1, was echoed in Experiment 2 by a similar trend in pre-lift preparation time, suggesting that the physical action was replaced by mental planning. The comparison of Experiment 1 results with the constrained condition in Experiment 2 also allows us to investigate how the posture selection for the lifting action affects other aspects of the movement, such as timing and robustness of task performance. Our results show that direct lifting (when successful) requires less time for task completion, but lifting after pre-grasp rotation is more robust for completing the task trials successfully.

5.3 Experiment overview

The main portion of the study, Experiment 1, examined what factors affected the selected lift-off postures when participants used pre-grasp object rotation. In addition, we investigated the relation between the lifting capability and the amount of object rotation to a new lift-off angle.

Experiment 1 consisted of two parts. In Part I, participants had to lift a canister, selected from a set of 4, to uncover a token on the table surface. The canister was presented at multiple initial orientations, and it was expected that participants would use pre-grasp object rotation before lifting the canister. The dependent variable of interest from Part I was the selected object orientations at the lift-off time. The distribution of lift-off postures was measured for different levels of object mass and task precision.

In Experiment 1 Part II, participants performed similar grasping actions on a weighted canister. Lifting capability was measured for different object orientations to investigate whether increase in lifting capability was correlated with the amount of pre-grasp object rotation in Part I.

The second portion of the study, Experiment 2, examines whether the token retrieval task is feasible when the participant is restricted to using a direct grasping strategy without object rotation. In addition, we compared the timing of the successful trials to examine whether either strategy offered an advantage in completing the task more quickly.

In contrast to the initial study presented in §4, this study observes the pre-grasp object rotation strategy for a highly-constrained set of task conditions. The task is constrained in order to reduce inter-subject variability in the choice of standing position and grasp type. Participants attempt to perform a canister lifting task while standing in place before a table. The grasp contact area on the canister is specified. The task must be completed while maintaining the vertical orientation of the canister within a specified angular precision. Geometrically-identical handled

canisters are used to test the response to the controlled object mass and task precision requirements.

5.4 Unimanual lifting task of a handled canister

In this section we describe the experimental apparatus and environment which is used for both Experiment 1 and Experiment 2. There were two types of tasks tested in the experiments. In the token retrieval task for Part I, participants had to lift a canister, selected from a set of 4, to uncover a token on the table surface. In the lifting interaction task for Part II, participants repeated the grasping motion on a weighted canister which was too heavy to lift. In Part II, lifting capability was measured for different object orientations to investigate whether lifting capability influenced the pre-grasp object rotation in Part I.

5.4.1 Apparatus

The object for the manipulation task was a plastic cylindrical canister (Fig. 5.1). Two indentations in the side of the canister formed a handle which could be grasped by inserting the thumb in one indentation and the fingers in the other indentation (Fig. 5.1(a)). A lightweight plastic ball (diameter 3.8 cm) was balanced on a cylindrical ring which was affixed to the canister lid. The diameter D of the ring determined the required angular precision in the object's vertical orientation for keeping the ball balanced (Fig. 5.1(b)).

The canister was presented to the participant on a flat platform at a height of 0.9m from the ground (Fig. 4.1). A screen was placed between the subject and the platform (Fig. 5.2(a)). Participants stood facing the screen with their hands relaxed at the side. Before a trial, the screen started in the lowered position so that the participant could not see the object placed on the platform. When a trial commenced, the screen was completely raised to show the object on the platform.

A T-shaped marking on the floor indicated the standing position (Fig. 5.2(b)). Participants stood in place with their feet centered on the vertical line of the T and their toes just behind the horizontal line of the T. The T marking was offset from the table such that the canister position was 0.1m to the right and 0.2m in front of the T intersection. The offset position was chosen for comfortable lifting of the canister using the right hand (Fig. 5.2(b)).

Passive reflective markers were used for tracking the object motion in Part I. The screen was tracked by 5 markers and each canister object was tracked by 6 markers. The 3D positions of the markers were recorded at 120 Hz using a Vicon camera system (Vicon Motion Systems, Los Angeles, California, USA). The 3D marker trajectories were reconstructed using ViconIQ software and were not filtered or processed in any way before the data analysis described in Section 5.6.1.

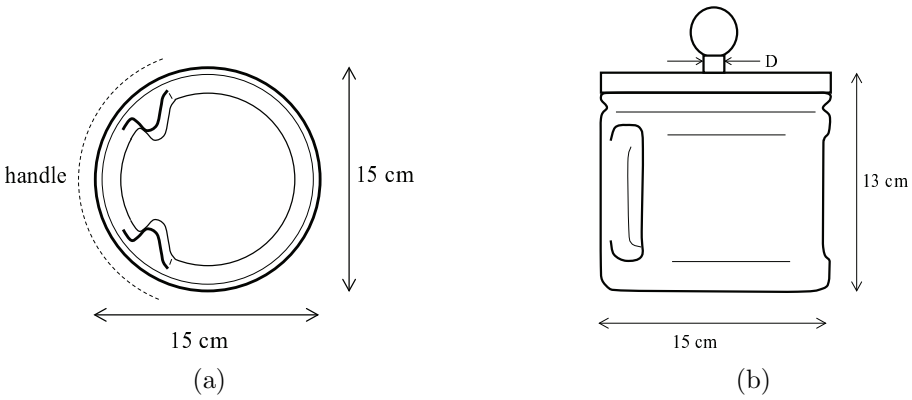

Figure 5.1: (a) Top view of the canister interior. The handle, denoted by the dotted arc on the left, is formed by two indentations in the body. (b) Side view of the canister, with a ball balanced on the lid. The handle is on the left in the diagram. There were two levels for the diameter D of the balance ring, which was used to manipulate the task difficulty factor of required angular precision with respect to the vertical orientation.

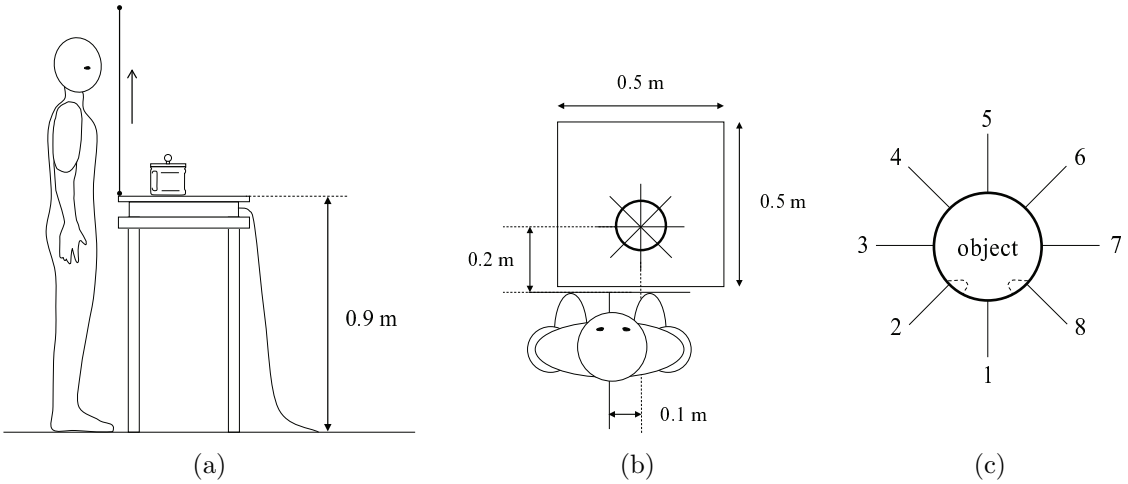

Figure 5.2: General layout of the experimental setting. (a) Participants started in a standing position facing the tabletop platform. A vertical screen blocking the view of the object was raised at the beginning of each trial. A scale under the platform measured the vertical load due to the object weight. (b) Participants stood slightly to the left of the platform and lifted the canister by the handle with their right hand. A T marking on the floor denoted the foot placement relative to the platform. (c) In each trial, the handled object started in one of eight orientations defined by the handle direction. In the figure, the handled object is in orientation 1, where the handle directly faces the participant.

Underneath the platform was a uni-directional measuring scale (0.02kg resolution) for digitally recording the lifting capability measurement in Part II (Fig. 5.2(a)).

The scale data were recorded separately from the kinematic data using MATLAB software (Mathworks, Inc.; Natick, MA). The scale output was queried repeatedly via the serial port interface such that the time periods between data samples were approximately 0.2-0.3 seconds.

5.4.2 Part I: Lifting motion for token retrieval task

In Part I of Experiment 1, a flat disc (diameter 2.5cm) was placed on the platform underneath the canister. The token retrieval task was to lift the canister in order to uncover the disc, while keeping the ball balanced on the cylindrical ring. Participants were instructed to use only the right hand to contact the canister and were restricted to contacting the canister at the handle area (Fig. 5.1(a)). Only the left hand could be used to the contact the token. Participants were free to move either the canister and/or the token from the original position, as long as at the end of the trial, both the canister and the token were placed on the platform and the canister did not cover the token. Participants were instructed to perform the task as quickly as possible while keeping the ball balanced.

There were 4 versions of the canister to test 2 levels of object mass and 2 levels of angular precision. For the object mass factor, the canister was filled with material to achieve a total mass of either 0.40kg (light) or 1.20kg (heavy). The canisters were covered by two different colors of tape to provide visual labels corresponding to the light canister (grey) and heavy canister (blue). For the angular precision, the cylindrical balance ring had diameter D of either 3.6cm (wide) or 0.8cm (thin). The diameter of the ring was visible to the participants when the object was presented on the platform.

The canister was presented on the platform in one of 8 orientations, as denoted by the handle direction (Fig. 5.2(c)). It was expected that the participant would slide or rotate the canister before lifting to adjust the orientation when the handle faced away from the participant. In the verbal instructions for the token retrieval task, there was no suggestion of pre-grasp rotation or sliding.

5.4.3 Part II: Measuring lifting capability

In Part II of Experiment 1, the maximum lifting capability was measured for different handle orientations in order to investigate whether lifting capability was a criterion for the pre-grasp object rotation observed in Part I. The canister object was placed on the platform and the participants stood at the T-marking as before. The task was to repeat the lifting grasp on the object in order to measure the lifting capability for a specific handle orientation. Participants were instructed to attempt to reach the presented configuration of the canister handle using the right hand and apply as much upward lifting force as was comfortable. Unlike Part I, participants were instructed not to move or reorient the canister from its presented configuration. Participants

were allowed to skip or abort the trial at any time if they could not reach the handle or apply upward force comfortably.

The filled canister had a total mass of of 13.7kg, which was beyond the unimanual lifting capability of any participant. The lifting capability was measured as the difference between the initial scale output and the minimum scale output during the grasp. To mimic the perception of the canister as experienced in the Part I lifting trials, a ball was balanced on the canister lid as in Part I. As before, the cylindrical balance ring had diameter D of either 3.6cm (wide) or 0.8cm (thin). There was no token to retrieve with the left hand.

The canister was presented in 1 of 10 handle orientations: the 8 orientations from Part I (Fig. 5.2(c)) and 2 additional orientations at the midpoint between angle 7 and 8 (denoted 7.5) and the midpoint between angle 8 and 1 (denoted 8.5). The additional orientations were sampled because it was hypothesized that the maximum lifting capability would occur within the region of selected lift-off angles. The lift-off angles were expected to be in the region between handle orientations 1 and 7, where the handle faced toward the participant's right hand.

5.5 Experiment 1: Pre-grasp rotation strategy

Experiment 1 examined the participants naturally-selected manipulation strategy. In Part I, the canister was presented at multiple initial orientations, and it was expected that participants would use pre-grasp object rotation before lifting the canister. In Part II, the participant's lifting capability was measured for different initial object orientations.

We hypothesized that task timing would increase for a more difficult manipulation task involving a heavier object or requiring more angular precision during the lifting. It was also expected that increased task difficulty would result in a higher constraint on the choice of object orientation for lifting due to the more demanding task conditions, which would affect the observed object configuration in two ways. First, we hypothesized that the amount of object rotation prior to lifting would increase for increased object mass or increased angular precision. Second, we hypothesized that the variability in the lift-off angle would decrease for increased object mass or increased angular precision. These hypotheses are investigated in Experiment 1 Part I.

5.5.1 Participants

A total of 12 adults (6 male, 6 female) volunteered for the study (age = 26.4 ± 4.9 years [mean ± standard deviation], height = 1.70 ± .01 m, mass = 64.1 ± 12.4 kg). All participants were right-handed by self-report. All participants signed informed consent forms approved by the Institutional Review Board of Carnegie Mellon University.

5.5.2 Experimental procedure

Before the trials recorded in Part I, the participants were first introduced to the canisters to become familiar with the mass levels and angular precision levels. Each of the 4 canisters was held, one at a time, with the ball balanced on the ring. There were no other practice trials for task training before the main experiment. The experimenter drew attention to the color coding of the object mass levels at this time.

There were a total of 32 trials for Part I, with one trial for each possible combination of the 2 mass levels, 2 angular precision levels, and 8 initial handle orientations. The 32 conditions were presented in a randomized sequence. If the ball was dropped before the canister was returned to the platform, the trial was aborted. A new randomized sequence was generated for the remaining conditions, including the failed condition. Thus the participant could not anticipate the trial conditions before the screen was raised at the beginning of each trial. After every 8 trials, participants were given a break where they were seated for approximately 1 minute. During the break, the experimenter asked the participant to report any fatigue from prolonged standing or repeated lifting.

In Part II for the measurement of the lifting capability, there were 10 trials for each of the 2 angular precision levels, as determined by the balance ring diameter. The first set of 10 trials tested the lifting capability for one angular precision level, followed by a second set of 10 trials for the other angular precision level. The order of angular precision levels alternated based on the order of participation. Thus, 6 participants were presented with the wide diameter ring first, and 6 participants were presented with the thin diameter ring first. Within each set, the 10 possible handle orientations were presented in a randomized sequence. We did not expect that the ball would drop due to the canister weight and the restriction to leave the canister configuration unchanged from its presented condition. However, if the ball was dropped, the failed condition would be repeated within a newly randomized sequence of the remaining conditions. A break was given between the two sets of 10 trials.

In addition to the scheduled breaks, the protocol included immediate rest if the participant reported fatigue at any time during the study. The entire experiment required approximately 90 minutes for a single participant.

5.6 Analysis of task response

5.6.1 Representative event time points

The selected object orientation at the lift-off time was the variable of interest from the lifting trials in Part I. In addition, the timing of the task completion was measured to characterize the difference in task difficulty for the tested conditions. The lift-off

angle and timing measures were computed from segmentation of the motion data for the canister and the screen. To obtain this motion data, the recorded 3D marker trajectories were first manually labeled and then registered to a rigid-body model of each object to find the object position and orientation at each time frame. To reduce noise in the tracking estimates, the model registration used all available markers (up to 5 markers for the screen and 6 markers for each canister) per time frame. The object configuration data was always computed from at least 4 markers.

For each trial, three time intervals comprise the total time (T_0) for task completion (Table 5.1). First, pre-interaction (T_1) consists of perception and reaching for the object. T_1 starts when the screen is raised and ends at the initial canister movement. Then, object interaction (T_2) occurs if there is optional adjustment of the canister configuration on the surface. T_2 starts at the initial canister movement and ends at the object lift-off from the surface. Finally, object lifting and token retrieval (T_3) starts after the object lift-off and ends when the object is returned to the surface.

The four key time points defining the time intervals were estimated automatically from the kinematic data based on manually-selected thresholds. (1) The start of the trial was determined as the time frame when the lower edge of the screen was raised above the canister to reveal the task conditions of mass level, angular precision level, and handle orientation. (2) Initial canister movement was detected when the average object marker difference from the starting configuration exceeded 5mm. (3) Object lift-off from the surface was detected when the upward vertical displacement of the canister position exceeded 5mm change from the initial vertical position. (4) Task completion when the object is set down was detected when the vertical displacement of the canister fell within 5mm of the ending vertical position.

5.6.2 Object rotation and lift-off angle

The third key time point, the object lift-off from the surface, was used to analyze the amount of pre-grasp object rotation. The object lift-off angle is the orientation of the object in the horizontal plane at the lift-off time frame. Object rotation was measured as the change in orientation between the initial handle angle and the lift-off handle angle. We computed the absolute amount of rotation so that there was no distinction between clockwise or counterclockwise rotation.

We also measured the variability of the selected lift-off angles over the set of 8 initial handle orientations for each of the 4 canisters. The average absolute deviation (AAD) is computed as the mean deviation of a lift-off angle from the mean of the set of 8 lift-off angles. It was hypothesized that for tasks which are more difficult, the AAD of the lift angle would be smaller because participants would respond to task constraints with more pre-grasp rotation, in order to grasp the handle from a preferred orientation.

Table 5.1: Response metrics observed for each trial of the token retrieval task in Experiment 1 Part I, where the participant lifted the canister from a presented object start orientation. Two groups of response metrics examined the object configuration choices for the pre-grasp rotation strategy and the timing of the task completion.

Response metric	Computation notes
Object configuration	
Lift-off angle	Handle angle in horizontal plane at lift-off frame
Pre-grasp rotation	Angle change between starting angle and lift-off angle
Timing	
Pre-interaction, T_1	Raising of screen to initiation of object movement (includes object perception and reaching)
Object interaction, T_2	Initiation of object movement to object lift-off from surface (pre-grasp adjustment)
Object lifting, T_3	Object lift off to object return to surface (includes token retrieval)
Total time, $T_0 = T_1 + T_2 + T_3$	Raising of screen to object return to surface

5.6.3 Statistical analysis of object mass and task precision

We analyzed the object orientation and timing metrics with linear mixed-effects (LME) models [Verbeke and Molenberghs, 2000] using the NLME (nonlinear mixed-effects) package [Pinheiro and Bates, 2000] for R 2.6.2 [R Development Core Team, 2008]. Similar to ANOVA tests for repeated measures, LME models account for the correlated errors between the dependent, repeated observations from an individual participant. The advantage of LME models is the greater statistical power in estimating the significance of the fixed effects compared to a repeated measures ANOVA, which is due to the LME models' simplified parameterization of the random effects' correlation structure. In our study, the data were grouped by participant such that the LME model accounts for the correlation between the repeated observations by fitting an individual intercept value for each participant. LME models can also handle missing observations without discarding all observations for one participant. This feature allowed us to include data from participants even if the lifting task were not completed for all handle orientations.

The t-test results in an LME model indicated which explanatory variables were statistically significant as fixed effects. In addition, the significance of each random effect in the final LME model was checked using a likelihood ratio (L ratio) test comparing the selected LME to the linear ANOVA model without the random effect

[Pinheiro and Bates, 2000].

The explanatory variables that were tested as fixed effects in the LME models included the object mass, angular precision as determined by balance ring diameter, the initial handle orientation, and the square of the initial handle orientation. The square of the initial handle orientation was included because it was expected that the time and amount of pre-grasp rotation would increase for initial angles further from some central preferred lift-off angle. It was also expected that the preferred lift-off angle would be around orientation 8, where the handle faced toward the participants' right side. Thus, the scale for the handle orientation variable was centered at orientation 8 for all LME models. The centering was achieved by re-coding orientations 1 through 8 as values -1, -2, -3, 4, 3, 2, 1, 0.

Task difficulty was measured by the total time for task completion. It was expected that the token retrieval task would be more difficult for increased object mass and/or for increased task precision. That is, even for the central orientation 8, the time duration would be longer for lifting the heavier canister and/or balancing the ball on the thin diameter ring. Thus the LME models for timing tested the task difficulty factors of object mass and angular precision as additive main effects.

For pre-grasp rotation prior to lift-off, we hypothesized that the amount of rotation would increase with task difficulty more for the initial handle orientations which are further from the central orientation. That is, for the central orientation, the object rotation would be small regardless of the canister mass and the balance ring diameter due to the handle presentation already being a preferred lifting configuration. For initial orientations outside the preferred region, there would be more rotation for increased task difficulty due to a stronger incentive to move the handle into the preferred region. Thus the LME model for object rotation included the interaction effects of the object mass with the squared orientation factor and the angular precision level with the squared orientation factor.

A corresponding hypothesis for object lift-off configuration was that the participants would select the lift-off angle of the handle with more consistency for more demanding tasks. That is, with increased task difficulty, the variability of the lift-off angle would be smaller. The AAD measure was used for this test of lift-off angle variability. Since a single AAD measure was computed over the set of 8 initial handle angles for each of 4 canisters, only object mass and angular precision were tested as additive fixed effects in the LME model.

5.6.4 Lifting capability

The purpose of measuring the canister lifting capability in Part II was to investigate whether the amount of rotation in Part I was correlated with the change in the lifting capability for a given handle orientation.

The lifting capability for a given handle orientation was computed as the difference between the initial scale measurement (when only the object mass contributed

to the scale reading) and the minimum scale measurement (when the participant applied the most upward force on the canister). A lifting capability profile for the entire 360-degree orientation range was computed by linear interpolation of the 10 samples measured in Part II. There were two profiles for each participant because the lifting capability was measured for the two levels of balance ring diameter. The maximum lifting capability was the largest measurement over the 10 samples, and it was used to normalize the profile magnitude. Thus, all lifting capability measures were expressed as a percentage of the maximum capability.

The lifting capability at the 8 initial handle angles was directly measured as part of the 10 samples. In contrast, the selected lift-off angles were different for each individual, and the lifting capability was not measured directly. For one level of balance ring diameter, there were 8 actual lifts in Part I for each of two object mass levels. The lifting capability associated with each actual lift angle was estimated from the linearly-interpolated lifting capability profile with the matching task precision level.

We hypothesized that the pre-grasp rotation action in Part I changed the orientation to improve the lifting capability for the primary lifting action. Thus, a larger amount of rotation would correspond to a larger increase in the lifting capability. For each individual trial, we computed the difference in the lifting capability at the selected lift-off angle relative to the lifting capability at the initial object angle. The data tested by the correlation computation were the mean values averaged over the 12 participants. For each of the 32 task conditions, the mean amount of rotation was paired with the mean difference in lifting capability.

In addition, for each of the 4 combinations of object mass and task precision, we computed the minimum and the AAD variability of the capability measures at the lift-off angles for the 8 trials. We expected that increased task difficulty would result in greater restrictions on the preferred lifting capability at the selected lift-off angles. Thus, similar to our hypothesis that the lift-angle variability from Part I would decrease with increased task difficulty, we hypothesized that the variability in the lifting capability measures for the lift-off angles would also decrease with increased task difficulty. This hypothesis was tested by a LME model with object mass and angular precision as additive fixed effects.

Finally, the normalized lifting capability profiles were used to compute what we call the potential lift-angle region. The potential lift-angle region is the set of all object orientation angles whose corresponding normalized lifting capability measure was greater or equal to the minimum capability. In this case, the minimum lifting capability ratio was the smallest computed ratio over the 16 trial conditions for the specific level of angular precision. All 16 of the actual lift-offs occurred at a handle angle corresponding to a lifting capability greater or equal to the minimum ratio. All potential lift-off angles also corresponded to a lifting capability ratio greater or equal to the minimum ratio. We compared the potential lift-off range to the actual lift-off range to investigate whether the minimum lifting capability was sufficient to

predict the actual lift-off angles.

5.7 Experiment 1 results

For both parts of the experiment, no participant reported fatigue or requested immediate rest in addition to the scheduled breaks.

5.7.1 Failures at task completion

For the lifting task in Part I, participants sometimes failed to maintain the angular precision constraint such that the balanced ball was dropped during the trial. Out of the 12 participants, 4 participants failed for 1-3 trials, 6 participants failed for 6-8 trials, and 2 participants failed for 10-12 trials. The conditions for the failed trials were repeated in a random order as described in the procedure. For two participants, there was no recording of successful task completion for 1 of the 32 conditions due to experimenter error. For one participant, the recording of a successful lift was incomplete for 1 condition due to equipment error. These unavailable conditions were omitted in the LME model such at that there are a total of 382 (instead of 384) data points for the object rotation, time T_1, and time T_2. There are a total of 381 data points for time intervals T_0 and T_3 due to the incomplete trial recording.

5.7.2 Task difficulty increased completion time

The total time T_0 for task completion (Fig. 5.3) was 3.46 seconds on average for the baseline light object with wide diameter ring at initial orientation 8 (Table 5.2). Increasing the object mass increased T_0 by 0.73 seconds ($t(365) = 4.00$, $p < 0.001$) and increasing the angular precision increased T_0 by 2.96 seconds ($t(365) = 16.22$, $p < 0.001$). The task also took longer when the handle orientation started further from the central orientation 8, as indicated by the positive coefficient for the squared orientation term ($t(365) = 17.60$, $p < 0.001$). The significant effect on the linear orientation term ($t(365) = -5.66$, $p < 0.001$) indicates that the quadratic trend is not symmetric around orientation 8. From the linear orientation and squared orientation coefficients, we computed that the minimum of the quadratic trend occurs at about 20 degrees counterclockwise from orientation 8 (Fig. 5.3).

5.7.3 Task difficulty decreased variability in lift-off angle

The amount of pre-grasp object rotation (Fig. 5.4) increased by 7.7 degrees for the increased angular precision corresponding to the thin diameter balance ring ($t(364) = 2.21$, $p = 0.0276$) (Table 5.3). The positive coefficient for the squared orientation ($t(364) = 16.62$, $p < 0.001$) indicates that the amount of rotation increased as the initial object orientation was further from the central orientation 8. The positive

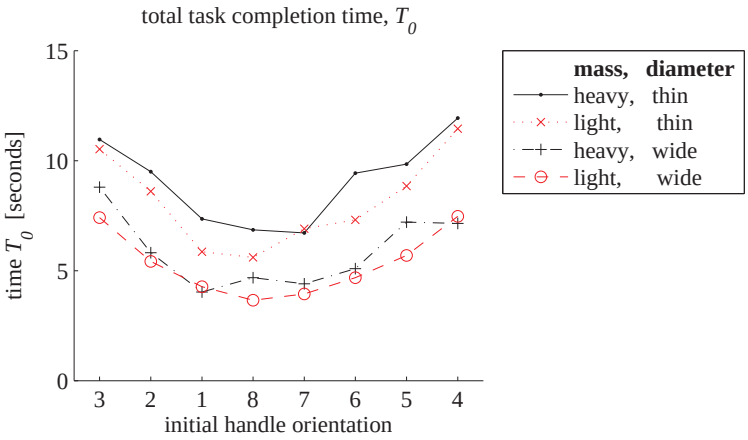

Figure 5.3: Mean data for the total time T_0 for task completion with pre-grasp rotation. The LME model results (Table 5.2) indicate that task completion time was longer for the heavy object mass and/or the increased angular precision corresponding to the thin diameter balance ring. The task also took longer if the handle orientation started farther from the central orientation 8, as indicated by the quadratic trend with respect to the initial handle orientation.

Table 5.2: Effects on task completion time (seconds) resulting from the linear mixed-effects (LME) analysis. The t-test and L ratio test results indicate the significance of the fixed effects and random effects, respectively. Significant effects ($p < 0.05$) are denoted by asterisks (*). In the model, the baseline mass level was light (L) and the baseline diameter for angular precision was wide (W). Thus the mean value for the mass and angular precision are the additive effects for the heavy (H) level and the thin (T) ring diameter.

Main effects	Mean (s)	Std. Error (s)	t(DF)	p
T_0: total task completion time			$t(365)$	
Baseline (L=0, W=0)	3.46	0.61		
Object mass (H=1)	0.73	0.18	4.00	<0.001*
Angular precision (T=1)	2.96	0.18	16.22	<0.001*
Orientation^2 coefficient	0.35	0.02	17.60	<0.001*
Orientation coefficient	-0.25	0.04	-5.66	<0.001*

Random effects	Baseline intercept (s)	Residual (s)	L ratio	p
Participant	2.00	1.78	248.85	<0.001*

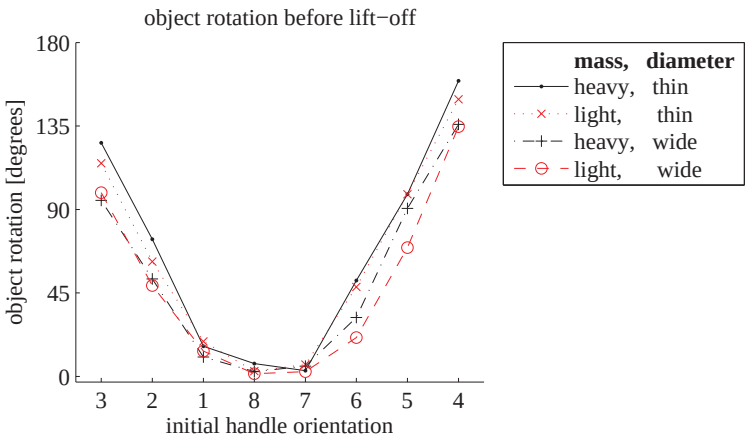

Figure 5.4: Mean data for the amount of pre-grasp object rotation. The LME model results (Table 5.3) indicate that at the central orientation, there was 7.7 degrees more object rotation for the increased angular precision corresponding to the thin diameter balance ring. In addition, amount of rotation had a significant quadratic trend with respect to the initial handle orientation, which is consistent with expectations for increased object rotation with increased distance from the preferred central orientation.

coefficient for the interaction effect of angular precision with squared orientation ($t(364) = 2.33$, $p = 0.0206$) indicates that there was an increase in the trend's quadratic curvature for the thin diameter balance ring compared to the wide diameter balance ring.

The average angular deviation (AAD) from the mean lift angle (Fig. 5.5) was 43.0 degrees for the baseline light object with wide diameter balance ring. The lift-off angle variability decreased by 4.6 degrees for increased object mass ($t(34) = -2.12$, $p = 0.0414$) and decreased 12.9 degrees for increased angular precision ($t(34) = -5.90$, $p < 0.001$) (Table 5.4).

We also performed the test on lift-angle variability on a subset of the data, excluding the two participants who failed and repeated the most trials (10 and 12 trials) in Part I. For the subset of the remaining 10 participants' lift-angle AAD measures, the lift-angle variability decreased by 7.1 degrees for increased object mass ($t(28) = -5.62$, $p < 0.001$) and decreased by 11.0 degrees for increased angular precision ($t(28) = -3.29$, $p = 0.0027$).

5.7.4 Correlation of object rotation with lifting capability increase

The maximum lifting capability (Fig. 5.6(a)) ranged from 1.5kg to 8.2kg. The maximum lifting capability for the wide diameter ring was greater than that for the thin diameter ring for 10 of the 12 participants. The maximum measurement

Table 5.3: Interaction effects on amount of pre-grasp object rotation (degrees) resulting from the linear mixed-effects (LME) analysis. The t-test and L ratio test results indicate the significance of the fixed effects and random effects, respectively. Significant effects ($p < 0.05$) are denoted by asterisks (*). In the model, the baseline mass level was light (L) and the baseline diameter for angular precision was wide (W). Thus the mean value for the mass and angular precision are the additive effects for the heavy (H) level and the thin (T) ring diameter.

Fixed effects	Mean (deg)	Std. Error (deg)	$t(DF)$	p
Main effects			$t(364)$	
Baseline (L=0, W=0)	-1.3	3.1	-0.41	0.6818
Object mass (H=1)	2.1	3.5	0.61	0.5427
Angular precision (T=1)	7.7	3.5	2.21	0.0276*
Orientation^2 coefficient	9.6	0.6	16.62	<0.001*
Orientation coefficient	-4.6	0.6	-7.94	<0.001*
Interaction effects				
Object mass * Orientation^2	0.4	0.5	0.76	0.4503
Angular precision * Orientation^2	1.1	0.5	2.33	0.0206*

Random effects	Orientation^2 coefficient (deg)	Residual (deg)	L ratio	p
Participant	1.3	23.3	38.28	<0.001*

occurred in the region between orientations 1 and 7.5 for all participants for the wide diameter angular precision level. For the thin diameter angular precision level, the maximum capability measurement occurred between orientations 1 and 7.5 for 10 of the 12 participants. In the two exception cases, the second largest measurement occurred between orientations 1 and 7.5, and it was at most 0.18kg less than the maximum measurement. Figure 5.6(b) shows that the average lifting capability profile is asymmetric around the region of maximum capability, since the lifting capability decreases more steeply for orientations to the left of the central region.

Figure 5.7 shows the correlation between the amount of object rotation and the normalized lifting capability. The data are divided into two groups due to the asymmetry of the left region (initial orientations 1-3) and right region (initial orientations 4-8) of the object orientation scale . For the left region, where the object rotation was counter-clockwise, the Pearson correlation coefficient was $r(10) = 0.9319$ ($p < 0.001$).

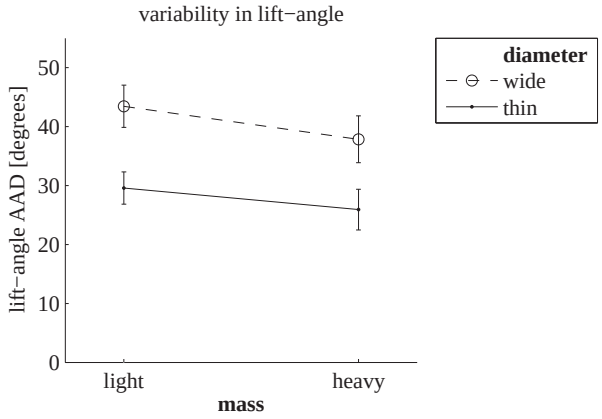

Figure 5.5: Mean data for lift-angle variability, as measured by the absolute average deviation (AAD). The error bars indicate ±1 standard error of the mean estimate. The LME model results (Table 5.4) indicate that the lift-angle variability decreased 4.6 degrees for the more massive object and 12.9 degrees for the increased angular precision corresponding to the thin diameter balance ring.

Table 5.4: Main effects on lift-angle variability resulting from the LME analysis. Lift-angle variability is measured by the average absolute deviation (AAD) of lift-angles over the set of 8 initial handle orientations. The t-test and L ratio test results indicate the significance of the fixed effects and random effects, respectively. Significant effects ($p < 0.05$) are denoted by asterisks (*). In the model, the baseline mass level was light (L) and the baseline diameter for angular precision was wide (W). Thus the mean value for the mass and angular precision are the additive effects for the heavy (H) level and the thin (T) ring diameter. The negative values for the additive effects indicate the lift-angle variability decreases for increased object mass and increased angular precision.

Main effects	Mean (deg)	Std. Error (deg)	$t(DF)$	p
Lift-angle variability			$t(34)$	
Baseline (L=0, W=0)	43.0	3.3		
Object mass (H=1)	-4.6	2.2	-2.12	0.0414*
Angular precision (T=1)	-12.9	2.2	-5.90	<0.001*

Random effects	Baseline intercept (deg)	Residual (deg)	L ratio	p
Participant	9.3	7.6	19.16	<0.001*

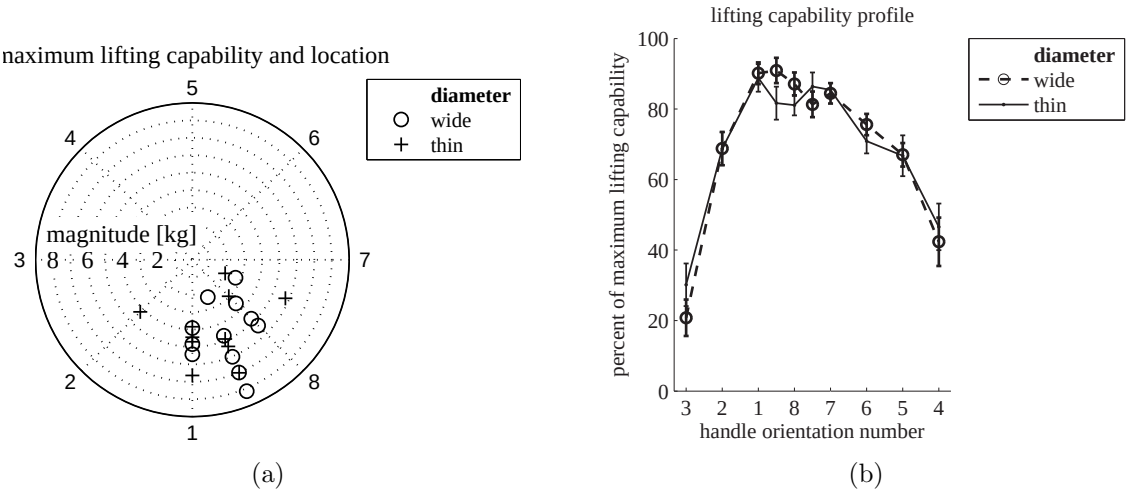

Figure 5.6: (a) The maximum lifting capability measurements are plotted at the corresponding handle orientations for each of the angular precision levels (wide or thin diameter). For each data point, the direction relative to the circle center represents the handle orientation of the measurement taken at the sampled handle angles (Fig. 5.2(c)). The radial distance from the circle center represents the magnitude of the maximum lifting capability. For all but two cases, the maximum measurement occurred between orientations 1 and 7.5, facing toward the participant's right side. These measurements were used to normalize each participant's measurements to determine the lifting capability ratios for the selected lift angles. (b) Mean values for the normalized lifting capability profile over the sampled handle angles.

For the right region, where the object rotation was clockwise, the Pearson correlation coefficient was $r(18) = 0.9678$ ($p < 0.001$). The positive correlation value indicates that a larger amount of rotation before lifting corresponded to a larger change in the lifting capability at the selected lift-off angle relative to the initial handle angle.

For all participants, the minimum lifting capability measure for a set of 8 selected lift-angles was greater than 50% of the individual's maximum lifting capability. The mean of the minimum lifting capability measure varied from 70% to 74% for the 4 combinations of object masses and angular precision levels (Table 5.5). The variability of the lifting capability (Fig. 5.8) measured by the AAD decreased by 1.5% for the increased angular precision ($t(34) = -2.39$, $p = 0.0225$) (Table 5.6).

The capability profiles measured for some participants had multiple local maxima, such that the potential lift-angle region determined only by the minimum lifting capability ratio could be discontinuous (Fig. 5.9). In general, the range of the potential lift-angle region based on minimum lifting capability was larger than the actual lift-angle region for several participants for both angular precision levels (Fig. 5.9). Thus, although the actual lifts occurred at angles associated with lifting capability ratios which were greater than the minimum ratio, some participants still rotated

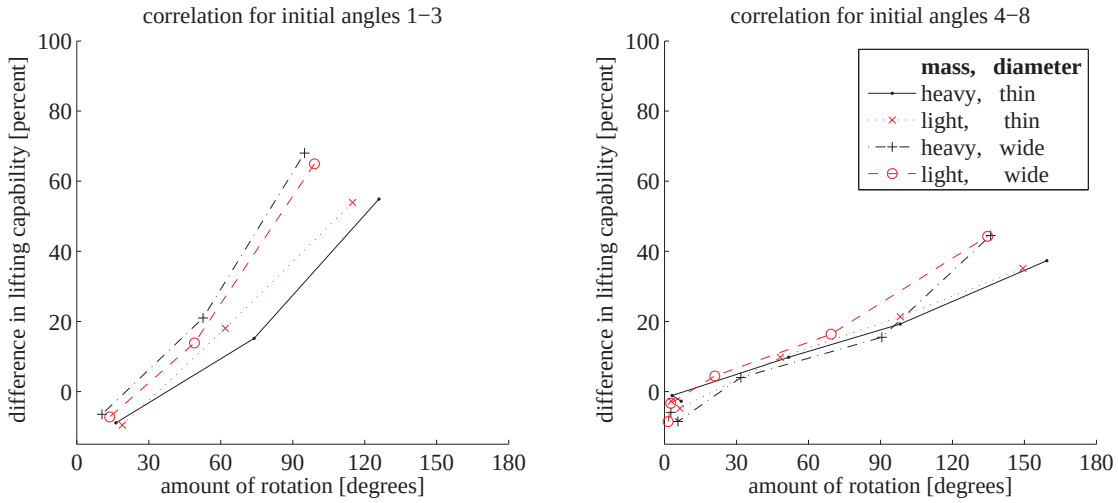

Figure 5.7: Correlation between amount of rotation and the difference in lifting capability. The data are divided due to the asymmetry of the (a) left and (b) right sides of the object orientation scale. The positive correlation indicates that for more object rotation prior to lifting, there is a larger difference in the lifting capability at the selected lift-off posture and the initial lift-off posture.

Table 5.5: Mean and standard error for the minimum capability ratios (percent) associated with the selected lift-angles of the four tested task conditions.

Mean ± Standard Error (percent)	Object mass	
	Light canister	Heavy canister
Angular precision level		
Wide diameter ring	70.3 ± 2.4	71.4 ± 2.6
Thin diameter ring	72.3 ± 3.0	73.7 ± 3.2

the objects more than was necessary to be just within the boundary of the potential lift-angle region.

5.8 Experiment 2: Direct grasping strategy

The purpose of the second experiment was to investigate the alternative strategy of direct grasping, for comparison with the pre-grasp object rotation strategy. In a direct grasping strategy, an object is grasped from its presented configuration without any prior adjustment in the environment.

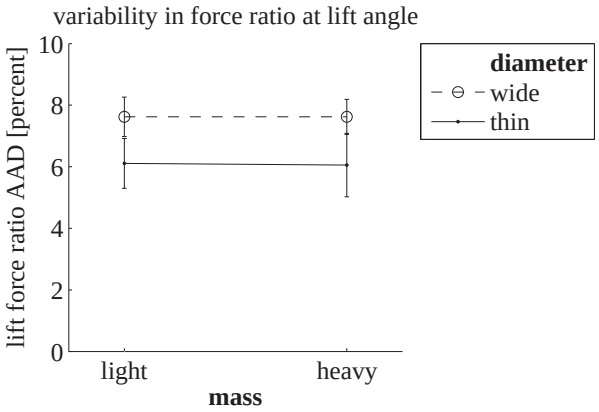

Figure 5.8: Mean data for variability of the lifting capability ratio, as measured by the absolute average deviation (AAD). The error bars indicate ±1 standard error of the mean estimate. The LME model results (Table 5.6) indicate that the variability of the lifting capability ratio decreased 1.5 percent for the increased angular precision corresponding to the thin diameter balance ring.

Table 5.6: Main effects on lifting capability variability from the LME analysis. Lifting capability variability is measured by the average absolute deviation (AAD) of lifting capability ratios (in percent) corresponding to the set of 8 initial handle orientations. The negative values for the additive effects indicate the variability decreases for increased object mass and increased angular precision.

Main effects	Mean (percent)	Std. Error (percent)	$t(DF)$	p
Lifting capability variability			$t(34)$	
Baseline (L=0, W=0)	7.6	0.7		
Object mass (H=1)	0.03	0.6	-0.04	0.9671
Angular precision (T=1)	-1.5	0.6	-2.39	0.0225*

Random effects	Baseline intercept (percent)	Residual (percent)	L ratio	p
Participant	1.5	2.2	5.10	0.024*

First, Experiment 2 tested whether the canister lifting task in Experiment 1 Part I was still achievable without pre-grasp object rotation. That is, was pre-grasp rotation necessary to complete the task successfully, or was it possible to perform the task using direct grasping without rotation? Second, Experiment 2 provides additional insight into the trade-offs in task performance between the two movement strategies.

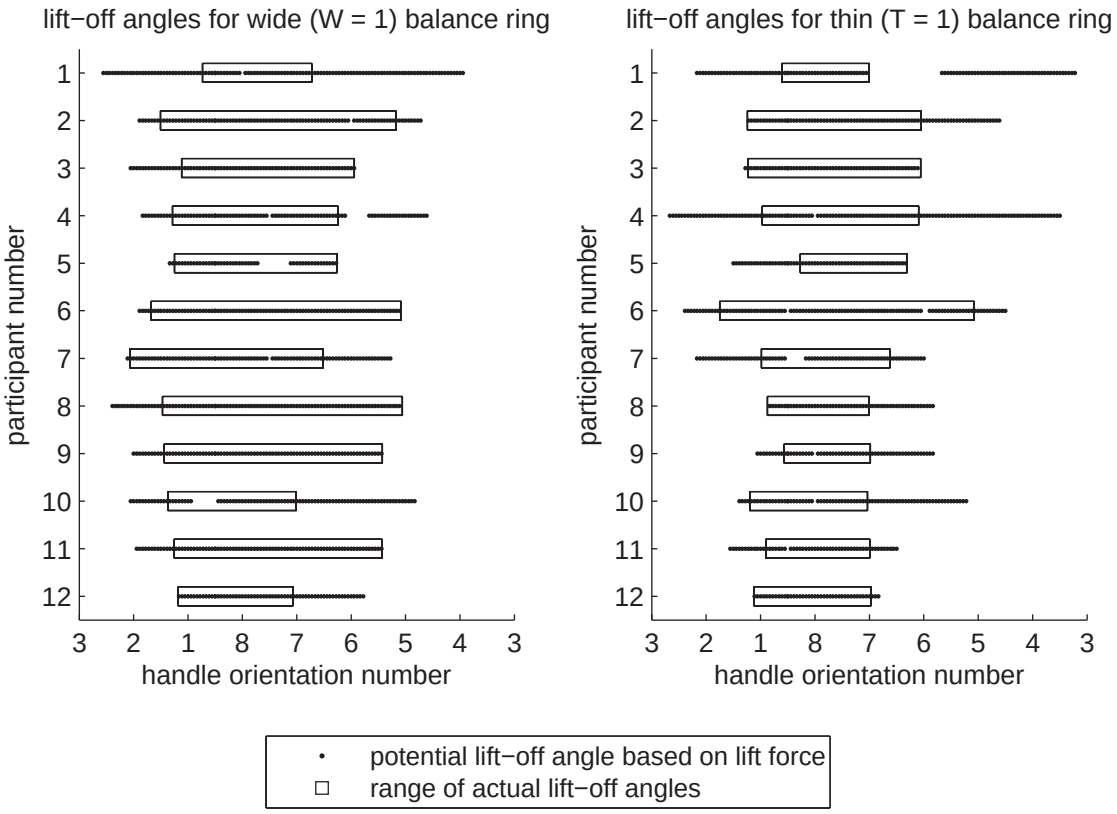

Figure 5.9: Comparison of actual lift-off angles from Part I to the potential lift-off angles based on the lifting capability measures from Part II. The potential lift-angle region computed from the minimum lifting capability is usually larger than the actual lift-angle region, suggesting that the lifting capability alone does not predict the selected lift-off angle. For some participants, there are discontinuities in the capability-based potential lift-angles because the measured capability profiles had multiple local maxima.

Was one strategy more robust for repeated task success? If the task was possible to complete by direct grasping, did it require more time than the spontaneously selected pre-grasp rotation strategy? If so, what time-period of the action would be lengthened?

5.8.1 Apparatus and experimental procedure

The procedure for Experiment 2 was the same as for Experiment 1 Part I, except for an additional task constraint for the lifting trials. The experimental setting and set of canisters remained the same.

As in Experiment 1, there were a total of 32 trials, testing each possible combination of the 2 mass levels, 2 angular precision levels, and 8 initial handle orientations.

The 32 conditions were presented in a randomized sequence, and a new randomized sequence was generated for the remaining conditions if there were any failed trials where the ball dropped. It was expected that it might not be feasible to lift the canister from all of the presented handle configurations. Thus, a condition was not repeated in the sequence once there were 3 failures for that specific condition.

The participant was instructed to retrieve the token under the canister as quickly as possible, without disturbing the balanced ball and without sliding the canister on the platform surface prior to lifting. The participant was informed that there would be a maximum of 3 attempts for any one set of task conditions.

In all, 5 participants (4 male, 1 female) who participated in Experiment 1 volunteered for Experiment 2 as a follow-up study.

5.8.2 Data analysis

Because the task timing would be compared between the constrained condition in Experiment 2 and the unconstrained condition in Experiment 1, we first verified that the 5 follow-up participants were a representative subset of the 12 original participants with respect to the task completion time for Experiment 1. An LME model tested the total task completion time for the pre-grasp rotation trials in Experiment 1 for the $N = 5$ subset of data.

The number of unsuccessful trials was compared between the two movement strategies for the 5 participants in Experiment 2. The total number of unsuccessful attempts included trials of any repeated task conditions. It was expected that there would be fewer unsuccessful attempts with the unconstrained pre-grasp rotation strategy in Experiment 1, because direct-reach-to-grasp actions would be less robust to the variety of task conditions.

The task timing for the direct grasping trials in Experiment 2 was determined using the same time segmentation procedure from Experiment 1 (see §5.6.1). The time courses of the natural pre-grasp rotation strategy and the direct grasping strategy were compared using LME models. The input data were the time intervals for the trials from Experiment 1 Part I and Experiment 2. Only the data from the 5 follow-up participants were considered, so that the strategies were compared trial-by-trial for corresponding task conditions and participants.

We hypothesized that the unconstrained pre-grasp rotation strategy of Experiment 1 could result in less total time for task completion relative to Experiment 2 due to the advantage of lifting the canister from a preferred object orientation. It was expected that the additional time required for the optional rotation action would be offset by decreased time for the pre-interaction and lifting action intervals. The LME model tested the fixed effect of the strategy on the time interval duration in a model which also included a random effect for the participant and fixed effects for the object mass, angular precision, the initial handle orientation, and the square of the initial handle orientation.

In addition, separate LME models tested the fixed effects on timing intervals for each strategy. In this case, the model for pre-grasp rotation tested the entire set of data from the 12 participants in Experiment 1. Another model for direct grasping tested the set of data from the 5 participants in Experiment 2.

5.9 Experiment 2 results

5.9.1 Increased task completion failures for direct grasping

No participant reported fatigue or requested immediate rest in addition to the scheduled breaks.

Three participants failed to complete the token retrieval task for 6-8 trials (including repeated task condition attempts), one participant failed for 11 trials, and one participant failed for 22 trials. For all five participants, there were more failed attempts with the direct grasping strategy in Experiment 2 than with the natural pre-grasp rotation strategy in Experiment 1 (Fig. 5.10(a)).

The token retrieval task was feasible for all five participants if it involved the lower angular precision requirement corresponding to the wide diameter balance ring (Fig. 5.10(b)). Failures to complete the task within 3 attempts were otherwise observed for all participants (two failing at one condition, two at two conditions, and one at five conditions). One participant was not able to directly lift the light canister with the thin diameter balance ring at handle orientation 4. Four of the five participants were not able to directly lift the heavy canister with the thin ring at handle orientation 4. The other infeasible tasks occurred at handle orientations 3, 5, and 7 for the heavy canister with the thin ring.

5.9.2 Direct grasping changed pre-interaction timing

The first LME model tested whether the 5 follow-up participants were representative of the original 12 participants. It was based on the total task time for the pre-grasp rotation trials in Experiment 1. The fixed effects which were significant for the complete set (Table 5.2) were also significant for the subset of 5 follow-up participants. The average total time was 3.96 seconds, which was higher than the average for all 12 participants (3.46 seconds).

The next LME model compared the timing between the two strategies for the successfully completed trials (Table 5.7). The total time for task completion T_0 was on average 1.49 seconds shorter using direct grasping instead of pre-grasp rotation ($t(298) = -6.00$, $p < 0.001$). The pre-interaction time T_1 of perception and reaching before initial object movement was 0.75 seconds longer for direct grasping over pre-grasp rotation ($t(298) = 5.41$, $p < 0.001$). The second time interval T_2 for object interaction between initial object movement and object lift-off was 2.30 seconds shorter for direct grasping without pre-grasp rotation ($t(298) = -11.43$, $p < 0.001$).

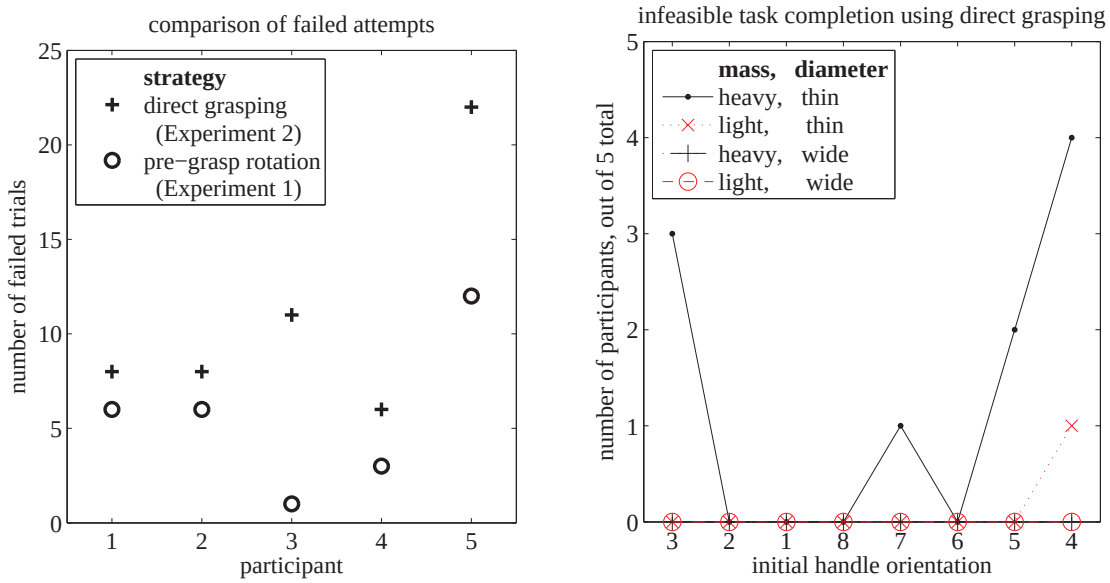

Figure 5.10: Failures at task completion using the direct grasping strategy. (a) Total number of failed lifting attempts (including repeated task conditions) for the token retrieval task. All five participants failed more trials when restricted to a direct grasping strategy without pre-grasp rotation. (b) The direct grasping task was considered infeasible for a particular set of task conditions if the participant could not complete the task within three attempts. All infeasible tasks conditions involved the increased angular precision level corresponding to the thin diameter balance ring. The task which was infeasible for the most participants was the token retrieval for the heavy canister with thin diameter balance ring at handle orientation 4, which is opposite the lift-off angle region observed when the natural pre-grasp manipulation is allowed.

The difference in the object lifting time T_3 between the two strategies was not significant.

For the direct grasping trials in Experiment 2, the total time T_0 for task completion was 3.51 seconds on average for the baseline light object with wide diameter ring at initial orientation 8 (Table 5.8). Increasing the angular precision increased T_0 by 2.58 seconds ($t(140) = 7.92$, $p < 0.001$). Similar to the Experiment 1 trials, the task took longer for initial handle orientations further from the central orientation 8 (Fig. 5.11), as indicated by the positive coefficient for the squared orientation term ($t(140) = 9.88$, $p < 0.001$).

The main difference in timing between the two strategies (Fig. 5.12) is the quadratic trend with respect to initial handle orientation (Table 5.9). For the natural pre-grasp rotation strategy in Experiment 1, the object interaction time T_2 had the largest coefficient (0.33) for the squared orientation effect ($t(366) = 21.35$, $p < 0.001$) . The squared orientation coefficients were 0.04 ($t(366) = 4.17$, $p < 0.001$)

Table 5.7: Fixed effect of manipulation strategy on movement time (seconds) from the LME analysis. Significant effects ($p < 0.05$) are denoted by asterisks (*). In the model, the baseline strategy was the natural pre-grasp rotation strategy observed in Experiment 1. Thus the mean value for the strategy factor is the additive effect for the direct rotation strategy observed in Experiment 2.

Main effect of direct grasping strategy	Mean (s)	Std. Error (s)	$t(298)$	p
T_0: total task completion time	-1.49	0.25	-6.00	<0.001*
T_1: perception and reaching	0.75	0.14	5.41	<0.001*
T_2: manipulation and grasping	-2.30	0.20	-11.43	<0.001*
T_3: object lift and token retrieval	0.06	0.11	0.54	0.5863

Table 5.8: Effects on task completion time (seconds) from the LME analysis of the direct grasping trials in Experiment 2. Significant effects ($p < 0.05$) are denoted by asterisks (*). Expect for object mass, the fixed effects which are significant are the same as those for the pre-grasp rotation trials (Table 5.2).

Main effects	Mean (s)	Std. Error (s)	t(DF)	p
T_0: total task completion time			$t(140)$	
Baseline (L=0, W=0)	3.51	0.92		
Object mass (H=1)	0.44	0.32	1.35	0.1807
Angular precision (T=1)	2.58	0.33	7.92	<0.001*
Orientation^2 coefficient	0.36	0.04	9.88	<0.001*
Orientation coefficient	-0.64	0.08	-7.94	<0.001*

Random effects	Baseline intercept (s)	Residual (s)	L ratio	p
Participant	1.92	1.96	70.44	<0.001*

for the pre-interaction time T_1 and -0.02 ($t(366) = -2.13$, $p = 0.0335$) for the object lifting time T_3. However, for the direct grasping strategy in Experiment 2, the pre-interaction time T_1 had the largest coefficient (0.22) for the squared orientation effect ($t(140) = 6.32$, $p < 0.001$). The squared orientation coefficients were 0.05 ($t(140) = 5.06$, $p < 0.001$) for the object interaction time T_2 and 0.09 ($t(140) = 5.02$, $p < 0.001$) for the object lifting time T_3. In other words, the response to initial handle orientation occurred primarily in the object interaction time for the pre-grasp rotation strategy but occurred primarily in the pre-interaction time for the direct

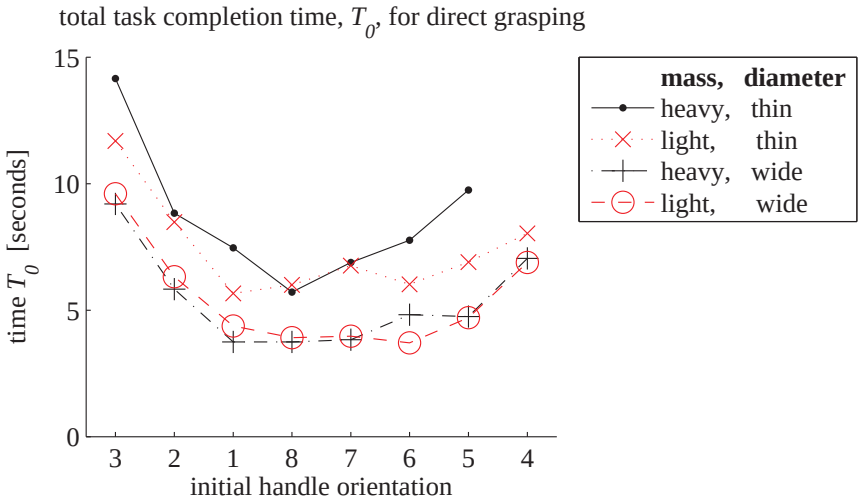

Figure 5.11: Mean data for the total time T_0 for task completion when constrained to direct grasping. The point at initial orientation 4 for the heavy canister with the thin ring is not shown because only 1 of the 5 participants completed this task condition. The LME model results (Table 5.8) indicate that task completion time was longer for increased angular precision corresponding to the thin diameter balance ring. The task also took longer if the handle orientation started farther from the central orientation 8, as indicated by the quadratic trend with respect to the initial handle orientation.

grasping strategy.

5.10 Discussion and summary

Pre-grasp object rotation, as opposed to direct-reach-to-grasp, was investigated as a movement strategy. Of particular concern were the conditions that lead to pre-grasp rotation and its utility. Selecting pre-grasp rotation over the alternative of direct grasping requires a choice in the amount of object rotation, which determines the object angle at the lift-off time. The experimental results suggest that the lifting capability associated with a particular object orientation is one criterion which underlies the selected movement strategy. The choice of lift-off angles, in turn, is sensitive to difficulty factors of object mass and task precision. In almost all of the task conditions, pre-grasp rotation is optional, because it is possible to use direct grasping even for object lift-off angles associated with lower lifting capability. One potential advantage of selecting a posture with higher capability is the improved robustness of completing the task successfully, at the cost of longer time for task completion.

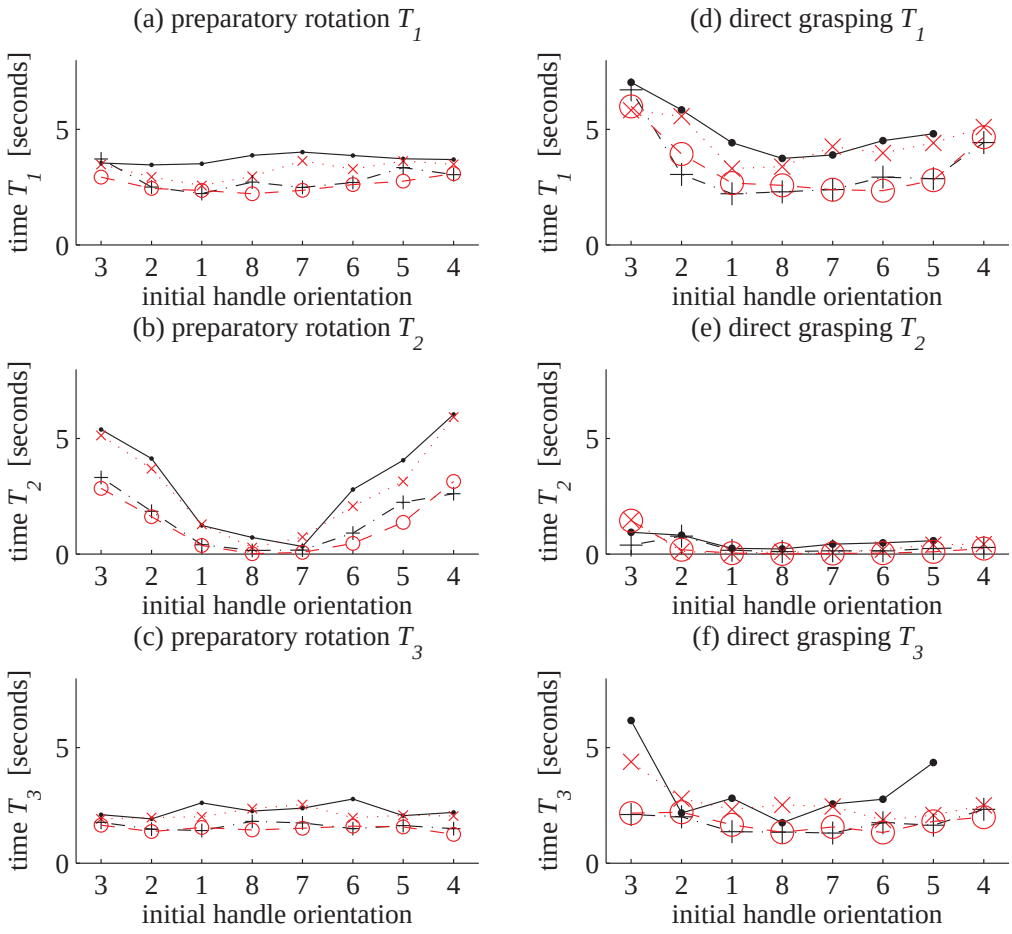

Figure 5.12: Mean data for the task components for both (a, b, c) the natural strategy observed for 12 participants in Experiment 1 and (d, e, f) the direct grasping strategy observed for 5 follow-up participants in Experiment 2. For the direct grasping condition, the point at initial orientation 4 for the heavy canister with the thin ring is not shown because only 1 of the 5 participants completed this task condition. Table 5.9 presents LME results of the fixed effects for each time interval.

5.10.1 Selection criteria for object rotation and lift angle

This study replicated previous observations that pre-grasp rotation shows a quadratic trend with respect to the initial handle orientation of the lifted object [Chang *et al.*, 2008]. That is, people tend to rotate the object in the shorter direction toward a desired lift-off orientation. More important, this investigation extends beyond the previous work by examining posture-dependent lifting capability as a criteria for the lift-off angle selection.

The results of Experiment 1 showed that the maximum lifting capability occurred around the central object orientation number 8, in the same region of initial handle

Table 5.9: Fixed effects on action component times (seconds) resulting from the LME analysis. Significant effects ($p < 0.05$) are denoted by asterisks (*). In the model, the baseline mass level was light (L) and the baseline diameter for angular precision was wide (W). Thus the mean value for the mass and angular precision are the additive effects for the heavy (H) level and the thin (T) ring diameter.

Main effects	Natural strategy with pre-grasp rotation (Experiment 1, $N = 12$)				Direct strategy without pre-grasp rotation (Experiment 2, $N = 5$)			
	Mean (s)	Std. Err. (s)	t(DF)	p	Mean (s)	Std. Err. (s)	t(DF)	p
T_1: pre-interaction			$t(366)$				$t(140)$	
Baseline (L=0, W=0)	2.35	0.21	11.36	<0.001*	2.26	0.64	3.55	<0.001*
Object mass (H=1)	0.35	0.08	4.51	<0.001*	0.21	0.07	0.97	0.3314
Task precision (T=1)	0.76	0.08	9.76	<0.001*	1.37	0.22	6.32	<0.001*
Orientation^2 coeff.	0.04	0.01	4.17	<0.001*	0.22	0.02	9.06	<0.001*
Orientation coeff.	0.01	0.02	0.48	0.6344	-0.37	0.05	-6.85	<0.001*
T_2: object interaction			$t(366)$					
Baseline (L=0, W=0)	-0.45	0.36	-1.26	0.2072	-0.01	0.10	-0.07	0.9472
Object mass (H=1)	0.25	0.14	1.73	0.0837	0.10	0.09	1.06	0.2930
Task precision (T=1)	1.57	0.14	11.00	<0.001*	0.17	0.10	1.87	0.0640
Orientation^2 coeff.	0.33	0.02	21.35	<0.001*	0.05	0.01	5.06	<0.001*
Orientation coeff.	-0.29	0.03	-8.23	<0.001*	-0.12	0.02	-5.23	<0.001*
T_3: object lifting			$t(365)$					
Baseline (L=0, W=0)	1.55	0.19	8.26	<0.001*	1.26	0.35	3.59	<0.001*
Object mass (H=1)	0.14	0.07	2.18	0.0296*	0.13	0.16	0.82	0.4142
Task precision (T=1)	0.64	0.07	9.81	<0.001*	1.04	0.16	6.56	<0.001*
Orientation^2 coeff.	-0.02	0.01	-2.13	0.0335*	0.09	0.02	5.02	<0.001*
Orientation coeff.	0.02	0.02	1.22	0.2214	-0.15	0.04	-3.94	<0.001*

angles where participants performed the least amount of rotation. The lifting capability might be considered a biomechanical measure of the comfort associated with a lifting posture, especially for a somewhat demanding task as the token retrieval task presented here. In accordance, the set of selected lift-angles with high lifting capability could be described as a "comfort zone" for the lifting grasps.

In particular, there was a strong correlation between the amount of object rotation and the change in the lifting capability that rotation would produce. This correlation suggests that the sub-action of pre-grasp rotation acts to increase the

posture-dependent capability for the primary lifting task. The object reconfiguration, which anticipates the strength capability for the final action, may be an externalization of the end-state comfort effect first observed by Rosenbaum *et al.* [1990]. Instead of only choosing a body posture in response to a task condition, the participant also adjusts the object in the environment to obtain task conditions which afford preferable postures.

The correlation curves (Fig. 5.7) contained points where there was a small negative difference in the lifting capability for a small amount of rotation. These corresponded to trials for initial handle orientations 1, 7, and 8, where the lifting capability is already near the maximum. The change in handle orientation may be due to the local adjustment resulting from the grasping interaction at lift-off time. The small decrease in lifting capability may be acceptable because the handle angle remains within the preferred comfort zone and does not need to be at the optimum posture. The values of negative change in lifting capability could also be due to limitations in the measurement technique of lifting capability, which depended on participants to perform the lifting interaction with a consistent amount of maximum effort for all samples.

The measurement of individual lifting capability over multiple object orientation angles was also used to investigate manual lifting strength as a possible predictor of the selected object lift-off angles. For all participants the lift-off angles in Experiment 1 Part I corresponded to at least 50% of the maximum lifting capability measured in Experiment 1 Part II. However, it is not clear that the lifting capability alone determines the angle selection. Although the minimum capability ratio is achievable for the entire potential lift-angle region, the actual lift-angle region may be more narrow due to additional factors such as preferred arm posture for lifting or the grasp comfort.

The selection criteria for the lift-off angle may be a combination of both costs associated with goal postures at the lift-off states and costs associated with the rotation sub-action. For example, the variability of lift-off angles over different initial angles may be due to a trade-off between the pushing effort required for the amount of rotation and the lifting capability at the final selected lift-off angle. Future investigation is required to determine whether goal postures alone are sufficient for predicting the selection of preparatory movements, as in a posture-based planning model [Rosenbaum *et al.*, 1993a,b, 1995]. It is possible that for complex manipulation involving pre-grasp object adjustment, action costs are required to model the dynamic object interaction in addition to the goal grasping postures.

5.10.2 Sensitivity to task difficulty factors

Experiment 1 also investigated how the task difficulty affects the specific strategy of pre-grasp object manipulation. Our intent was to examine whether the lift-off angle selected by the pre-grasp rotation was invariant across task conditions, which

might be the case if individuals had a fundamental preference for a particular set of geometric task conditions. We expected, however, that the selection of object lift-off angles would vary in response to the task demands, such that more demanding tasks would result in a narrower range of angles which resulted in successful task completion.

Object mass and required angular precision, as manipulated by the balance ring diameter, were tested as two specific task difficulty factors which could affect task performance in Experiment 1 Part I. The angular precision is similar to the geometric task difficulty index proposed by Latash and Jaric [2002] as a parameter for describing postural coordination for a drinking task. We additionally investigated object mass as a strength requirement for the manual lifting task in our experiments, which may be more physically demanding than a drinking task. The differences in total time of task completion confirmed that it was more difficult to lift the heavier object and/or balance with the thinner diameter ring. Both factors contributed to longer time for the token retrieval task, even for the same initial handle orientation.

We first examined the effect of the task difficulty factors on the amount of object rotation. Only angular precision, and not object mass, was a significant factor for the amount of rotation. When greater angular precision was required for the lifting task, the canister was rotated more from the initial orientation.

We also examined the effect of the two task difficulty factors on the variability of the object orientation at the lift-off time. Increased object mass and increased angular precision resulted in decreased lift-angle variability. This relationship indicates that although object mass did not affect the amount of object rotation for a specific starting angle, mass did affect the selection of the preferred angles prior to lifting. The variability of the lift-angles represents the size of the comfort zone for a particular combination of object mass and task precision. Increased variability indicates a less restrictive selection of the preferred lifting postures. The lower levels of object mass and/or task precision resulted in a relaxation of the constraint on the region of selected lift-angles. For the task precision factor (but not the object mass factor), the increased variability in the lift-angles also corresponded to increased variability in the lifting capability within the selected comfort zone. This result suggests that the externalization of the lifting capability criterion is diminished for less precise tasks, which adds to the previous finding by Rosenbaum *et al.* [1996] that the end-state comfort effect is decreased for less precise dowel wielding tasks.

The difference in the effect of object mass on the specific amount of rotation and the lift-angle variability could be due to the specific two levels of object mass tested in the experiment. The canister mass for the "heavy" level was specifically chosen to be light enough to avoid fatigue from repeated lifting. Given the trend in the predicted direction, we hypothesize that the object mass would reach significance as a factor for the amount of object rotation if a heavier canister could be reasonably tested in a future experiment.

5.10.3 Pre-grasp rotation as a movement strategy

The strategy of pre-grasp rotation is an example of movement planning for a sequence of sub-actions to complete a task. In particular, pre-grasp rotation involves not only the anticipation of body posture but also the selection of the intermediate object configuration that is an adjustable task condition. If the intermediate object configuration is similar for different task conditions, then a well-practiced or comfortable lifting action can be reused. Achieving similar intermediate task conditions for reuse of the lifting action requires adaptation of the rotation action for different initial object orientations. The rotation of the object handle to a particular region suggests that for some tasks there may be a preference for adapting only the pre-grasp action component for reuse of the lifting posture, compared to planning an entirely new direct grasping action. The present findings indicate that the choice to adapt the rotation motion instead of the lifting action may be influenced by the posture-dependent lifting capability as well as the difficulty factors of object mass and task precision.

In some cases, the task conditions may be sufficiently difficult that pre-grasp rotation may be required to complete the task because direct grasping is infeasible. The results of Experiment 2 suggest that this case occurs for the lifting a heavy object with a strict angular precision requirement when the handle is opposite the preferred lift-angle region. In less demanding tasks involving a lighter object or less angular precision, direct grasping may be a feasible strategy for successful completion. The participants' choice in Experiment 1 to use pre-grasp rotation even for these task conditions indicate there are advantages to reconfiguring the object to a new orientation with higher lifting capability. The results from Experiment 2 suggest that lifting from the preferred comfort zone improves the robustness of the lifting action such that the task can be completed successfully in fewer attempts. The increase in robustness may be due to the increased efficiency or safety margin of the higher lifting capability relative to the minimal lifting capability required to complete the task.

Pre-grasp rotation did not result in a clear timing advantage over direct grasping, because the total time was shorter for direct grasping. The largest difference between the two strategies is the relative contribution of the different time intervals for pre-interaction, object interaction, and object lifting. The time T_3 for object lifting was not significantly different between the two strategies. The results from Experiment 2 suggest pre-grasp rotation may be selected over direct grasping due to the shorter planning and reaching time T_1 before initial object movement, even if the increase in the object interaction time T_2 results in an overall longer time for task completion.

We further note differences in how the component times were affected by the object orientation. For the unconstrained pre-grasp rotation strategy, the pre-interaction time was relatively constant across different initial handle orientations (Fig. 5.12). That is because the goal posture for initial object movement must only contact the

canister handle for the adjustment motion. The main contribution to the quadratic effect of initial orientation is from the object interaction time T_2 where the pre-grasp rotation occurs.

However, when pre-grasp rotation is not permitted in the direct grasping trials, the main contribution to the quadratic orientation trend arises in the pre-interaction time, T_1. The increase in the pre-interaction time for direct grasping suggests that the restriction for no pre-grasp adjustment is compensated for by longer planning and reaching times before object contact. The planned posture for initial contact with the object must be capable of lifting the canister to complete the token retrieval task. This difference in the pre-interaction time may be partially due to a mental rotation of the object for posture planning, to compensate for the lack of physical rotation. The average increase of 0.75 s for time T_1 is on the order of the delay times for mental rotation (0.4-1.1 s) reported by Cooper and Shepard [1973], for example. Although the pre-grasp rotation strategy requires longer absolute time for task completion, the longer pre-interaction time for direct grasping may affect the participants' perceived task time due to the complexity of initial planning phase.

Overall, pre-grasp object rotation involves several components which may be studied to understand how humans choose between alternative movement strategies. We investigated the changes in object configuration prior to lifting, which represented the selection of the preferred task-space conditions for the manual lifting action. Further investigation of posture and grasp differences may lead to greater insight into the factors determining an individual's choice between a direct grasping strategy and a pre-grasp rotation strategy. Factors such as initial and final object locations may also play a role in the selection of a manipulation strategy. In addition, we observe the pre-grasp rotation is just one of many strategies that people use for object interaction prior to grasping. Other pre-grasp manipulation actions include sliding, pivoting, and complex tumbling maneuvers to reconfigure the the object before the final grasp. Additional research of pre-grasp manipulation strategies beyond pre-grasp rotation can complement the existing literature which has often focused on direct reaching and grasping tasks.

6 Workspace analysis

For some scenarios, it may be possible to complete the manipulation task using a direct grasping strategy without any preparatory object rotation. In these cases the pre-grasp rotation strategy may provide some benefit over the direct grasping strategy even though it is not required. This benefit was examined in the human motion studies presented in §4 and §5. In the human studies, we also found that the hand grasp on the object changed when using the direct grasping strategy. It is possible that the change in grasp was to accommodate the increased difficulty of reaching the object handle in the orientations facing away from the participant.

In other cases, the task may be more demanding such that grasping is infeasible without pre-grasp object rotation. For example, the object start position may be near the boundary of the arm's workspace such that some handle angles are unreachable by direct grasping. The object may be unreachable especially when a specific grasp of the object handle is required.

We investigated the workspace for reachable grasps of an example pan grasping task. This chapter presents kinematic analyses for both a model of the human arm and a model of a robotic manipulator system. This workspace analysis was used to compute the set of pan configurations which are directly graspable by the manipulator in the example object acquisition task.

6.1 Example grasping task scenario

The example task scenario is to grasp a cooking pan from a presented position on a tabletop surface. The goal of the workspace analysis is to determine the initial pan configurations on the tabletop where it can be grasped. The boundary of the workspace will indicate potential areas where a preparatory manipulation strategy such as object rotation can extend the applicability of the grasp.

In the kinematic analysis, the manipulator's base is fixed at the origin, and the $+z$ axis points upwards such that the table surface spans the x-y plane (Fig. 6.1). We consider the workspace in front of the manipulator ranging from a height of 0.5m to 1.5m. Results are shown in particular for a selected table height of 0.765m, based on the dimensions of a physical table used in the laboratory for demonstrations.

The cooking pan is a typical example of a handled object for which pre-grasp

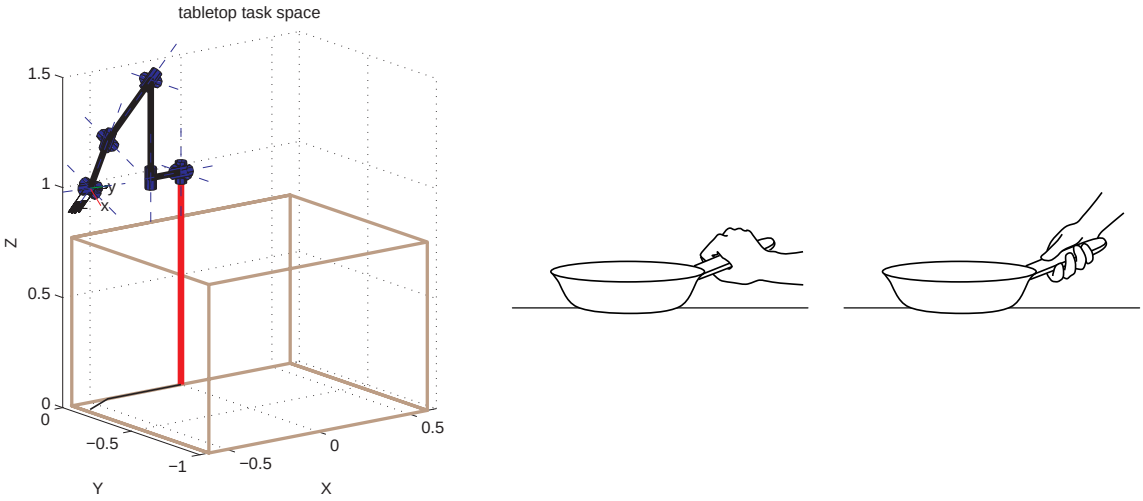

(a) (b) (c)

Figure 6.1: Example pan grasping task for workspace analysis. (a) Tabletop task space for example scenario. (b) Overhand, straight cylindrical grasp of cooking pan by its handle. (c) Underhand, oblique cylindrical grasp of cooking pan by its handle.

rotation may facilitate grasping. The modeled object dimensions are based on a small cooking pan with 17cm-diameter body and 12cm-length handle. The specific grasps considered in the analysis are an overhand, straight cylindrical grasp where the thumb wraps under the handle and an underhand, oblique cylindrical grasp, where the fingers wrap underneath the handle with the thumb above (Fig. 6.1). These grasps are based on the two main styles of grasps observed for many of the participants in the studies of natural pre-grasp rotation in §4. The following kinematic analysis focuses on achieving the relative transformation between the pan handle and the hand dorsum necessary for the oblique grasp, and it is not concerned with the hand pose as determined by the finger configuration.

6.2 Human manipulator model

The kinematic model of the human right arm consists of 10 degrees-of-freedom (DoFs) representing the kinematics of the trunk, shoulder, elbow, and wrist joints. The joint axes are described relative a reference configuration where arm is extended to the right of the body and the palm is face-up (Fig. 6.2). Starting proximally from the trunk and proceeding distally to the wrist, the joint axes are:

1. trunk rotation (torsion) – axial rotation about the spine,
2. trunk lateral bending (lateral flexion) – bending at the waist to the left and right sides,

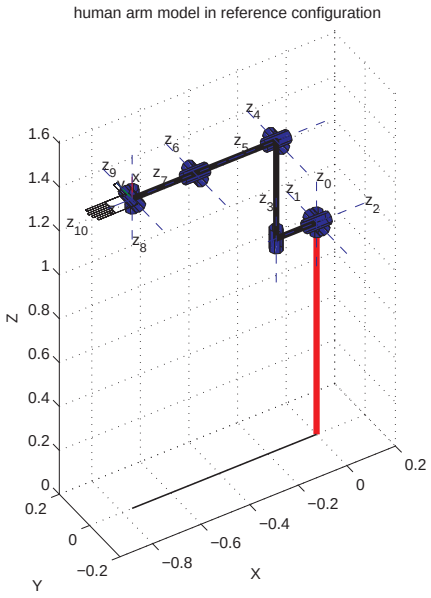

Figure 6.2: Kinematic model of the human right arm in the reference configuration. In the reference configuration, the palm of the hand faces up ($+z$) and the thumb points behind the body ($+y$).

3. trunk flexion – bending at the waist to the front and back,
4. shoulder azimuth (adduction/abduction) – rotation of the upper arm in the horizontal plane,
5. shoulder elevation (flexion/extension) – rotation of the upper arm in the vertical plane,
6. shoulder axial rotation (internal/external rotation) – axial rotation about the upper arm,
7. elbow flexion – bending between the upper arm and forearm,
8. forearm axial rotation (pronation/supination) – axial rotation of the forearm segment,
9. wrist deviation (adduction/abduction) – bending of the hand to the thumb (radial) and pinky (ulnar) sides, and
10. wrist flexion – bending toward the back (dorsal) or palm (volar) sides of the hand.

The human manipulator is modeled as a serial chain of single-axis joints using the Denavit-Hartenberg (DH) convention for describing kinematic chains [see, e.g., Spong *et al.*, 2006]. The DH parameters for the human arm model (Table 6.1) are expressed in "standard" DH notation, as opposed to the "modified" DH notation. First, the base link of the model is located at height d_0 above the origin to model

Table 6.1: Kinematic model of the human right arm, expressed in standard Denavit-Hartenberg notation. All 10 joints are revolute DoFs. The base is offset vertically by a height d_0 to model the person in standing position. The non-zero joint offsets correspond to average limb lengths for the human arm (see text).

axis	joint description	joint angle θ [rad]	link length a [m]	link twist α [rad]	joint offset $d[m]$
-	lower body length	-	-	-	$d0 = 0.966$
z_0	trunk rotation	θ_1	0	$-\frac{\pi}{2}$	0
z_1	trunk lateral bending	θ_2	0	$-\frac{\pi}{2}$	0
z_2	trunk flexion	θ_3	0	$+\frac{\pi}{2}$	$d_3 = -0.166$
z_3	shoulder azimuth	θ_4	0	$+\frac{\pi}{2}$	$d_4 = 0.445$
z_4	shoulder elevation	θ_5	0	$+\frac{\pi}{2}$	0
z_5	shoulder axial rotation	θ_6	0	$-\frac{\pi}{2}$	$d_6 = 0.334$
z_6	elbow flexion	θ_7	0	$+\frac{\pi}{2}$	0
z_7	forearm axial rotation	θ_8	0	$-\frac{\pi}{2}$	$d_8 = 0.259$
z_8	wrist deviation	θ_9	0	$-\frac{\pi}{2}$	0
z_9	wrist flexion/extension	θ_{10}	0	$+\frac{\pi}{2}$	0

the human in a standing position. The non-zero joint offset parameters d_i in Table 6.1 correspond to the body segment dimensions of the trunk and arm. The lengths are based on the average anthropomorphic data reported by Clauser *et al.* [1969] as follows:

- d_0 is the lower body length (to represent the offset from the ground to the trunk),
- d_3 is half the chest breath length (to represent the offset from the trunk center to the right shoulder),
- d_4 is the trunk height,
- d_6 is the upper arm length, and
- d_4 is the forearm length.

For convenient description of the reference configuration (Fig. 6.2), constant offsets are added to the joint angle values, as listed in Table 6.2. With these offsets, all joint angles are zero in the reference configuration. In addition, the model limits the joint rotation to anatomic ranges of motion based on the measurements reported by Doriot and Wang [2006] for the trunk and upper body. The joint range values in Table 6.2 are the average values of maximum voluntary range of motion for male and female adults.

Table 6.2: Joint angle parameters for the model of the human right arm. The reference angle offset is added to the joint angle value for convenience in describing the reference configuration. Joint limits are based on human arm measurements reported in the literature. The forward kinematics pre-computation samples the joint space at discretized points. In most cases, the range of motion is sampled in increments of 15 degrees, denoted by the span notation of θ_{low}:15:θ_{high}.

joint axis description	reference offset [deg]	joint limits [deg]	discretized sample points [deg]
trunk rotation	$\theta_{1r} = 0$	[-64, 64]	{-30:15:30}
trunk lateral bending	$\theta_{2r} =$ -90	[-39, 36]	{-39, -30:15:30, 36}
trunk flexion	$\theta_{3r} = 90$	[-17, 61]	{-15:15:60}
shoulder azimuth	$\theta_{4r} =$ -90	[-15, 136]	{-15:15:135}
shoulder elevation	$\theta_{5r} = 90$	[-76, 60]	{-75:15:60}
shoulder axial rotation	$\theta_{6r} = 0$	[-118, 24]	{-118, -105:15:15, 24}
elbow flexion	$\theta_{7r} = 0$	[21, 143]	{21, 30:15:135, 143 }
forearm axial rotation	$\theta_{8r} = 90$	[-145, -30]	{-145:5:-30}
wrist deviation	$\theta_{9r} =$ -90	[-25, 21]	{-25:5:20}
wrist flexion/extension	$\theta_{10r} = 90$	[-63, 46]	{-60:5:45}

6.3 Robot manipulator model

The robot manipulator system consists of a Mitsubishi PA-10 7-DoF arm with a 24-DoF Shadow Hand C3 end effector (Shadow Robot Company, London, UK). The kinematic model for this workspace analysis only considers the degrees of freedom required to the position the hand segment without modeling the finger joints. The two wrist DoFs of the Shadow Hand are added to the 7 DoFs of the PA-10 arm for a total of 9 DoFs (Fig. 6.3). The robot manipulator's joints correspond approximately to anthropomorphic joint axes, with one less trunk DoF compared to the human model (Table 6.3). In the reference configuration, all the limbs are extended in the vertical direction with fingers pointing up ($+x$) and the thumb pointing to the back ($+z$) (Fig. 6.3).

As with the human model, the serial-chain kinematics of the robot manipulator are expressed in standard DH notation (Table 6.3). The base link of the model is located at height d_0 above the origin to represent the vertical offset of the platform supporting the manipulator. All other non-zero joint offset values represent the part lengths.

Constant offsets are added to the joint angle values Table 6.4 to achieve consistency with the robot programming interface. With these offsets, the reference config-

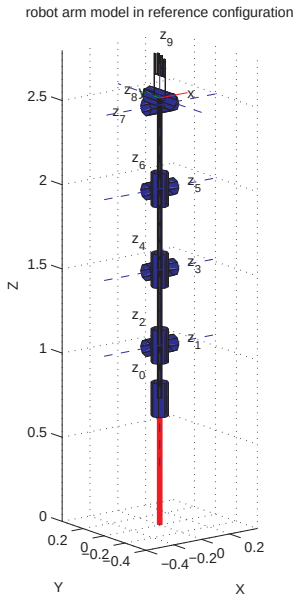

Figure 6.3: Kinematic model of the robot arm in the reference configuration. In the reference configuration, the fingers point up ($+z$) and the thumb points behind the body ($+y$).

Table 6.3: Kinematic model of the robot manipulator, expressed in standard Denavit-Hartenberg notation. All 9 joints are revolute DoFs. The base is offset vertically by a height d_0 to model the person in standing position. The non-zero link lengths and joint offsets correspond to the part lengths of the mechanism.

axis	joint description	joint angle θ [rad]	link length a [m]	link twist α [rad]	joint offset $d[m]$
-	base offset height	-	-	-	$d0 = 0.75$
z_0	trunk rotation	θ_1	0	$-\frac{\pi}{2}$	$d_1 = 0.317$
z_1	trunk flexion	θ_2	0	$+\frac{\pi}{2}$	0
z_2	shoulder azimuth	θ_3	0	$-\frac{\pi}{2}$	$d_3 = 0.45$
z_3	shoulder elevation	θ_4	0	$+\frac{\pi}{2}$	0
z_4	shoulder axial rotation	θ_5	0	$-\frac{\pi}{2}$	$d_5 = 0.48$
z_5	elbow flexion	θ_6	0	$+\frac{\pi}{2}$	0
z_6	forearm axial rotation	θ_7	0	$-\frac{\pi}{2}$	$d_7 = 0.4995$
z_7	wrist deviation	θ_8	$a_8 = 0.04$	$-\frac{\pi}{2}$	0
z_8	wrist flexion/extension	θ_9	0	$+\frac{\pi}{2}$	0

Table 6.4: Joint angle parameters for the model of the robot manipulator. The reference angle offset is added to the joint angle value for convenience in describing the reference configuration. Joint limits are based on the technical specifications of the hardware. The forward kinematics pre-computation samples the joint space at discretized points. In most cases, the range of motion is sampled in increments of 15 degrees, denoted by the span notation of θ_{low}:15:θ_{high}.

joint axis description	reference offset [deg]	joint limits [deg]	discretized sample points [deg]
trunk rotation	$\theta_{1r} = 45$	[-177, 0]	{-165:15:-105}
trunk lateral bending	$\theta_{2r} = 0$	[-94, 94]	{-90:15:90}
trunk flexion	$\theta_{3r} = 0$	[-90, 90]	{-90:15:90}
shoulder elevation	$\theta_{4r} = 0$	[0, 137]	{0:15:135, 137}
shoulder axial rotation	$\theta_{5r} = 0$	[-135, 120]	{-135:15:120}
elbow flexion	$\theta_{6r} = 0$	[-120, 0]	{-120:15:0}
forearm axial rotation	$\theta_{7r} = 112.5$	[-180, 188]	{-172.5:15:172.5}
wrist deviation	$\theta_{8r} = -90$	[-25, 25]	{-25:5:25}
wrist flexion/extension	$\theta_{9r} = 90$	[-40, 30]	{-40:5:30}

uration (Fig. 6.3) corresponds to the joint angle vector $(-135, 0, 0, 0, 0, 0, 67.5, 0, 0)$. The joint range values in Table 6.4 are the based on the hardware limits on the joint rotation.

6.4 Reachable object configurations in workspace

For each manipulator, we compute the possible initial object orientations for which the the oblique cylindrical grasp is achievable. This section describes the computation method for the kinematic analysis of the high-dimensional configuration space of the manipulator systems. In particular, to find inverse kinematics (IK) solutions for the high-DoF arm pose which also satisfy the joint limits, we pre-compute arm pose samples as possible IK initialization points. The palm segment of the hand is the end effector, and the analysis does not consider the configuration of the finger joints. First, the configuration space is sampled at discretized joint angle values to pre-compute a representative set of possible end-effector configurations in the workspace. Then, given a desired hand configuration for the pan grasp, we search for an inverse kinematics (IK) solution for the arm configuration using the nearby pre-computed samples as initialization points.

6.4.1 Pre-computation to sample configuration space

A set of possible IK initialization points is first generated by computing the end-effector configurations for a set of discretized joint angle values using forward kinematics (FK). Because the manipulator models have 9 or 10 DoFs, it is prohibitive to store the FK pre-computation for a fine discretization in joint space. We address this concern with a combination of several methods for reducing the number of pre-computed arm configurations that are saved. First, most of the joint angles are sampled at only 15-degree increments. The values used for the FK pre-computation are listed for human model in Table 6.2 and for the robot model in Table 6.4. Second, range of motion of the base joint (z_0) is limited to a subset of the full range to reduce the number of samples. The subset range is symmetric and centered around the reference configuration where the manipulator "faces" the table area in the $-y$ half-plane. Limiting the base joint rotation is equivalent to reducing the trunk torsion about the waist. Third, we use a decomposition approach by computing the wrist kinematics separately from the rest of the chain, such that two smaller pre-computation tables are stored. We describe the decomposition approach in more detail below.

The final three joint DoFs, collectively referred to as the wrist, are sampled separately from the other joints to decompose the FK pre-computation into two parts of the end-effector position and the wrist orientation. The end-effector position is computed for the discretized values described above for all DoFs except the final three joint axes, which are fixed at the reference configuration values for a "neutral" wrist configuration. In this way, the FK pre-computation only samples 6 or 7 DoFs instead of 9 or 10 DoFs, respectively, for the robot and human models. The resulting set of pre-computed hand poses are representative of the reachable end-effector positions, but not the end-effector orientations. In human manipulator model, the final three axes compose a spherical wrist because the forearm axial rotation, wrist deviation, and wrist flexion axes intersect at a common point. Thus, rotation about any of these axes will not change the end-effector position but only its orientation. The last three axes of the robot manipulator do not form a spherical wrist because the wrist deviation and wrist flexion axes are separated by non-zero link length $a_8 = 4$cm (Table 6.3). Due to the limited range of wrist deviation, the hand position computed for the neutral wrist orientation will differ at most by $2a_8\sin(25) = 3.4$cm from the possible hand positions at other wrist orientations. The wrist contribution to the end-effector orientation is pre-computed separately by sampling the range of motion of the three wrist DoFs (Table 6.2, Table 6.4).

6.4.2 Inverse kinematics for feasible arm configurations

A particular object configuration on the tabletop is graspable if there exists a manipulator configuration which achieves the end-effector configuration relative to the object for the specified grasp. We use the pre-computed sampling of the configuration

space described in §6.4.1 to select initialization points for the IK search.

Given a specific desired end-effector configuration, we determine a set of IK initialization points by choosing samples which approximately match the desired position and exactly match the desired orientation. For the position, the candidate samples of the neutral wrist configuration must lie within a 10cm cube centered around the desired end-effector position. For the orientation, the candidate samples of neutral wrist orientations must be within the allowable set of neutral-wrist orientations which could reach the desired end-effector orientation with respect to the three wrist DoF joint limits. This allowable set of neutral-wrist orientations is computed by transforming the desired end-effector orientation by the inverse transformations of the pre-computed wrist orientation contributions. For an admitted IK initialization points which satisfies both the position and orientation criteria, the three wrist angles can be computed directly such that the initialization point is the full 9 or 10 DoF configuration rather than the 6 or 7 DoF neutral-wrist sample. We found that initializing with the wrist angles which achieved the desired end-effector orientation was important for successfully finding an IK solution which satisfies the wrist joint range of motion.

Once the initialization set has been determined for the desired end-effector configuration, the IK solution is computed iteratively using a pseudo-inverse Jacobian method [Corke, 1996]. The implementation the pseudo-inverse Jacobian search does not respect joint limits, so that it is possible for the returned solution to violate the specified joint ranges of motion. This limitation is one reason why the success of finding a feasible IK solution that satisfies the joint angle limits is highly-dependent on the initial guess for the iterative search. Any IK solution which does not satisfy all the joint limits is discarded. All configurations in the initialization set are tested such that the overall search may generate multiple IK solutions which reach a single desired end-effector pose.

6.5 Results and discussion

For both manipulator systems, the reachable pan orientations are computed for both the straight and oblique cylindrical grasps for several pan positions on the tabletop. The kinematic analysis only determines which pan configurations are definitely graspable but cannot prove that solutions for the other configurations do not exist. Due to the sensitivity of the IK solution on the initialization, it is possible that the analysis is limited by the particular choice of discretized joint values for the FK pre-computation. To test the sensitivity to discretization, we perform the same IK analysis with the modification that joint angles for each candidate IK initialization is perturbed by a random value up to sample increment size. Including the perturbations of the IK initialization from the original discretization points did not expand the range of graspable orientations for any of the cases.

For the human manipulator, there are few pan positions for which more than 75% of the possible handle orientations can be reached by either grasp (Fig. 6.4). In particular, near the boundary of the workspace, only half or less of the pan orientations are graspable. For most pan positions, there is a larger capture region for the oblique grasp compared to the straight cylindrical grasp. The pan positions near the coordinates (-0.3m, -0.6m) are of particular interest since this region corresponds to the approximate x-y pan position relative to the waist for the lifting postures from the human motion capture experiments described in §4. The limited reachability of the straight grasp compared to the oblique grasp for this region may be a reason that there were fewer examples of this grasp observed from informal inspection of the human pan grasping. The capture region size of about 210 degrees and direction toward the body of the oblique grasp in this area is consistent with the large variability in the observed lift angles observed from the human examples (Fig. 4.5).

The robot manipulator has a larger workspace than the human manipulator due to the longer limb lengths (Fig. 6.5). Even though the robot has fewer degrees of freedom (9 versus 10 DoFs), the large range of motion of the PA-10 arm joints results in reachable pan configurations which are qualitatively similar in capture region direction as that for the human manipulator. Similar to the results for the human manipulator, the capture region size for the straight cylindrical grasp is generally smaller than that for the oblique grasp for radii greater than 0.5m. The area of interest for the lab demonstration setting are the pan locations about 1m in front of the robot ($(x, y) = (0m, -1.0m)$) due to the position of the table. In this area, the capture region points toward the manipulator's right side ($-x$) for the straight grasp and faces toward the manipulator ($+y$) for the oblique grasp, which is similar to the human manipulator. In terms of the potential for workspace extension due to pre-grasp rotation, there are also several pan positions for which only about half of the pan orientations are graspable. This case occurs particularly for the outer regions of the workspace with greater than 0.7m radii length from the base torso position.

6.6 Future work for workspace analysis

The workspace analysis presented in this chapter illustrated the workspace limits of a direct grasping strategy for a particular specified grasp. We hypothesize that pre-grasp rotation can extend the workspace of a particular grasp to include initial object orientations which were unreachable by direct grasping. This extension is possible because the grasp used for pre-grasp rotation is likely to be less constrained than that for direct grasping because the object load is partially supported by the surface. In the pan example considered in this chapter, non-prehensile contact with the handle could be used to pivot the object to a new orientation which is then graspable.

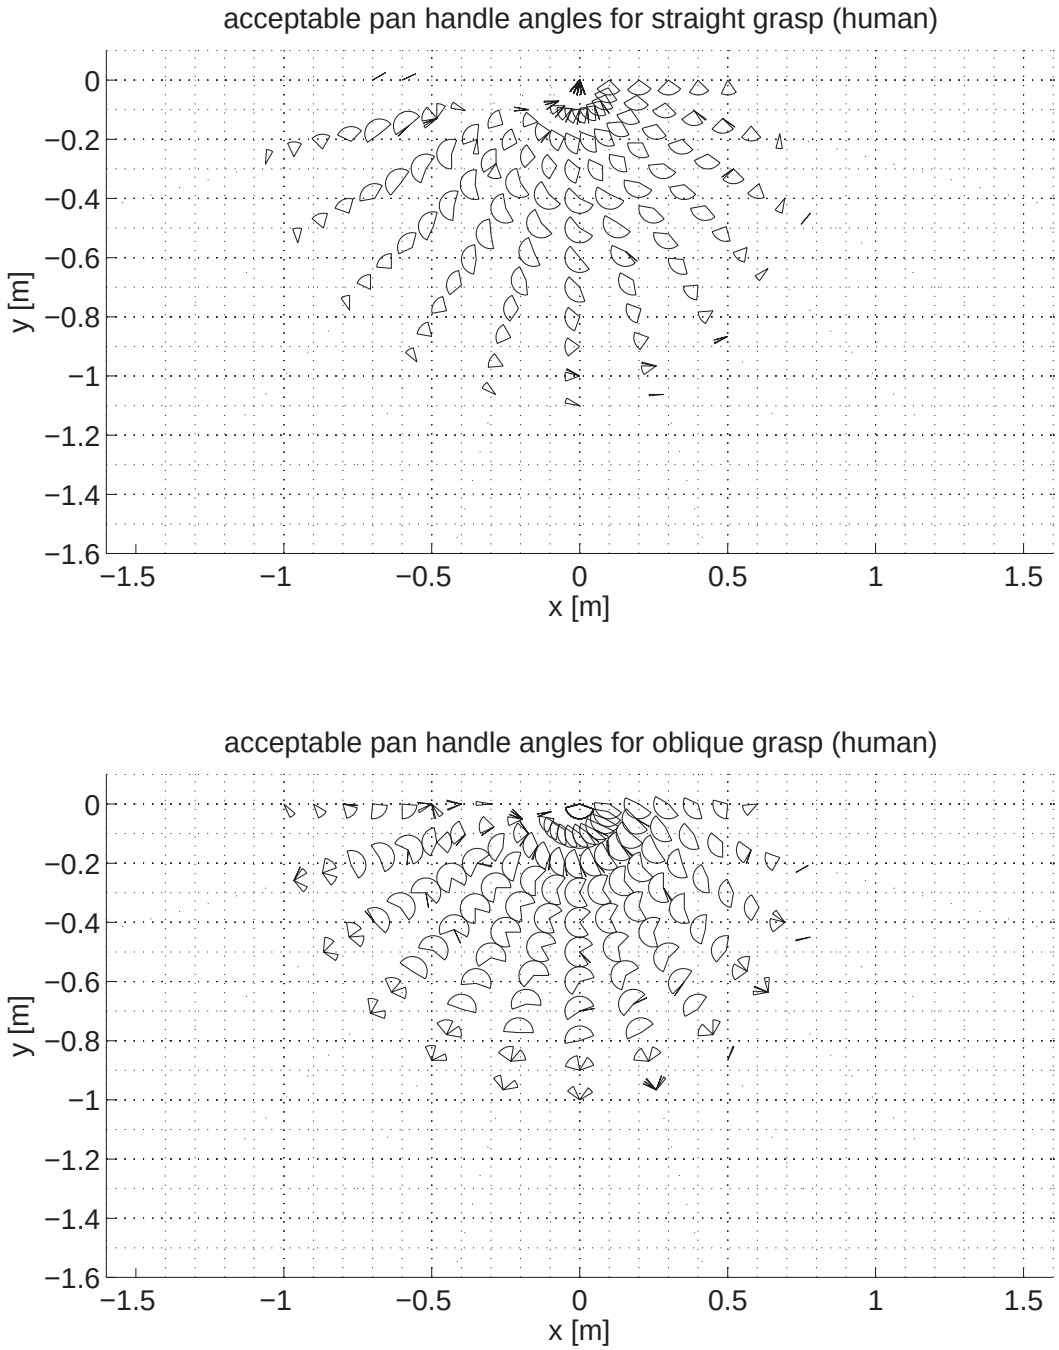

Figure 6.4: Visualization of the pan orientations reachable by the human manipulator. The sectors at each selected tabletop positions denote the range of the graspable handle directions.

The next chapter presents an implementation of a simple pre-grasp rotation action for grasping the pan object from multiple orientations. The implementation results

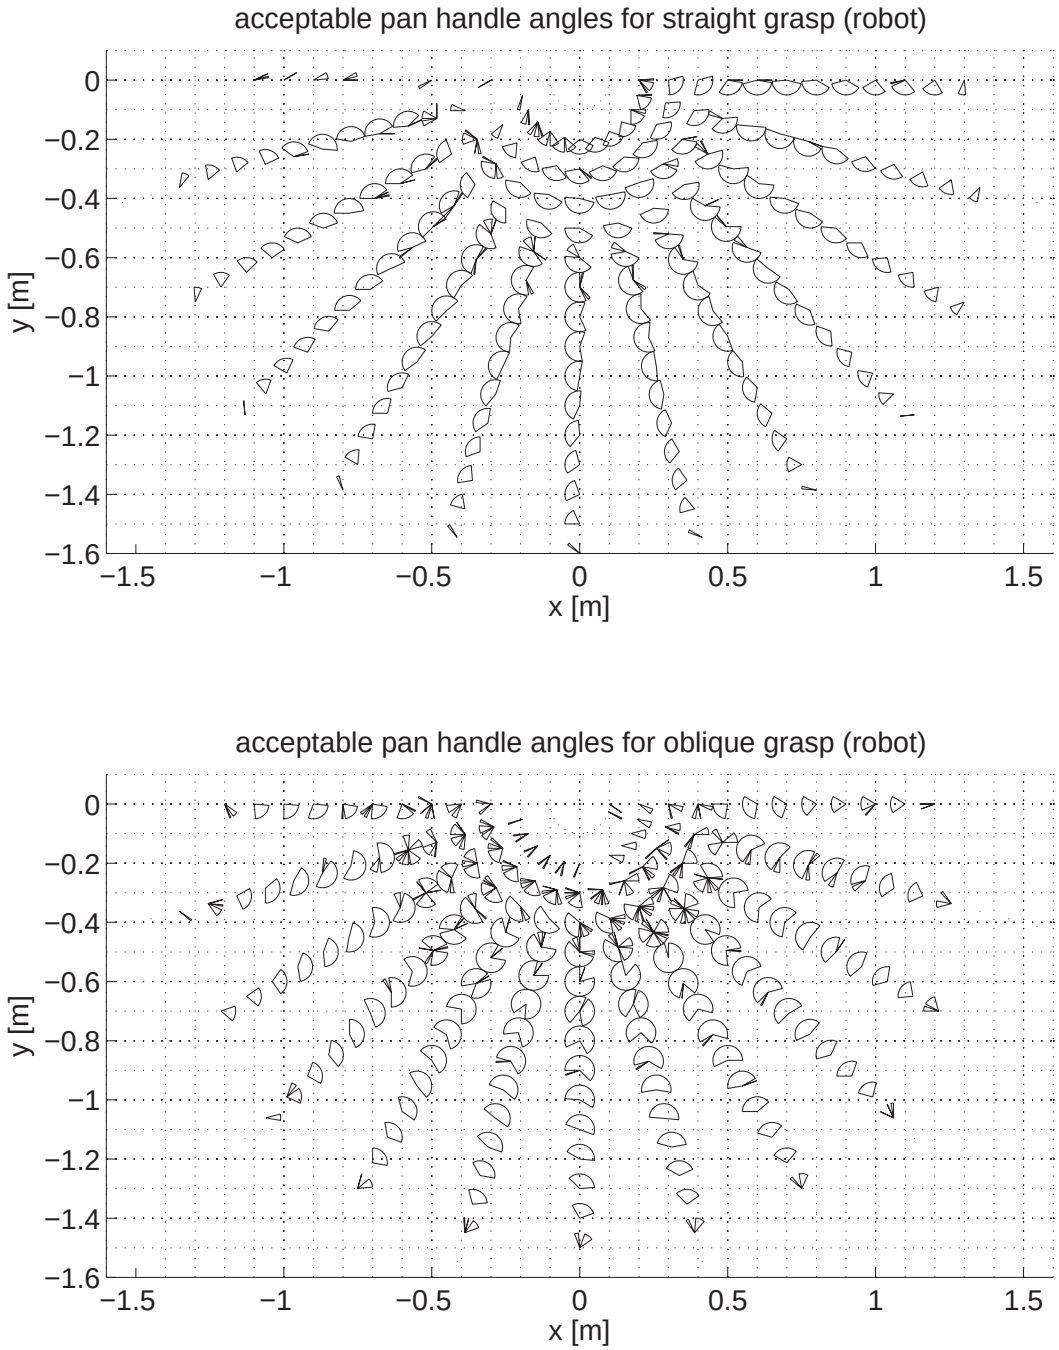

Figure 6.5: Visualization of the pan orientations reachable by the robot manipulator. The sectors at each selected tabletop positions denote the range of the graspable handle directions.

in §7 demonstrate how the pre-grasp rotation can extend a grasp's capture region to the full 360-degrees of possible object orientations for a particular object location in

the workspace.

For other object locations, pre-grasp rotation is also likely to extend the capture region. Here we suggest approaches for complete evaluation of the workspace, which are beyond the scope of this thesis but could be explored in future work.

First, pre-grasp rotation could be evaluated by quantifying the increase in the reachable workspace due to including non-prehensile pushing as a preparatory action. This quantification could be achieved by extending the workspace analysis of direct grasping to evaluate the possible object configurations which are reachable by any arm pose which achieves a pushing contact. It is expected that for several object locations, the allowable initial object orientations will be greater than that for direct grasping only.

Second, pre-grasp rotation could be evaluated according to the improvement of some grasp quality metric, such as the manipulator joint torques, even for initial pan configurations that are directly reachable. In this case, the focus would be on the increase in the number of arm poses for the grasping action following pre-grasp rotation. The multiple arm poses corresponding to a specific object orientation could be evaluated with respect to metrics such as manipulability and/or joint torque capability. This evaluation would demonstrate that even for situations when the initial object configuration is within the direct grasping capture region, pre-grasp rotation to a new object orientation may be able to achieve a higher quality arm pose for grasping.

7 Grasp reuse strategy for a robot manipulator

One potential benefit of the pre-grasp rotation strategy is that the object reconfiguration allows for reuse of a canonical grasp. This benefit was suggested by the analysis of the hand grasp orientation in the human motion studies in §4. Additionally, the workspace analysis in the previous chapter (§6) suggested how pre-grasp rotation might be used for object configurations on the boundary of the workspace to bring the object within a specific grasp's capture region.

In this chapter we examine the concept of pre-grasp object rotation as a strategy for re-using a single well-tuned grasp routine programmed for a robotic manipulator system. Here the grasp routine consists of the complete arm reaching action and the grasping movement, not just the final transform of the hand coordinate frame relative to the object. Such a prototype grasping routine may only successfully lift a specific object from a small set of initial orientations. Pre-grasp object rotation can extend the effective workspace of the single grasping prototype which is time-consuming to program.

We implemented a specific example of pre-grasp rotation: rotating a handled pan using a single-finger pushing contact prior to a lifting grasp routine. This rotation action was implemented on an anthropomorphic robot manipulator to investigate how pre-grasp rotation can extend the effective workspace of a well-tuned grasping action. The results show that the pre-grasp rotation of an object allows for a single grasp action to be reused for a much wider range of object configurations compared to direct grasping without rotation.

Work in this chapter has been previously published as part of Chang *et al.* [2008].

7.1 Related literature

The example presented in this chapter uses a simple pushing strategy to rotate an object prior to grasping it. The benefits of utilizing the support surface as a passive manipulator to increase workspace and load limits are familiar from the push-planning literature [Mason, 1986]. Much of that work is concerned with the automatic planning of push manipulations to reorient objects on a supporting surface with either known or uncertain contact conditions. For example, Lynch and Mason [1995] ex-

plore the conditions for complete control of a rigid sliding object while pushing to an arbitrary pose on a plane. In contrast, our emphasis is on the identification of specific heuristic action sequences which humans choose in response to prototypical situations. Thus, rather than modeling the general push dynamics, we have deliberately chosen an example with simple contact conditions and high error tolerance which admits a trivial ad-hoc solution to the push-planning problem.

Planning methods for pushing manipulation have been demonstrated on other humanoid platforms in recent work [Stilman and Kuffner, 2005; Hauser *et al.*, 2007; Ng-Thow-Hing *et al.*, 2007]. Our focus is on the benefits of using pushing to reuse grasping strategies for a wide range of object orientations, rather than on the planning of the pushing action alone.

The implementation presented in this chapter is also a simple example of sensorless manipulation. It does not measure object state, but assumes the object lies within a bounded set of initial configurations. It also uses a contact strategy chosen to reduce the uncertainty of the object state until it lies within the much smaller set of graspable states. Again, this type of strategy has been well-studied in the motion planning literature Lozano-Perez *et al.* [1984], but our emphasis is on identifying small sets of primitives which can be used to quickly construct adequate solutions to typical grasping problems rather than general optimal planning.

7.2 Hardware and task scenario

Our anthropomorphic manipulator system consists of a Mitsubishi PA-10 7-DOF manipulator with a 24-DOF Shadow Hand C3 end effector (Shadow Robot Company, London, UK). The kinematics of this system were previously described in the workspace analysis in §6. Foam padding is attached to the Shadow Hand to modify the palm geometry. The task specified for the robot implementation (Fig. 7.1) mimicked the task conditions in the human studies previously described in §4. The object starting position on the table is located 0.9m in front of the manipulator base. The object goal position is located on the table 0.35m to the left of the start position. As in the human studies, the object could start in one of several configurations. Twenty-four handle directions were selected to sample the full 360-degree orientation space at intervals of 15 degrees. The object in the robot experiments was a small cooking pan with a handle. The pan was empty and had a total mass of 0.46kg. Optical markers were attached to the pan in order to track the object configuration using a Vicon camera system for motion capture.

7.3 Open-loop routines for grasp reuse

The grasping strategy using pre-grasp rotation was implemented as two manually-programmed open-loop actions (Fig. 7.2). One action is the pre-grasp rotation

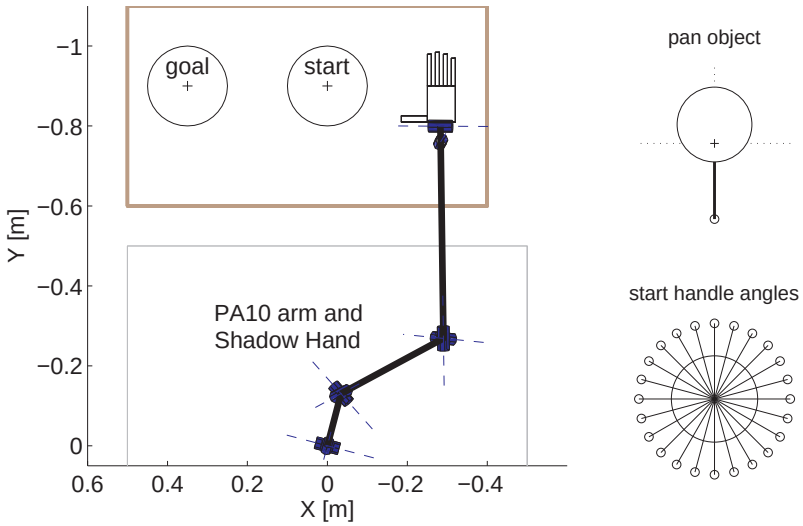

Figure 7.1: Layout of the manipulator experimental setting. (a) The right-handed anthropomorphic robot arm transported a handled pan from the start position to the goal position on a table setting. (b) The pivot point of pan object is placed at the start position. The pan started in one of 24 orientations defined by the handle direction. The orientations are nominally spaced by 15 degrees.

action for reconfiguring the handle orientation prior to grasping (Fig. 7.2a). The other action is the grasping action for lifting and transporting the pan by its handle (Fig. 7.2b). The two actions are executed sequentially for a complete manipulation action which transports the pan from any initial handle orientation (Fig. 7.3).

The grasping action was intended to mimic the underhand grasp of the pan observed in the human study trials with the right-hand unimanual constraint. The intended handle orientation for the grasping action was the direction facing toward the manipulator and slightly toward the right, as observed in the human examples (Fig. 4.5). The grasping action is a sequence of three motion components: an approach motion, the grasp motion, and the transport motion (Fig. 7.2b). In the approach motion, the Shadow Hand maintains a relaxed open-hand pose. During the grasp motion, the PA-10 arm configuration remains fixed while the hand's finger joints close around the handle. The hand then maintains a tightly-closed pose during the transport motion while the PA-10 arm moves to lift and transport left to the goal position. All three motion components were manually programmed for the specific object, intended handle orientation, start position, and goal position. Initial hand contact with the object often occurred during the approach motion, when the palmar side of the fingers contact the right edge of the handle.

The rotation action was implemented as a pushing motion using single-finger contact with the object to rotate the cooking pan around its natural pivot point

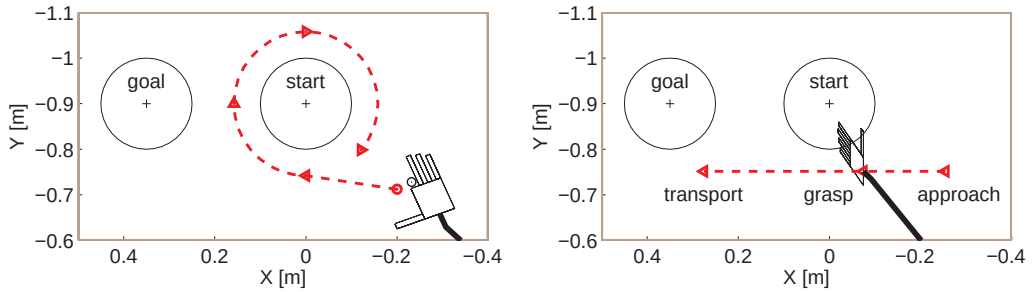

Figure 7.2: Schematic of the two open loop actions implemented on the robot manipulator. (a) In the pre-grasp rotation action, the hand traces a circular arc around the starting position to rotate the pan using a single index finger contact with the edge of the handle. (b) In the grasp action, the hand moves in a relatively straight path across the table to first approach the object from the right, grasp the handle, and then transport the object left to the goal.

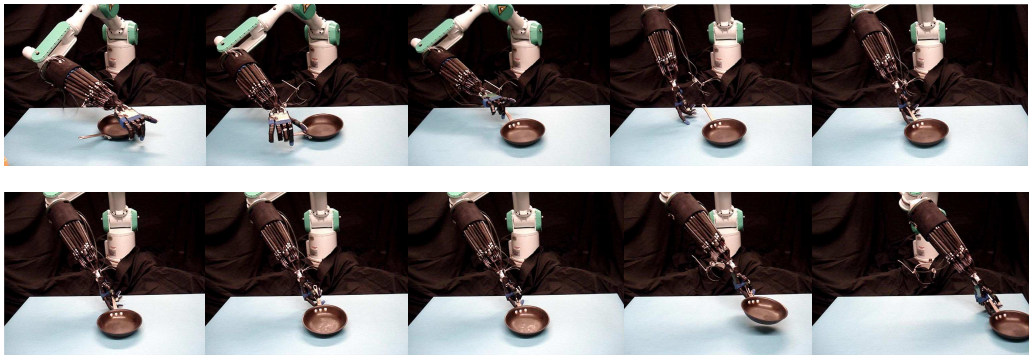

Figure 7.3: Implementation of pre-grasp rotation on the robot manipulator. The frames in the top row show the pre-grasp rotation action and approach prior to the grasp. The frames in bottom row show the grasp of the pan handle from the canonical orientation and the transport of the pan.

(Fig. 7.2a). The index finger was flexed 90 degrees such that it pointed normal to the palm. The thumb and other three fingers were extended in the plane of the palm, which remained parallel to the table during the rotation action. The index fingertip first approaches the object start position along an initial straight segment. Then the index fingertip traces a circular arc of 315 degrees in a clockwise direction around the object start position and ends within the intended grasp capture region. Tracing the full 315 degree arc path allows the open-loop rotation action to be executed identically regardless of the initial pan handle angle. The 45 degree gap in the circular path was deliberately designed to avoid contacting the object if the handle starts within the original capture region of the grasping action alone.

7.4 Empirical evaluation of workspace extension

To measure how well the pre-grasp rotation enables reuse of the grasping action, we compared the grasping action alone to the sequence of the preparatory rotation followed by grasping action. The two methods were each tested on the different initial handle orientations in a set of 24 consecutive trials. The manipulation was considered successful if the grasp lifted the pan off the table surface and transported the pan to the goal position.

In the 24 consecutive trials using the grasping action alone, the manipulator successfully grasped and transported the pan to the goal position for 4 of the 24 initial handle angles (Fig. 7.4(a).) The empirical capture region of the grasping action alone was 45 degrees. The grasp component works best for the two handle angles in the center of the capture region. For the two outer angles of the region, the handle was rotated toward the center of the region by either a clockwise push from the approach component or a counterclockwise push from the outstretched fingers in the grasp component.

In contrast, when the pre-grasp rotation action is used prior to the grasping action, the manipulator completed the transport task successfully for all 24 of the 24 consecutive trials (Fig. 7.4(b).) In all but one of the trials, the index finger made contact with the handle at some point during the rotation arc path and pushed the handle clockwise. The one exception was the second leftmost handle angle of the four orientations already within the capture region of the grasping action alone. Because the handle orientation was already centered in the original grasping capture region, the object remained stationary during the rotation action, as intended by the design of the circular push path.

The pre-grasp rotation action consistently rotated the pan into the grasping action's capture region (Fig. 7.5). The handle angles after rotation and before the grasp were all within a 15-degree range. The grasping action further reduced the uncertainty in object pose at the goal position. The handle angles after the grasping and transport action were all within a 6-degree range.

7.5 Comparison to alternative grasp reuse strategy

In our implementation we have focused on the idea of reusing an entire grasp action consisting of the approach motion, grasping motion, and transport motion components for both the arm and the hand. Each motion component was manually programmed for the specific task tested in our experiments. The grasping motion is the most critical component that is manually-programmed, because the hand pose must be carefully tuned in order to securely grasp the thin handle of the pan. During the grasping motion component, the manipulator arm configuration was stationary while the finger joints closed around the pan handle.

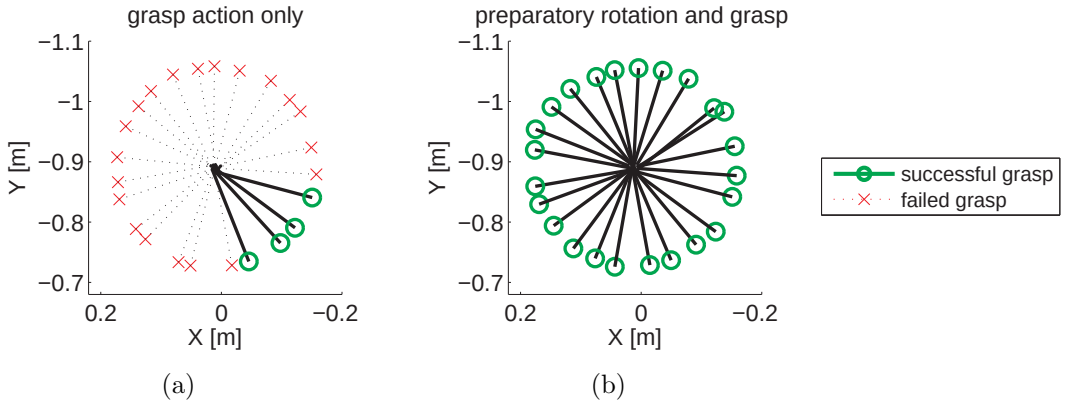

Figure 7.4: Empirical test results for two open-loop manipulation strategies for transporting the pan. The plotted handle directions are coded according to whether the trial resulted in a successful transport of the pan to the goal position. (a) Initial object pose for the trials using only the grasp action. The grasp sequence can grasp the pan from 4 of 24 tested angles for a 45 degree capture region. (b) Initial object pose for trials using the pre-grasp rotation action before grasping. The manipulator was able to grasp the pan from all 24 tested angles using rotation.

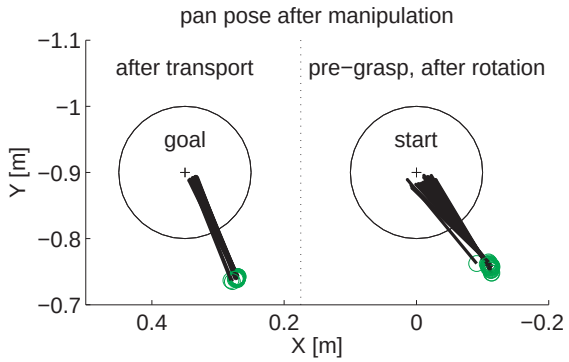

Figure 7.5: The object handle pose for the trials using pre-grasp rotation strategy. After the rotation action (right), the pre-grasp action has significantly reduced the uncertainty in object pose by pushing the handle into the grasping capture region (15 degree range in orientation). After the grasping action (left), the object has been transported consistently to the final goal position (6 degree range in orientation).

An alternative scheme to reuse a well-tuned grasp would be to re-plan the arm configuration and arm motion components without changing the hand motion during the grasp. In this way, as long as the same relative configuration is maintained between the palm of the hand and the object, the same tuned finger motion for grasping can be used with new arm configurations. This grasp-reuse scheme was considered previously in the workspace analysis in §6 and only considered the kinematic reachability of the specified hand pose relative to the object handle.

For each possible handle orientation of the pan given the same center position, we computed the required palm transform in the workspace required to maintain the same relative configuration to the object handle. Given the desired palm transform, we searched for an inverse kinematics solution of an arm configuration which would achieve the desired end-effector (palm) configuration. The inverse kinematics solution was computed iteratively using a pseudo-inverse Jacobian method [Corke, 1996]. Because the inverse kinematics solutions are highly-dependent on the initial guess for the iterative search, multiple initial configurations were tested for each desired palm pose. The guesses were selected from a database of pre-computed arm configurations discretized in joint space. Any pre-computed configuration whose end-effector position was in the neighborhood of the desired palm position was evaluated as an initial guess.

The results of the kinematic analysis (Fig. 7.6) show that the same relative transform between the object and hand is reachable for a wide range of handle orientations much larger than the empirical capture region of the single grasping action. However, about one-third of the possible handle orientations are still unreachable by the manipulator. Thus, even under the considered alternative grasping scheme, pre-grasp rotation could still be used to achieve successful grasps for unreachable handle directions.

7.6 Discussion

Our implementation of a specific open-loop rotation was intended to test the merit of the pre-grasp rotation strategy in terms of extending the effective grasp capture region. Under this scheme, the pan may be rotated in almost a full circle for some initial configurations. The length of the rotation for some of the handle directions facing toward the manipulator makes the strategy sensitive to the initial position of the pan. It is possible for the hand to lose contact with the handle during the rotation if the pivot point is not placed properly. The strategy we implemented can be extended by adding additional open-loop actions tuned for different sets of initial conditions, which might be determined by only a few bits of sensor data. A simple example would be to minimize the overall pan rotation by using a counterclockwise preparatory motion for handle orientations on the left. This modification would avoid the need to maintain contact for a long duration along the circular push-path and may improve the robustness of the rotation action in achieving the desired handle configuration prior to the grasp.

In addition to extending the capture region of a well-tuned grasping action, the pre-grasp rotation strategy may be desirable in humanoid robots by making the manipulator motion appear more human-like. Features such as bidirectional rotation and a more relaxed hand pose for multi-finger pushing contact could improve the manipulator's natural appearance.

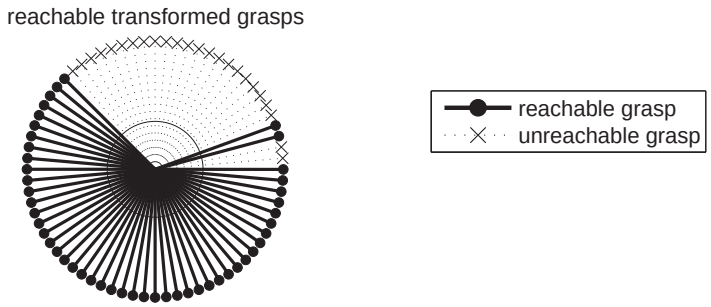

Figure 7.6: Reachable object handle orientations computed for a scheme where the reaching motion is re-planned but the relative configuration between hand and object during grasping is the same. If a novel reaching motion is planned for each possible handle orientation, the manipulator may be able grasp from a wider capture region then the single open-loop grasping action. However, about one-third of the orientations are still unreachable.

Other similar preparatory manipulation strategies include sliding, rolling, or tumbling maneuvers which re-configure the object prior to grasping. In the human subject experiments, translational sliding of the object was not constrained in any of the task scenarios. The human motion capture data does reveal that some of the unimanual manipulation resulted in both planar rotation and translation prior to lifting. In the robot experiments, we found that a simple pivoting motion was sufficient to increase the capture region without specifically programming a translational displacement action.

Directions for future work in robotics include examining which features of these preparatory manipulation strategies should be imitated in humanoid robots. Some features, such as the arm configuration, might be essential for a robot to appear human-like. Other features, such as optimization of the joint torques or grasp quality, might be important heuristics for achieving robust performance in difficult task conditions. We explore the ways that pre-grasp rotation can improve manipulation plan success and quality in §8 and §9.

8 Planning feasible pre-grasp interaction for transport tasks

The results from the human studies suggest that the potential benefits of pre-grasp manipulation include an increase in task success rate as well as an increase in the quality of the primary manipulation task performed after object acquisition.

In the previous chapter, we demonstrated that pre-grasp object manipulation, in the form of pre-grasp rotation of a handled cooking pan, can improve task performance by extending the effective workspace of the well-tuned, robust grasp. The pre-grasp rotation, grasp, and transport actions were all manually-programmed for a specific object that was presented at a particular location in the workspace.

For a task with less structured initial conditions, it is impractical to manually specify the manipulation actions for a large set of potential object placements. In this chapter, 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 for the load-supporting postures of the transport action.

The focus of this chapter is a method for finding feasible manipulation plans in a tractable manner. In the next chapter, we explore extensions to the core method presented here which could further improve the action quality of the selected task.

The material in this chapter has been previously presented in the proceedings of the *IEEE International Conference of Robotics and Automation* [Chang *et al.*, 2010b].

8.1 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. 8.1). Usually there are several ways to accomplish each action due to kinematic redundancy in the manip-

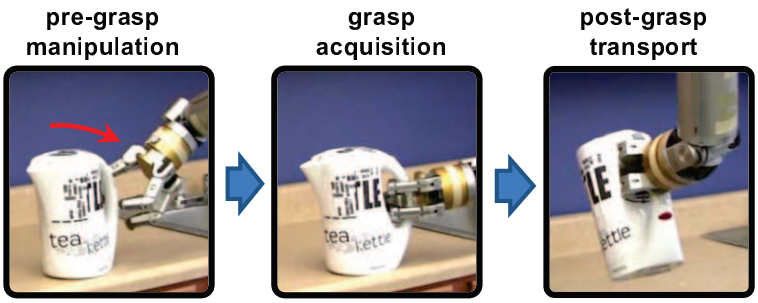

Figure 8.1: A fetching task requires a complete manipulation plan including grasp acquisition and transport. We present a method for incorporating pre-grasp manipulation, specifically object rotation, in the manipulation plan. The object adjustment resulting from pre-grasp manipulation increases the task success rate and improves the quality of the transport plan.

ulator as well as freedoms in the task specification. Examples of task freedoms are the graspable regions on the object [Gienger *et al.*, 2008] and the allowable object poses at the goal location [Berenson *et al.*, 2009c].

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 [Chang *et al.*, 2008], enabling more robust task completion and higher quality grasping postures at the time of object lifting [Chang and Pollard, 2009].

In this study, we augment the typical action sequence of grasp acquisition followed by object transport with an additional preparatory action component. This method results in a complete manipulation sequence of pre-grasp manipulation, grasp acquisition, and post-grasp transport, illustrated in Fig. 8.1. To achieve this sequence, a planner needs to automatically decide on the extent of pre-grasp manipulation. This decision 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 chapter is structured as follows: First, §8.2 discusses related work in manipulation path planning. §8.3 describes the action components composing the complete transport task sequence, and §8.4 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 §8.5. Finally, §8.6 discusses the method's limitations and directions for future work.

8.2 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 [Latombe, 1991]. A large body of work has focused on modeling the manipulator's configuration space (c-space) [Lozano-Perez, 1983] in a preprocessing step to find the free-space regions [Lozano-Perez, 1987] or connected paths for a static environment [Ahuactzin *et al.*, 1998; Kavraki and Latombe, 1994; Kavraki *et al.*, 1996]. Recent work has also developed several variants of sampling-based planning techniques, such as Rapidly-exploring Random Trees (RRTs) [LaValle and Kuffner, 2000], for path-planning in dynamic environments. These studies have focused on finding manipulation plans for reach-to-grasp motion [Bertram *et al.*, 2006; Weghe *et al.*, 2007; Berenson *et al.*, 2007; Diankov *et al.*, 2008b], transport from a starting configuration where the object is already grasped [Yamane *et al.*, 2004; Berenson *et al.*, 2009b], or sequences of alternating reaching and transport paths [Simeon *et al.*, 2002; Sahbani *et al.*, 2002; Simeon *et al.*, 2004].

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 [Lozano-Perez *et al.*, 1984; Mason, 1986; Lynch and Mason, 1995], and recent work also demonstrated push-planning methods on humanoid robots [Stilman and Kuffner, 2005; Hauser *et al.*, 2007; Ng-Thow-Hing *et al.*, 2007]. Other possible pre-grasp actions are toppling and tumbling [Lynch, 1999], regrasping [Tournassoud

et al., 1987; Simeon *et al.*, 2004], and whole-body pivoting maneuvers [Yoshida *et al.*, 2007, 2008, 2010]. 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 [Chang and Pollard, 2009]. We extend the previous work by presenting a method for autonomous planning of transport manipulation tasks. Similar to the approach in Hauser *et al.* [2007], our method is based on the selection of good transitions between actions to obtain a successful sequence.

8.3 Manipulation transport task

In this chapter 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. 8.1):

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.

8.3.1 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 [Chang *et al.*, 2008; Chang and Pollard, 2009]. 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.* [2008b] 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.

8.3.2 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 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 end-effector poses for grasping (points in the 6-DoF task-space) and/or as regions in task-space [Berenson *et al.*, 2009c]. 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 [Berenson *et al.*, 2009a].

8.3.3 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 regrasping. This transport action 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.* [2009a] 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.

8.4 Transition state selection for planning

Fig. 8.2 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. The object pose 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. 8.2, 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 $\mathbf{q}_t(\theta_t)$.

8.4.1 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:

$$c(\mathbf{q}) = \frac{1}{p(\mathbf{q})} \tag{8.1}$$

and has possible values in the range $[0, \infty]$. The overall manipulator payload p is computed from the load limits of the individual joints, $\tau_{\max}$, as follows. The joint torques $\hat{\tau}$ necessary to statically support a unit payload force $\hat{f}$ are computed from

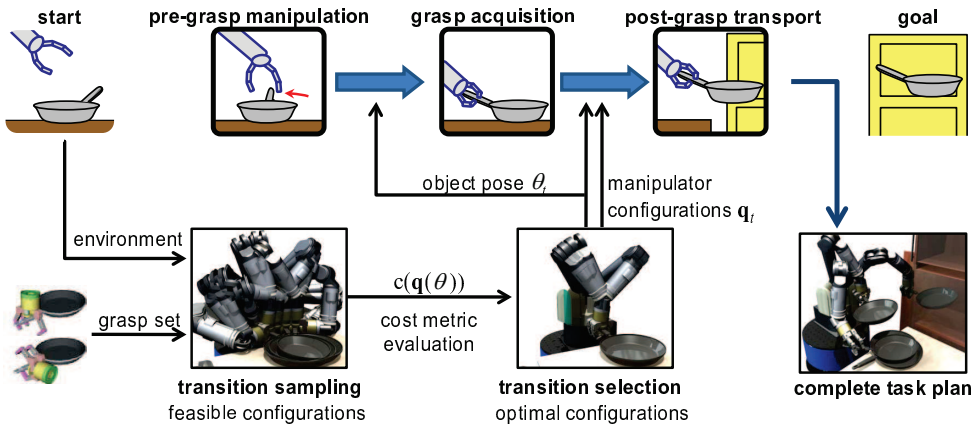

Figure 8.2: Overview of the transition selection method for planning a complete object transport task. The specification of the individual action components requires selection of the target object orientation θ_t and the manipulator configuration $\mathbf{q}_t$ at the grasp acquisition point. The object 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.

the Jacobian of the manipulator pose:

$$\hat{\tau}(\mathbf{q}) = J(\mathbf{q})^{\mathrm{T}} \hat{f} = J(\mathbf{q})^{\mathrm{T}} \begin{bmatrix} 0 \\ 0 \\ -1 \end{bmatrix}. \tag{8.2}$$

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:

$$p(\mathbf{q}) = \min_j(\mathbf{p}_{\max_j}(\mathbf{q})) = \min_j \left| \frac{\tau_{\max_j}(\mathbf{q})}{\hat{\tau}_j(\mathbf{q})} \right|. \tag{8.3}$$

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.

8.4.2 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 straightforward 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 [Berenson *et al.*, 2009c], 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 [Diankov and Kuffner, 2008]. 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 $\mathbf{c}(\theta)$.

8.4.3 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 feature by normalizing the threshold percentile cost by the number of candidate manipulator configurations N:

$$s(\theta) = \frac{\text{percentile}(\mathbf{c}(\theta), \text{t})}{N(\theta)}. \tag{8.4}$$

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

$$\theta_t = \arg\min_{\theta} s(\theta) \tag{8.5}$$

is selected as the target object pose for pre-grasp rotation. These target transition variables θ_t and $\mathbf{q}_t$ determine the sub-goals necessary for planning the individual action components described in §8.3.

8.5 Validation

8.5.1 Simulation experiments

The approach was tested on multiple transport task scenarios, shown in Fig. 8.3. 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 sampling 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 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.* [2008b].

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.

Figure 8.3: 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.

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 [Berenson *et al.*, 2009a]. 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 C 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:

$$C = \sum_{i=1}^{N} c(\mathbf{q}_i)\|\mathbf{q}_i - \mathbf{q}_{i-1}\|. \tag{8.6}$$

8.5.2 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. 8.4). 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.

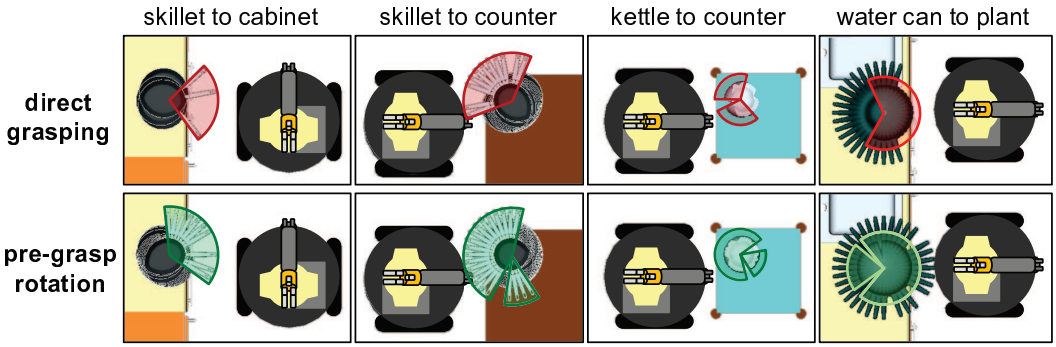

Figure 8.4: Comparison of plan success over different initial object orientations. The qualitative annotations highlight the general range of the handle orientations for which a feasible transport plan was found. In the last example problem, the watering can spout extends further than the handle, thus the highlighted region is opposite the spout spokes.

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. 8.5 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 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. 8.6, 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 difference in potential cost gains at the lifting grasp transition point is shown in Fig. 8.7. The variation in the best grasping posture cost (Fig. 8.7) is the inverse of maximum payload values (Fig. 8.8). The maximum payload curve is analogous to the lifting capability profile measured earlier in the human studies (Fig. 5.6(b)), where the preferred lifting region corresponded to higher lifting capability. For the robot manipulation scenarios tested here, the range is greatest for the grasps of the skillet pan, suggesting that pre-grasp rotation together with posture selection could be more effective in lowering the transport cost compared to the other scenes. For the postures in the best 10% threshold set (Fig. 8.9), the limiting joint was either the second shoulder joint or the distal wrist joints

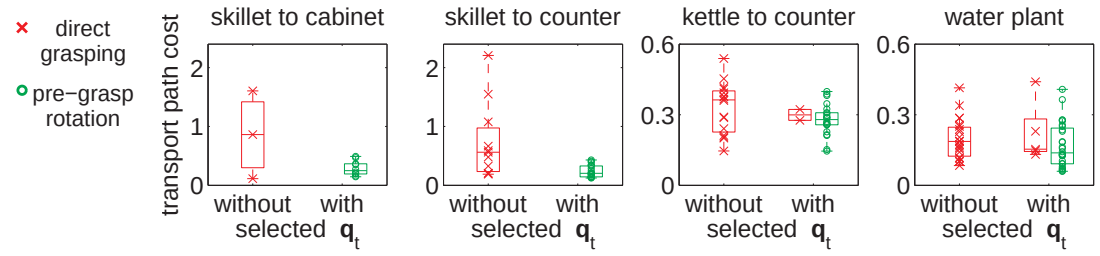

Figure 8.5: Comparison of transport path cost for the successful plans. Costs of the individual trials are denoted by × and ∘, and the box plot overlay indicates the median and quartile values for each column. On the left of each plot, the costs of the direct grasping approach result from transport paths 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.

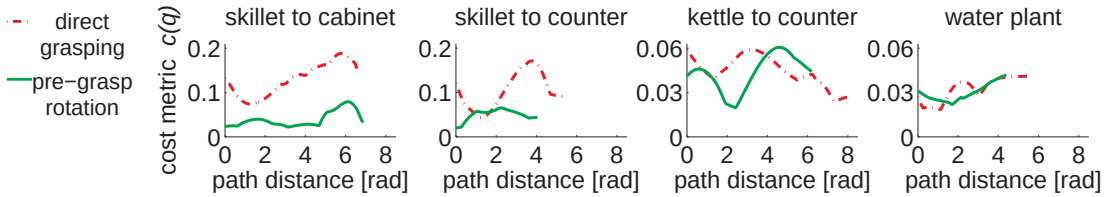

Figure 8.6: Evolution of the cost metric for the trials with median transport path cost. The path distance on horizontal axis is the Euclidean distance of the joint angles in the 7-DoF c-space. In the first two example problems, there is a large range of cost values, and paths which are initialized at 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.

for the skillet pan grasps, while the limiting joint was most often one of the distal wrist joints for the kettle and watering can grasps.

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 8.1 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 com-

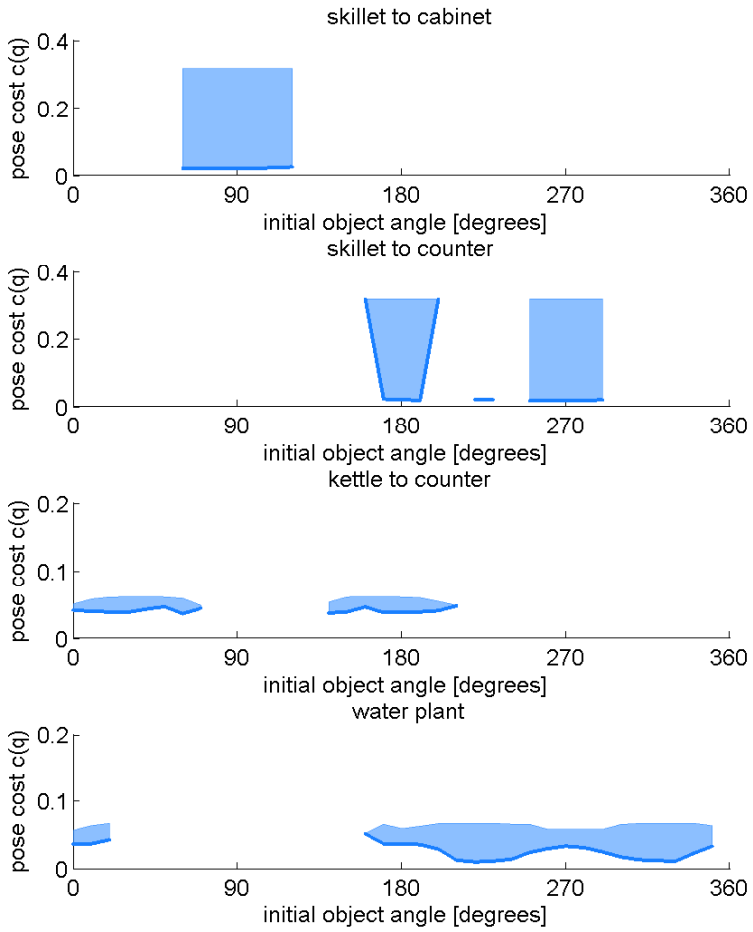

Figure 8.7: Cost range of lifting posture over different object orientations. The source postures contributing to the cost range are those in the best 10% threshold set for each object orientation.

pleted 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

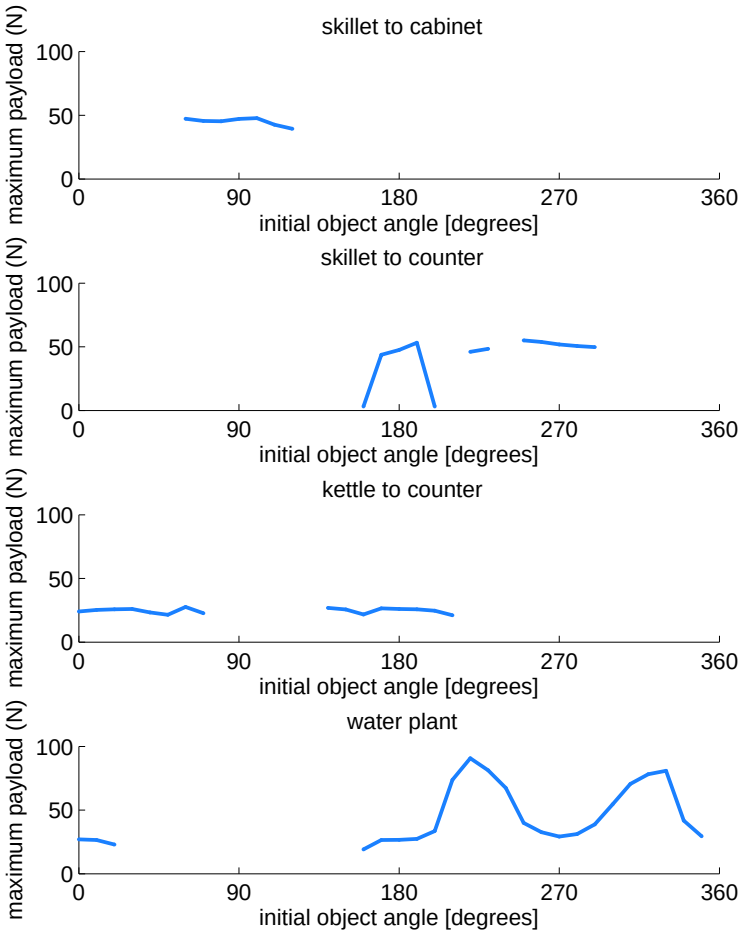

Figure 8.8: Maximum payload of the best posture over different object orientations. The maximum payload is the inverse of the payload safety margin cost.

pre-grasp interaction.

8.5.3 Physical demonstration

We demonstrated successful transport task plans on HERB [Srinivasa *et al.*, 2010], a manipulator platform consisting of the Barrett WAM 7-DoF arm and the BarrettHand end effector (Barrett Technology Inc., Cambridge, MA). Fig. 8.10 shows example sequences of the successful plans. Please also see videos of the physical task execution at:
http://personalrobotics.intel-research.net/projects/pregrasp.php.

The motion plans were executed open-loop on the robot, after the object was

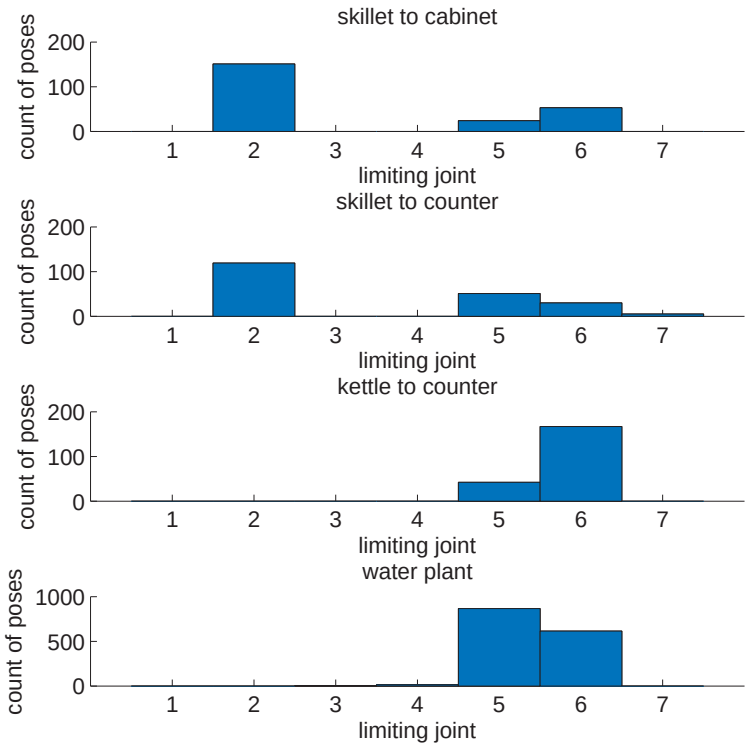

Figure 8.9: Distribution of the joint index which is the limiting joint for the lifting grasp postures. The source postures contributing to the histogram count are those in the best 10% threshold set for each object orientation.

Table 8.1: Additional planning time for pre-grasp manipulation compared to direct grasping approach.

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%

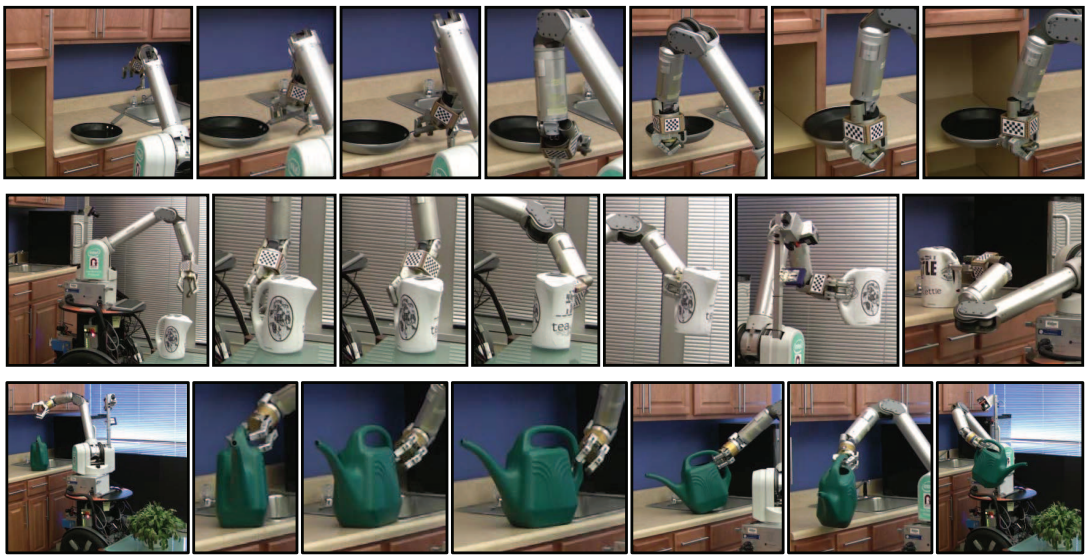

Figure 8.10: Example transport task plans demonstrated on the physical robot.

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.

8.6 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. A possible explanation for this result is that the limiting joints determining the maximum payload are often the distal joints at the manipulator "wrist." The constraint to maintain nearly-upright object orientations during transport may force the wrist configuration to remain similar to that of the selected transition state at 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.

9 Action quality and optimization

In the previous chapter, we presented a method for finding feasible object transport plans that automatically incorporate pre-grasp rotation. To make the planning computationally tractable, our method focused on the selection of key transition points in the overall manipulation, which determined the start and end point configuration constraints for the individual action path plans. One result of optimizing the selected grasp posture transition point was the indirect improvement in the following transport path. The effect was indirect because the path planner was agnostic to the path cost metric and the selected grasp posture was the only optimization control parameter.

This chapter further explores the relation between the optimization cost function and selected grasping posture for the pre-grasp rotation. First we review previous work investigating possible cost criteria for explaining motor action selection. We then present our experiments on the effects of particular cost functions on the pre-grasp object rotation. The experiments examine the effects of cost functions on the evaluation of individual isolated grasping postures, as well as the effects on the choice of whole motion plans over a sequence of manipulator postures.

9.1 Cost functions for action evaluation

Robot manipulators, like humans, are able to achieve a specified task with multiple alternative actions when there is redundancy with respect to the constraints. This redundancy can occur when the system is highly-articulated or over-actuated such that the task specification is under-constrained. When there are several alternatives available for feasibly completing an action, a quality or cost metric can be used to evaluate and select a preferred action to execute. Modeling such an optimization function for action selection can also be used to predict, construct, or retarget new motions and behavior.

Research in the biomechanics and neuroscience fields has investigated a number of possible optimization functions for explaining human motion. For the upper extremity, both minimum jerk [Flash and Hogan, 1985] and minimum torque change [Uno *et al.*, 1989] have been proposed and evaluated as underlying optimization criteria for arm reaching movements in free-space motions and unloaded tasks. Biomechanics

simulation of human motion has also examined energy or torque-based effort to describe full body postures for lifting heavy loads [Dysart and Woldstad, 1996; Chang *et al.*, 2001]. Other constraints that have been investigated as factors influencing motor behavior are joint motion limits [Zhang and Rosenbaum, 2008] and strength limits [Lee *et al.*, 1990; Burgess-Limerick and Abernethy, 1997]. Subjective perception of postural or action comfort may impact voluntary selection between alternatives, as illustrated by Rosenbaum's end-state comfort effect [Rosenbaum *et al.*, 1992; Short and Cauraugh, 1997; Fischman, 1998], for example. Comfort and discomfort may also be objectively correlated to biomechanical strength or actuation capabilities [Lee *et al.*, 1990; Rosenbaum *et al.*, 1996].

In robotics, algorithms for optimizing similar cost functions have been used to achieve specific application goals. In motion path planning for both mobile robot navigation and manipulation actions, collision-free paths of minimal length are of interest for efficient task completion. The path length may be defined by a spatial distance [Jarvis, 1985; Barraquand *et al.*, 1992; Shiller, 2000], the trajectory's time duration [Shiller and Dubowsky, 1987, 1991], or combination of both space and time [Jarvis, 1989]. A reference path which is feasible but not necessarily efficient can be modified with shortcutting methods [Chen and Hwang, 1998; Kavraki and Latombe, 1998] to reduce the path length while maintaining obstacle avoidance or improving path smoothness [Quinlan and Khatib, 1993; Brock and Khatib, 2002; Ratliff *et al.*, 2009]. Positioning error in the end-effector pose is also a relevant cost function in tasks requiring tool path accuracy. For example, Aboaf and Paul [1987] proposed an inverse Jacobian control method for minimizing orientation error near kinematic singularities while guaranteeing position accuracy. Other criteria of interest include control effort as defined by torque or energy [Hollerbach and Suh, 1985; Yen and Nagurka, 1987; Shiller, 1994]. In the presence of environmental modeling uncertainty, active sensing policies may select actions based on maintaining visibility [LaValle *et al.*, 1997; Torabi *et al.*, 2007; Diankov *et al.*, 2009] or tactile contact with the target [Lynch *et al.*, 1992].

Earlier in §4 and §5, we framed pre-grasp interaction in the context of alternative movement strategies and examined possible evaluation metrics that could consistently explain the preference for pre-grasp rotation over direct grasping. Our analysis focused on the choice of individual body postures and evaluated these postures by quasi-static, torque-based measures of the load on anatomical joints. In this chapter, we investigate the quality of pre-grasp rotation plans in terms of the torque-based payload metric on the grasping posture presented in §8. In addition, we investigate whole-path metrics measuring the payload margin over the path distance and the smoothness of a trajectory defined by the motion jerk. The payload margin metric is relevant for decreasing mechanical loads for a robot performing similar transport tasks repetitively. Motion smoothness is an example of one potential cost to minimize in order to create natural-looking movements for a robot operating in the presence of humans.

9.2 Overview of experiments

Overall this chapter presents four sets of experiments that investigate how variations in the optimization cost function change the choice of object lift-off angle and thus the amount of required pre-grasp rotation.

The first two sets of experiments examine how variations in the cost function change the choice of individual grasping postures. First, we consider the payload safety margin metric from the previous chapter and test whether changes in the manipulator joint torque limits result in consistent changes in the limiting joint determining the maximum payload (§9.3). Second, we show that in scenes where the same limiting joint restricts most of the feasible grasping postures, a simpler objective of minimizing a single joint torque value is a possible alternative cost function for locally improving a grasping posture (§9.4).

Optimizing only the grasping posture affects the key transition point in the overall manipulation path, but it is not the only control parameter for determining the entire motion plan. The third and fourth sets of experiments examine methods for optimizing cost functions which are evaluated over an entire motion plan instead of only a single posture or point. Such an optimization would be suitable in applications where additional computational resources can be devoted to improving the quality of a feasible plan. The purpose of the presented experiments is to evaluate the potential gains in plan quality for pre-grasp interaction when a direct grasping plan or an alternative pre-grasp rotation plan is also feasible for the task.

In the third set of experiments (§9.5), we continue to focus on the quality of the transport path after grasp acquisition, which is the only component action where the manipulator supports the entire object load. We present and evaluate optimization extensions at two levels in the path planning process. First, we consider optimization of a single feasible transport path at the level of path segment changes. Next, we also consider the optimization of the transport path cost over a set of multiple feasible transport paths for the same task scenario. For these experiments, the cost metric for an entire path is based on the individual configurations (points) that compose the path, irrespective of the path timing.

In the fourth set of evaluations (§9.6), we analyze the potential path quality gains when considering the overall manipulation, not just the transport portion. The entire action plan includes the pre-grasp rotation, reach-to-grasp motion, and the transport action. We evaluate the plan according to a minimum-jerk cost metric. The minimum-jerk measure depends not only on the sequence of points along the path, but also the assignment of a specific timing to achieve a particular trajectory.

9.3 Effect of different manipulator payload limits

We investigated the effect of changing the manipulator payload limits to gain further insight into the distribution of costs for the candidate grasping postures. One possibility is that the combination of weaker distal joints and end-effector orientation constraints produces a limited set of preferred grasping postures. Less restricting limits for the distal torques could result in decreased payload costs for the grasping postures meeting the threshold percentile. This change may correspond to a decreased range of costs over the possible object orientations, such that the pre-grasp rotation to a preferred angle has negligible advantage over a feasible direct grasp.

We examined this hypothesis by evaluating the payload safety margin for the same manipulator with altered joint torque limits. Table 9.1 shows the original manipulator torque limits used in all previous computations. These limits were based on the mechanical specifications of the Barrett WAM (Barrett Technology Inc., Cambridge, MA). The altered torque limits in Table 9.1 for a synthetic situation assigned identical torque limits to all 7 joints. The magnitude of 30 N-m torque limits was selected as a middle range limit similar to the actual joint 4 limit of 28.8 N-m. The identical individual limit values resulted in weaker proximal joints relative to distal joints, since each proximal joint must also support the distal links' masses as well as the end-effector payload.

The payload safety margin costs $c(\mathbf{q})$ were computed for the candidate grasping postures for the two task scenarios considered in this section. Fig. 9.1 shows how the torque limits change the costs for the postures within the threshold percentile. The number of total feasible grasp poses is also shown per candidate object orientation θ. For both task scenarios, there was generally a larger range of cost values for object orientations with a larger number of total poses. However, the effect of the altered payload limits was different for the two scenarios. For the kettle scene, the lowest 10% costs ranged from 0.01 to 0.06 over all possible object lift angles, and this small range suggests limited utility from handle reorientation by pre-grasp interaction. With the altered torque limits, the range was much larger, from 0.02 to 0.15, indicating larger potential improvements in grasp posture cost for handle reorientation. The change in cost range was the opposite for the pan scene, where the original torque limits correspond to a wider range of 0.02 to 0.3 of the lowest 10% cost values compared to the cost range of 0 to 0.01 for the altered torque limits. This difference in results shows that the utility of object rotation is not predictable from the manipulator capabilities alone but also depends on the specific grasping task and environment.

Fig. 9.2 shows that, with the original torque limits, the distal joints were most frequently the limiting joint in the payload safety margin cost. For both task scenarios, the altered torque limits shifted the limiting joint to the proximal joints. Because the change in the cost distribution was not the same for both cases, our results provide a counter-example to our hypothesis that weaker distal joints would necessarily correspond to increased cost range over different initial object orientations.

Table 9.1: Original and altered manipulator torque limits in comparison of grasping posture cost sensitivity. In each case, the individual torque limits are symmetric for positive and negative moments.

Individual joint torque limits [N-m]	Joint number (proximal to distal)						
	1	2	3	4	5	6	7
Original	75.6	50.4	50.4	28.8	6.01	6.01	0.6
Altered	30.0	30.0	30.0	30.0	30.0	30.0	30.0

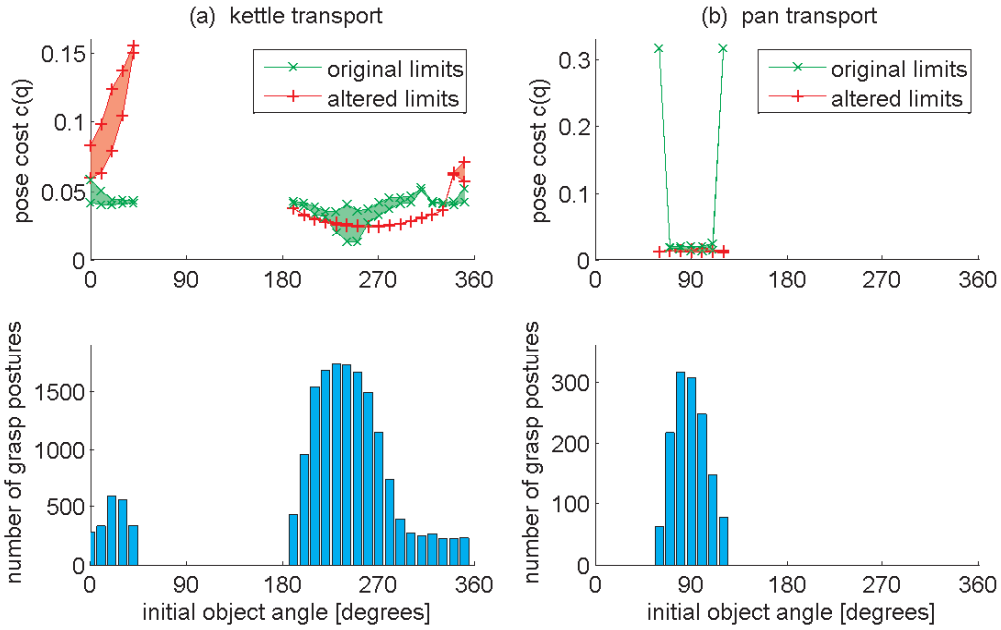

Figure 9.1: Payload safety margin cost for candidate grasp postures within the threshold percentile, i.e.the lowest 10% cost values per sampled inital object angle. The top row shows the minimum and maximum pose costs for candidate grasp postures over different object orientation angles. The blank regions correspond to object orientations for which there are no feasible grasp postures. The bottom row shows the total number of feasible candidate grasping postures over the object orientation angles. The number of grasp postures is the same regardless of the torque limits used in the evaluation.

Our examples suggest that the payload cost value distribution will be specific to the interaction between the manipulator torque limits, object grasp set, and object position relative to the robot. Thus, the sampling stage to evaluate candidate grasp postures is useful not only for selecting the optimal target orientation for rotation but also for predicting whether the gains over a direct grasp are a sufficient trade-off for the cost of computing a pre-grasp rotation maneuver.

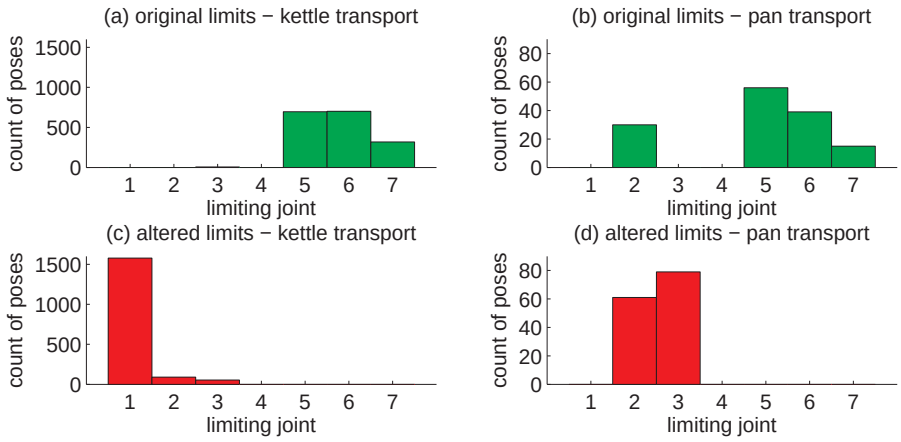

Figure 9.2: Distribution for the limiting joint for the set of selected grasp postures. The counts are shown only for the set of grasping postures that meet the threshold percentile for each set of payload limits.

9.4 Cost functions for grasp-posture selection

In the second set of experiments, we compare the optimization of an alternative torque-based cost function to the the payload safety margin metric. The analysis of sampled grasping postures in the previous sections indicated that the limiting joint for the payload metric was dominated by a few particular joint axes, rather than being uniformly distributed across the 7 joint axes of the example manipulator. Since the payload margin is determined by the limiting joint, an alternative cost function for evaluating the grasping postures is the torque magnitude of the predicted weakest joint. In the next experiments presented in this section, optimizing the positioning of the weakest or limiting joint did result in similar target object angles for the grasping postures.

Besides providing a simpler mechanistic alternative to the payload margin cost, the method in this section also presents a way to refine a feasible grasping posture with respect to a cost function. Local refinement is useful for exploration and optimization in the continuous configuration space of the grasping posture between the discretized sampling points.

We first describe the refinement method for optimizing a cost metric while maintaining pose constraints. Then we present the analysis results comparing optimization of the payload safety margin to minimization of a single joint torque magnitude.

9.4.1 Gradient-based optimization of joint torque

The configuration space for representing a specified grasp posture is $N+1$ dimensional for an N-degree-of-freedom (DoF) manipulator and a single pivot freedom for object rotation. Note that this count is for a single grasp posture only, defined by the

desired relative transform between the manipulator end effector and the object. The search space will have higher dimension when there are multiple allowable grasps or grasp freedoms such as the orientation or position of the end effector along the object handle.

We used a gradient-based search in the configuration space to optimize the grasping posture with respect to a particular cost metric. The grasping postures must satisfy end-effector pose constraints that achieve the desired grasp relative to the object configuration in the scene, and the object's orientation in the horizontal plane is variable due to possible pre-grasp rotation. The problem can be recast as refining the posture of an $(N+1)$-DoF manipulator whose most distal joint is the passive object rotation freedom and whose end-effector position is located on the object pivot axis. The recast end-effector constraint is the 6-DoF pose of a virtual frame attached to the passive object rotation joint. Thus the objective is to compute an optimal $N+1$ configuration with respect to the 6-DoF pose constraint on the virtual link.

A general method for a gradient-based optimization satisfying the end-effector constraint is to travel in the null space of the configuration constraint. For the $(N = 7)$-DoF manipulator described in the previous section, the full configuration space is 8-dimensional, and the null space with respect to the 6-DoF constraint is 2-dimensional. In practice we found that projecting gradient steps onto the local null-space of the current configuration Jacobian resulted in drift away from the pose constraint manifold that was difficult to mitigate in an automatic manner.

Instead, we used a gradient search in a 2-DoF subspace of the full configuration space in order to maintain the end-effector pose constraint without drift complications. We selected, for convenience, the first joint axis of the manipulator and the passive object rotation axis as the 2 subspace freedoms. For each given point in the 2-D subspace, we solved for the remaining 6 joint angle values which satisfy the 6-DoF end-effector constraint. The inverse kinematics was solved using an analytical solution that was specific to the WAM robot and that made use of the decomposition between orientation and position constraints to solve for the joint angle values.

Below are the algorithm steps for the cost optimization using the subspace gradient search:

1. Start with an initial point in the 2-DoF subspace and a corresponding full 8-DoF reference configuration.
2. Numerically compute the cost metric gradient in the 2-DoF subspace.
   (a) For a step in each direction in the subspace, compute the corresponding full 8-DoF configuration. If there are multiple analytic solutions, discard any which violate kinematic joint limits, and use the distance to the reference configuration to select from among any remaining alternatives.
   (b) Evaluate the cost metric for the full 8-DoF configuration.
   (c) Compute the gradient from a finite differencing scheme in the current subspace dimension.

3. Take a step opposite the gradient in the 2-DoF subspace (for minimization instead of maximization of the relevant cost function).
4. Resolve the full 8-DoF configuration for the new point using an analytical solution given the 2-D subspace point.
5. Repeat from step 2 if none of the following termination criteria are satisfied:
   (a) No full 8-DoF configuration can be found within the kinematic joint limits.
   (b) The cost value of the new point remains stationary within some threshold or has increased with respect to the last cost.
   (c) The maximum iteration or computational time limit has been exceeded.

Our implementation of the subspace gradient search in MATLAB did not include collision checking for the manipulator pose against environment obstacles. However, collision checking could be included as an additional collision criterion.

We used the described subspace gradient search to compare the optimization of two types of torque-based cost metrics for the grasping posture. The first is the payload safety margin cost, which depends on the torque limits of the manipulator's multiple joints. The second is a torque magnitude of a single selected joint expected to be the weakest joint.

9.4.2 Comparison of payload margin and single joint torque metrics

We tested three task scenarios: the kettle transport task and the two cooking pan transport tasks described previously in Fig. 8.3. In each case, a single grasp was selected from the original grasp set described in §8.5. The selected grasp was the canonical one where the end-effector axes of the hand/palm are aligned with the principal axes of the object frame. Thus the presented results exclude grasps where there is a roll freedom around the object handle axis.

For each scenario, we used the precomputed grasping postures sampled by the method described in §8.5 as initialization points for the optimization. Overall there were 331, 810, and 579 initial points, respectively, for the kettle grasps at the table, pan grasps at the table, and pan grasps at the counter. We used the subspace gradient search method to optimize the payload margin cost and the single joint torque cost for joints 2, 5, and 6, which were most often the limiting joint.

Figures 9.3, 9.4, and 9.5 show the change in the optimization cost function from the initialization point to the final posture after optimization. In the sampled initialization postures, there were multiple grasping postures per sampled object orientation angle. For clarity, the plots show only the final optimized posture with the lowest cost for each group of postures with the same initialization object angle.

For the kettle grasps (Fig. 9.3), the limiting joint for the payload safety margin cost was most often joint 6, the wrist flexion, and otherwise joint 5, the forearm

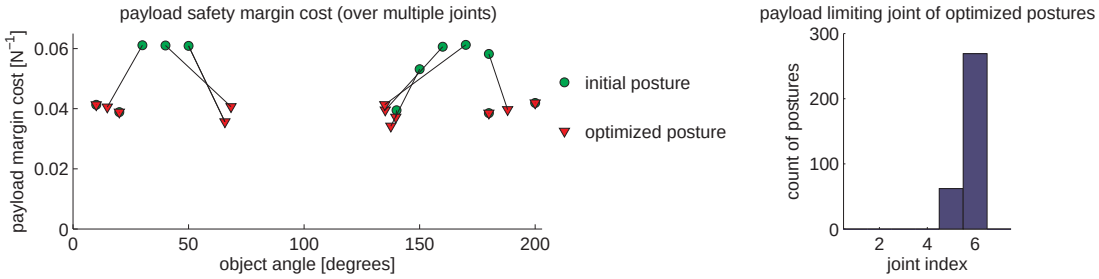

(a) Kettle on table: Payload safety margin optimization

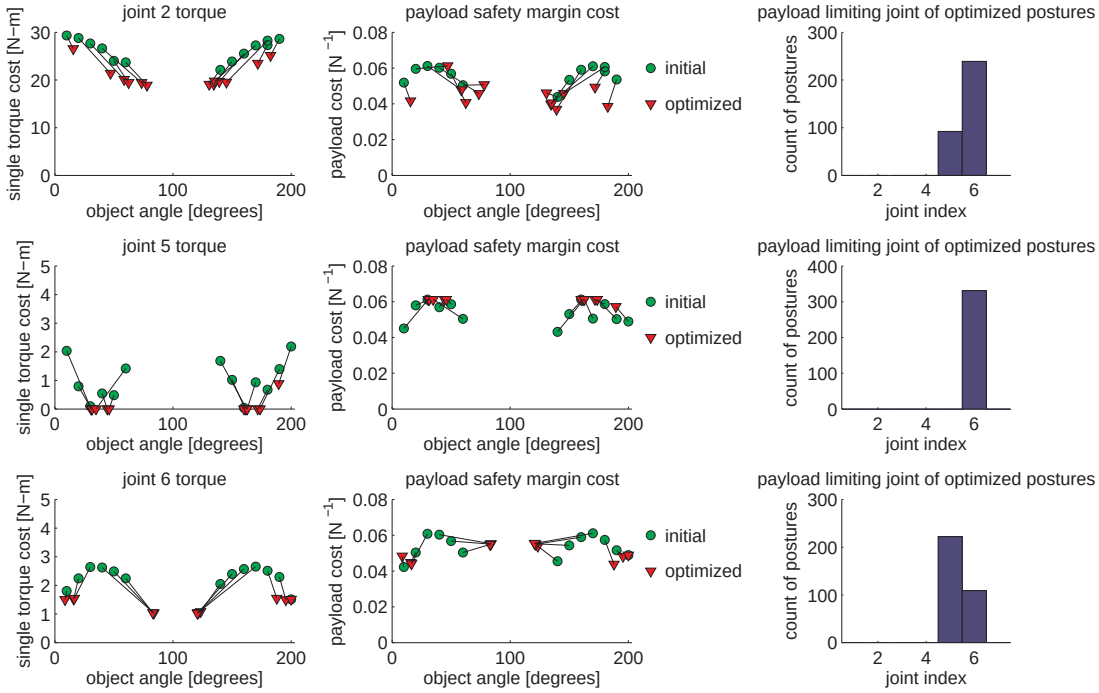

(b) Kettle on table: Single joint torque minimization

Figure 9.3: Optimization results for grasping the kettle from a table. (a) The change in the cost versus the passive object rotation freedom for optimizing the payload safety margin (left). For clarity, only the results with the lowest final cost from each group of postures with the same initial object angle are shown. Joint 6 was most often the limiting joint (right) for the entire set of final optimized postures for grasping the kettle. (b) The change in the joint torque cost versus object angle (left) when the optimization cost function is the torque magnitude at a single joint, for joints 2, 5, and 6. The corresponding changes in the payload margin cost are shown (center) for the postures resulting from single joint torque minimization. The limiting joint for the payload cost is shown (right) for the entire set of final postures resulting from single joint torque minimization.

roll or pronation. (Note the distribution of the payload-limiting joint in (Fig. 9.3) is different from the distribution shown in (Fig. 9.1) due to the restriction to a

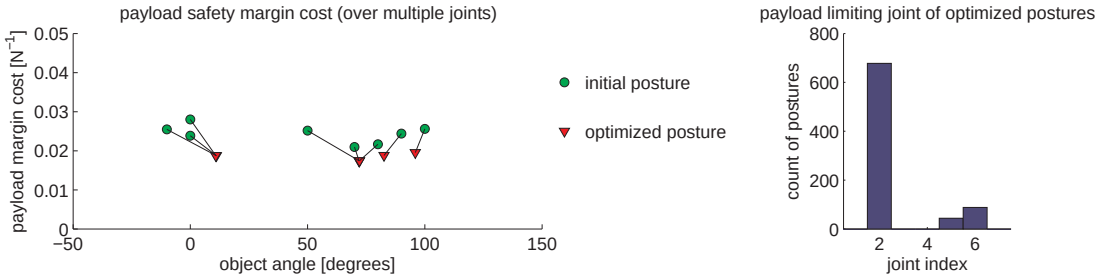

(a) Pan on table: Payload safety margin optimization

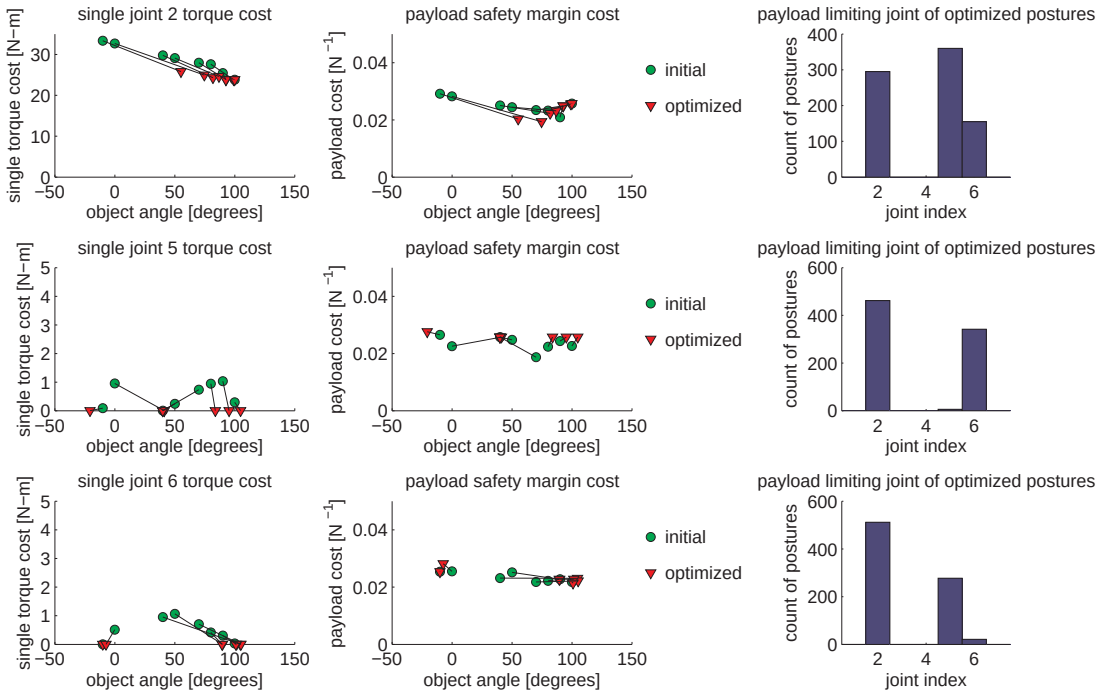

(b) Pan on table: Single joint torque minimization

Figure 9.4: Optimization results for grasping the cooking pan from a table. (a) The change in the cost versus the passive object rotation freedom for optimizing the payload safety margin (left). Only the lowest-cost result from each group of postures with the same initial object angle are shown. Joint 2 was most often the limiting joint (right) for the entire set of final optimized postures for grasping the pan. (b) The change in the joint torque cost versus object angle (left) when the optimization cost function is the torque magnitude at a single joint, for joints 2, 5, and 6. The corresponding changes in the payload margin cost are shown (center) for the postures resulting from single joint torque minimization. The limiting joint for the payload cost is shown (right) for the entire set of final postures resulting from single joint torque minimization.

single grasp within the grasp set.) Thus, joint 6 is considered the weakest joint for this task. When the single joint 6 torque cost was used as the cost function for

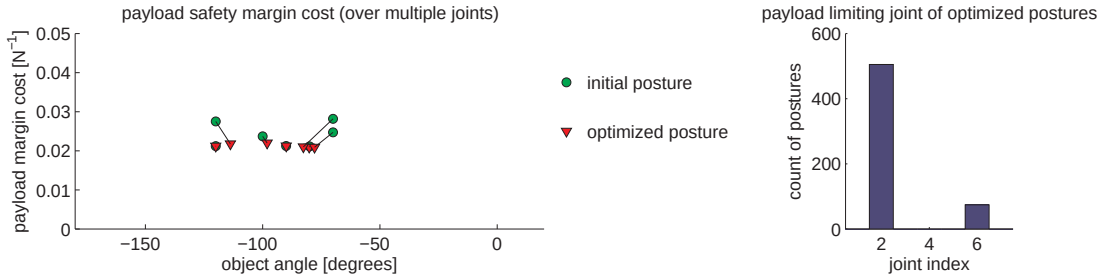

(a) Pan on counter: Payload safety margin optimization

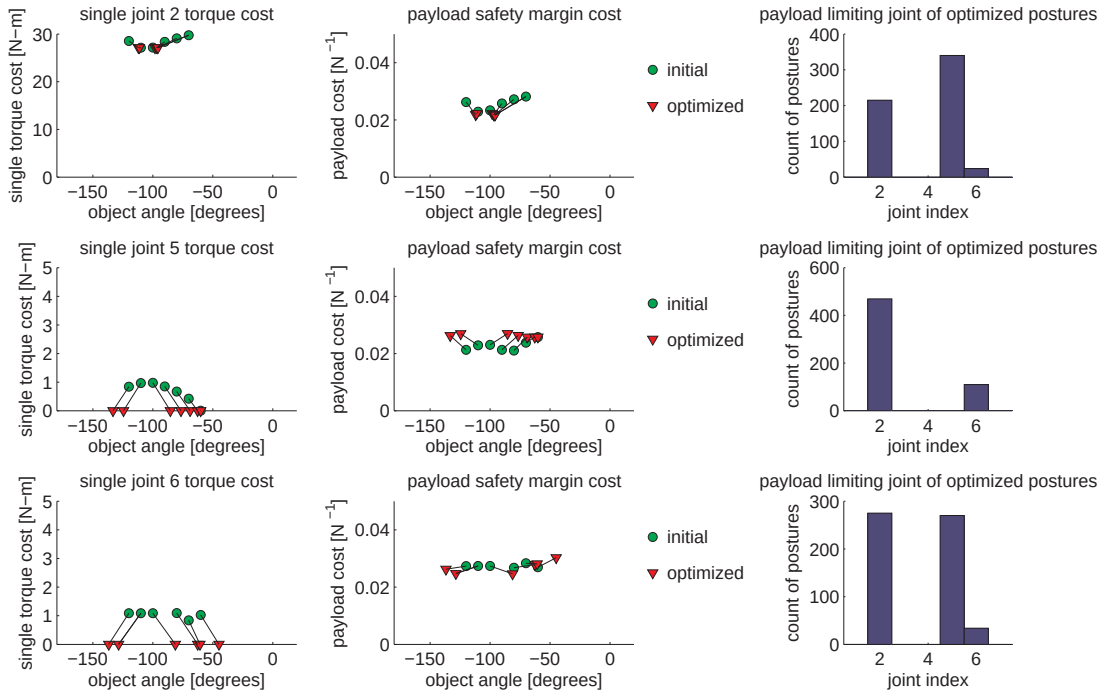

(b) Pan on counter: Single joint torque minimization

Figure 9.5: Optimization results for grasping the cooking pan from a counter. (a) The change in the cost versus the passive object rotation freedom for optimizing the payload safety margin (left). For clarity, only the results with the lowest final cost from each group of postures with the same initial object angle are shown. Joint 2 was most often the limiting joint (right) for the entire set of final optimized postures for grasping the pan. (b) The change in the joint torque cost versus object angle (left) when the optimization cost function is the torque magnitude at a single joint, for joints 2, 5, and 6. The corresponding changes in the payload margin cost are shown (center) for the postures resulting from single joint torque minimization. The limiting joint for the payload cost is shown (right) for the entire set of final postures resulting from single joint torque minimization.

optimizing the grasp postures, the resulting object angles and postures were similar to those resulting from the payload metric optimization (Fig. 9.6, Fig. 9.7).

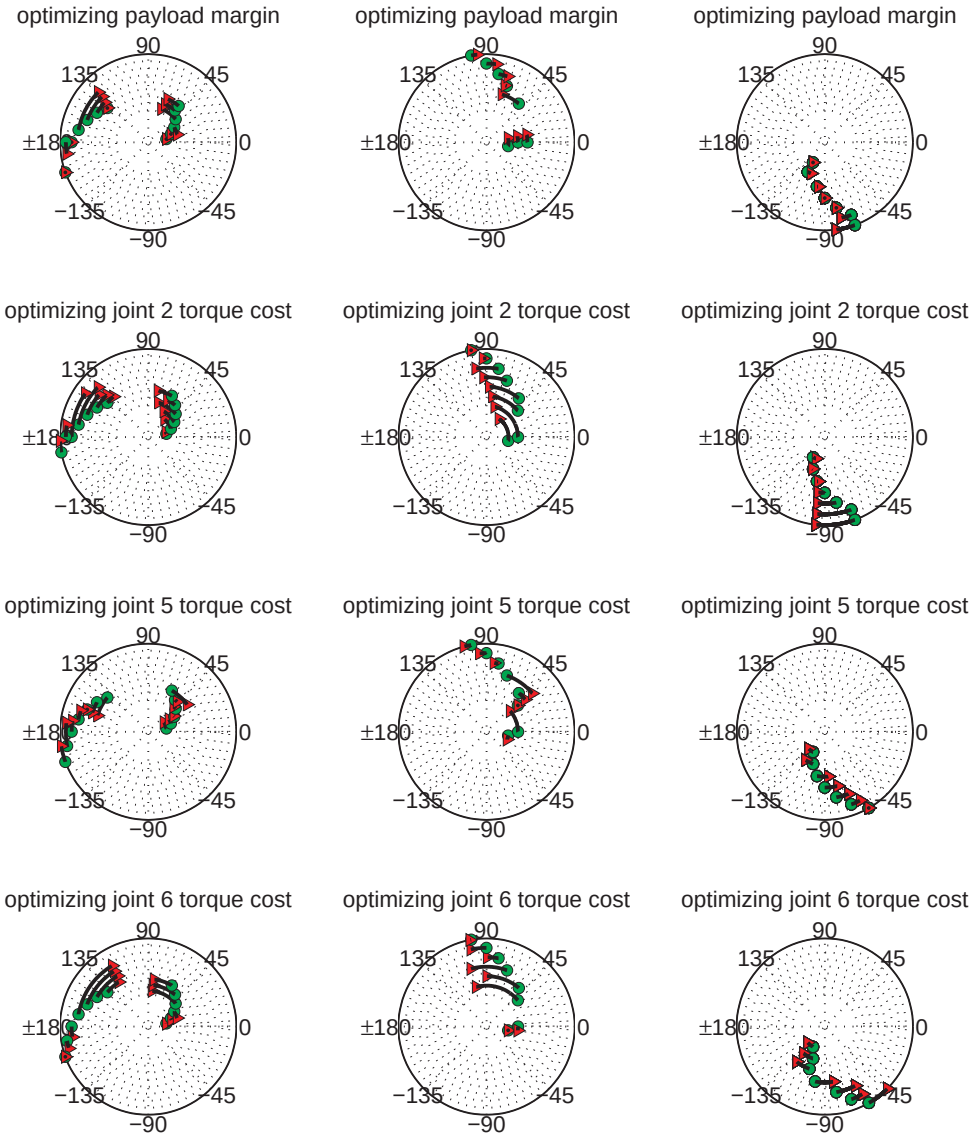

(a) Kettle on table (b) Pan on table (c) Pan on counter

Figure 9.6: Object angles corresponding to joint torque optimization results. The angular orientation indicates the object handle direction. Examples at larger radii have lower final costs for the payload margin cost, in the first row, or the single joint torque costs, in the bottom three rows. Green circles and red triangles indicate the object angle of the initialization and optimization postures, respectively. The direction to the robot base is at approximately (a) 90 degrees in the plot coordinates for the kettle grasps, (b) 90 degrees for the pan grasps at the table, and (c) −90 degrees for the pan grasps at the counter.

In contrast, while it was possible to achieve arm configurations with zero joint 5 torque, the resulting postures had higher payload margin costs than the original

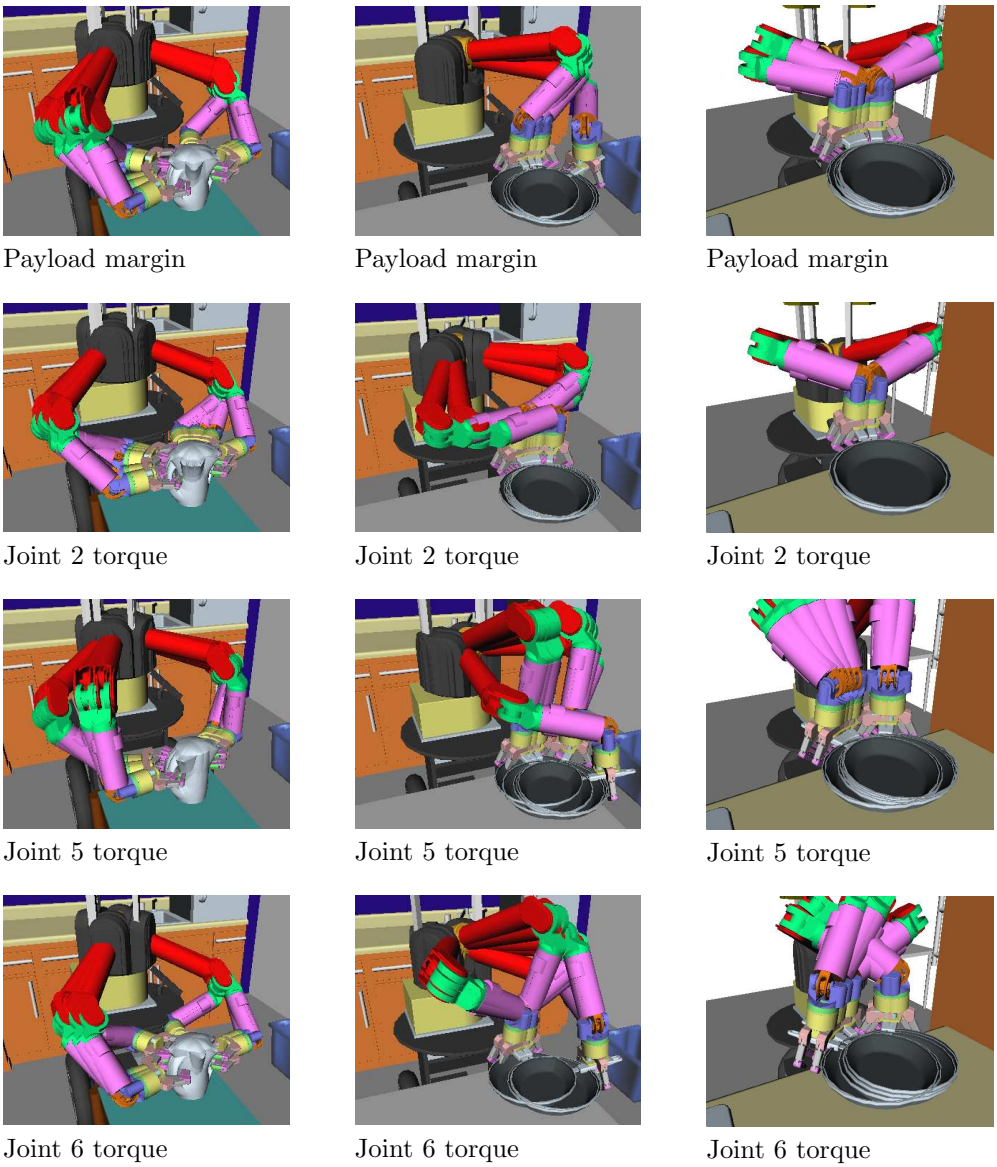

(a) Kettle on table (b) Pan on table (c) Pan on counter

Figure 9.7: Optimized grasping postures for different torque-based cost functions. The illustrated robot grasps correspond to the final optimized postures whose object angles are shown in Fig. 9.6. (a) For the kettle grasps, the object angles from payload margin optimization are most similar to those for minimizing the joint 6 torque. (See Fig. 9.3). (b, c) For the pan grasps, the object angles from payload margin optimization are most similar to those for minimizing the joint 2 torque. (See Fig. 9.4 and Fig. 9.5).

initialization postures. Optimizing the joint 2 shoulder elevation torque, however, yielded the lowest corresponding payload costs (see middle column of Fig. 9.3).

This result demonstrates that optimizing a single joint torque independently is not sufficient to replicate the payload margin optimization, but it can provide a simple alternative cost funciton that results in similar object angles.

The cost function comparison yielded similar results for the two pan grasping tasks in Fig. 9.4 and Fig. 9.5. The main difference was that for the pan grasping postures, the limiting joint for the payload metric optimization was most often joint 2, shoulder elevation, instead of joint 6, wrist flexion (Fig. 9.4 and Fig. 9.5). Thus the comparison of the single joint torque optimizations shows that minimizing the joint 2 torque magnitude yielded more similar final object angles compared to minimizing joint 5 or joint 6 torques (Fig. 9.6; Fig. 9.7). This difference indicates that the limiting joint is not necessarily easily predicted by the individual torque limits (Table 9.1) and also may not be consistent for the same manipulator for different objects. Thus, to substitute a single joint torque cost for the multi-joint payload metric cost, evaluation of an initial sample of feasible grasping postures is necessary to estimate which joint is the weakest in a given scene context.

The lack of convergence to a single minimum cost region (Figures 9.3, 9.4, and 9.5) illustrates the local nature of the optimization method and thus the requirement to consider multiple initialization postures.

9.5 Optimization of transport path quality

In this section, we present and evaluate extensions for further optimization of the transport path quality beyond the grasp posture selection. We evaluate two variations on the planning method presented in §8, path segment modification and multiple plan re-initializations.

In general, task feasibility is a priority over task quality because we do not wish to sacrifice basic task completion for want of an optimal solution. Thus, we examine methods for improving task quality that can be integrated into the pipeline for finding feasible paths in high-dimensional configuration spaces with obstacles.

We consider three aspects of the sampling-based planning process in our experiments (Table 9.2). These "control parameters" that influence the path quality are:

1. **Initial point selection**: For the transport path, the grasp posture at object lift-off is the initial point or configuration of the path. In the previous chapter we referred to the grasp posture as the "transition point" between the grasp acquisition and the load-bearing transport action. This control parameter could be selected at random from feasible postures in a naïve method. The initial method presented in §8 investigated how restricting the allowable grasp postures could indirectly affect the quality of the following path.
2. **Path segment modifications**: Given a feasible path with fixed end points, local modifications can be made to internal path segments. We examine changes at the path shortcutting stage, where the initial feasible path is progressively

Table 9.2: Extensions for transport path quality optimization. We evaluate two variations on the planning method presented in §8, namely path segment modification and multiple plan sampling. The extensions, shown in the last two columns, build upon the grasp posture optimization rather than testing the additional control parameters in isolation.

Control parameter	Planning method variation			
	Naïve planning method	Grasp posture selection	Path segment modification	Multiple plan sampling
Initial point selection	Random	Restricted	Restricted	Restricted
Path shortcutting metric	Agnostic	Agnostic	Consistent	Consistent
Search re-initialization	Single	Single	Single	Multiple

modified. Often the raw output from a sampling-based method such as an RRT planner is a circuitous path that is then refined by searching for shortcut path segments in the configuration space. We examine whether using a shortcut cost metric that is consistent with the desired optimization goal improves the path quality, compared to using c-space distance as a default metric.

3. **Search re-initialization**: There are potentially multiple paths between different sets of end-point configurations. This redundancy offers the opportunity for high level selection from among multiple paths. Instead of accepting the first feasible path, we can evaluate multiple plan attempts and select the lowest cost path from the available feasible plans.

The effects of the last two control parameters, path segment modification and multiple plan initializations, are evaluated as additions to the first control parameter of grasping posture selection. Table 9.2 shows the progression of these extensions that we add to the method described earlier. The evaluation of initial point selection, represented by the comparison between the first two columns in Table 9.2, was the focus of the previous chapter §8. This section investigates the methods in the last two columns, representing the addition of local path modification and selection from among multiple paths.

9.5.1 Experiment task scenarios

The manipulation tasks for the simulation experiments in this chapter are similar to the examples in the previous chapter's experiments (§8.5). The manipulator was the 7-DoF right arm of a bimanual platform with two Barrett WAM robots (Barrett Technology Inc., Cambridge, MA) (Fig. 9.8). The robot base location was fixed in the world frame per example scene. The pre-grasp rotation strategy was compared to direct grasping for two transport tasks, one involving the cooking pan object and

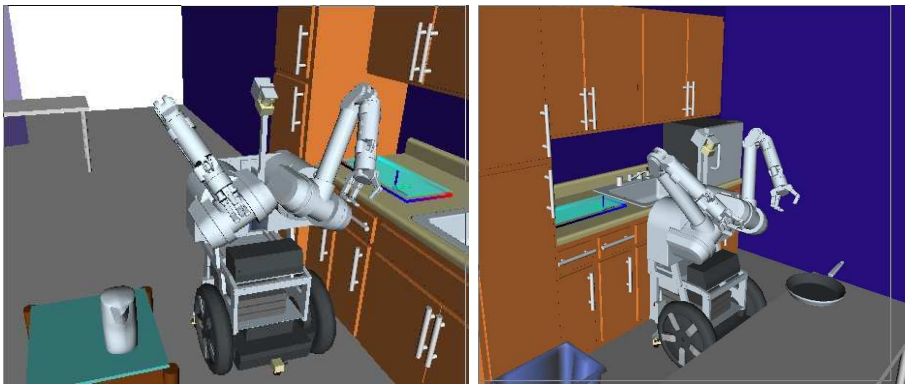

(a) (b)

Figure 9.8: Tasks scenarios for path quality optimization experiments. (a) The kettle should be transported from the table to the countertop. (b) The cooking pan should be transported from the table to the countertop. In both scenarios, only the right arm of the bimanual manipulator is considered available for movement.

one involving the kettle object. Fig. 9.8 shows the example scenes with the objects at the initial location and with the transport goal regions highlighted.

As before in §8.5, in the grasp posture computation stage, the object orientation θ was sampled at 10-degree intervals for 36 possible object poses. For each separate grasp region defined for the object, 3 end-effector poses were sampled within the bounds of the 6-D pose region. The grasp postures were restricted to the best 10% configurations by setting the percentile threshold to $t = 10$. The tolerance for upright object orientation was ± 10 degrees for the pan and ± 6 degrees for the kettle.

The pre-grasp rotation plan was attempted using the discretized depth-first search of the caging grasp planner by Diankov *et al.* [2008b]. In this chapter, both pre-grasp interaction and direct grasping plans were restricted to transition point grasp postures from the threshold set associated with the object orientation angle. For both methods, if no grasp acquisition plan is found, then no transport plan was 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 [Berenson *et al.*, 2009a].

9.5.2 Path shortcutting with a consistent cost metric

Standard path shortcutting methods accept a local path segment modification based on the path distance in the configuration space. Although the path distance between points are weights in the transport path payload cost in Eq. 8.6, a purely distance-based cost is not consistent with the optimization objective of payload safety margin.

One method for incorporating a consistent objective metric more directly into the

path selection is to use the desired objective as the evaluation metric during path shortcutting. Here we evaluate how using different shortcutting metrics could change the overall path quality when combined with the selection of low-cost transition (or initial transport path) points (§8.4).

In this set of experiments, we used the same grasp posture selection method described in §8.4 to decompose the overall manipulation task into the three action components. The only difference is in the final stage of refining the feasible paths first returned by the RRT search. Instead of shortcutting the original raw paths based on the configuration space distance, we consider two different shortcut cost metrics:

- workspace path length, and
- payload safety margin metric, weighted by configuration space distance.

Each of these cost-metrics are applied to the raw RRT path independently to yield two refined paths per feasible task. First, candidate short-cuts to the original path are accepted only if they result in lower workspace path length, measured from the end-effector position. Second, candidate short-cuts to the original path are accepted only if they result in decreased payload path cost given by Eq. 8.6. In each case, the shortcut paths respect the end-effector pose constraints for keeping the object upright during the transport action (using the method developed by [Berenson *et al.*, 2009a]).

Each of the resulting smoothed paths is then evaluated according to the payload cost metric, Eq. 8.6. Thus shortcutting by the workspace distance cost and evaluating by the payload cost is considered the agnostic or inconsistent path modification method. Shortcutting by payload cost is the consistent path modification method.

For both example task scenarios, we evaluated three methods: (1) direct grasping with unrestricted grasp postures (the naïve method), (2) direct grasping with restricted grasping postures, and (3) pre-grasp rotation with restricted grasping postures. In each case, transport plans were attempted for 36 initial object orientations θ evenly spaced at 10-degree intervals. Thus, there was a maximum of 36 successful plans for each of the three methods.

Table 9.3 shows the number of feasible plans for each experiment set. For both tasks, pre-grasp rotation with restricted grasping postures yielded the highest task success rate. There was marked difference between the scenes for the comparison between random grasp selection and restricted grasp postures for direct grasping. For the pan transport task, we hypothesize that the decrease in task success rate associated with restricting grasp postures is due to finger collision with the table surface for the pan grasps. Because the restricted grasp postures were sampled from three reference grasps, there could be dramatically fewer collision-free grasp postures if the hand itself was in collision for any of the reference grasps. The larger set of grasp postures from the full continuous grasp region may have been an advantage in this case.

For each feasible transport plan, the original path was refined by the two shortcut

Table 9.3: Task feasibility for the two example scenes. The total number of successful transport plans is listed out for the set of 36 task input conditions. The task feasibility is based on the original unrefined search output and is independent of the shortcut cost metric.

Scene	Total feasible plans (out of 36)		
	Direct grasping, unrestricted postures	Direct grasping, restricted postures	Pre-grasp rotation, restricted postures
Kettle transport	8	15	17
Pan transport	14	4	16

metrics. Fig. 9.9 shows a comparison of the two path refinements in the workspace. In each example, the two workspace paths were similar, as were the configuration space paths that shared identical end points. This similarity suggests that path refinements based on either of the tested shortcut metrics converged to similar regions. This similarity may be due to the dependency, albeit nonlinear, of workspace length on configuration space length that only vanishes for null-space movements. In addition, candidate short-cuts were only feasible if they satisfied the end-effector pose constraints, which further limited the set of allowable path-segment modifications.

Accordingly, the comparison of the final payload costs shows that there is little difference between path modifications using the two shortcut cost metrics (Fig. 9.10). This similarity suggests that the variation in the payload safety margin may contribute less to the overall transport cost than variation in path distance.

Comparison of the transport path costs between methods in Fig. 9.10 (i.e. between plots in the same row) shows similar results as the previous experiments illustrated in Fig. 8.5. For the kettle object, the distribution of transport path costs was similar for both direct grasping and pre-grasp rotation. In contrast, pre-grasp rotation had much lower transport path costs for the pan object. This result reproduces the previous observations in §8 for a different manipulator structure and different robot base orientation. It suggests that the difference could be due to constraints on the distal rather than proximal joints. It is also consistent with our earlier observation that the limiting joints for the payload metric were often the distal joints, such that the object handle axis was more influential than manipulator base orientation.

9.5.3 High-level optimization over multiple path choices

In this section we consider optimization of the payload cost at the level of entire paths (with different end points) rather than only path segment modifications (where the end points stay fixed). Additional computation time is spent to search for multiple feasible paths instead of limiting the plan solution to the first attempt. In these experiments, the focus is on whether multiple plan attempts could improve the trans-

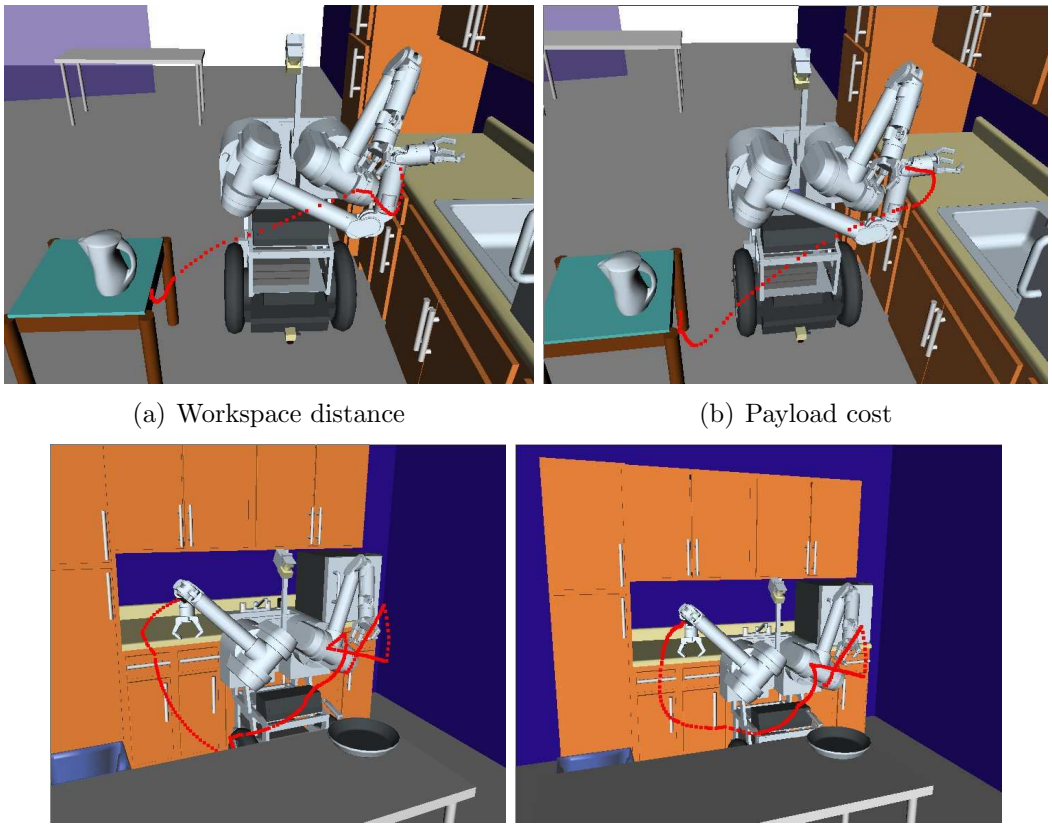

(a) Workspace distance (b) Payload cost

(c) Workspace distance (d) Payload cost

Figure 9.9: For the same reference path, the path refinement by the two cost metrics produces similar paths. (a) Kettle transport path shortcut by workspace distance. (b) Kettle transport path shortcut by payload cost. (c) Pan transport path shortcut by workspace distance. (d) Pan transport path shortcut by payload cost.

port path cost and change the trade-off between the direct grasping and pre-grasp rotation strategies. Recall that the transition sampling computation in the reference method from §8 determined not only a single but multiple low-cost grasping postures for a given target object orientation θ. Thus transport plans were not limited to a single grasp posture at the transition point after grasp acquisition. Likewise, multiple end configurations were possible within the transport goal region requirements for the object configuration.

In this section, we consider multiple grasp plan attempts as well as multiple transport plan attempts. Fig. 9.11 illustrates the structure of the overall planning method.

For the direct grasping strategy, there were $n_g = 5$ attempts to grasp the object from the presented configuration, and for each successful grasp plan there were $n_t = 5$ attempts to find a transport plan to the goal region. The target transition points

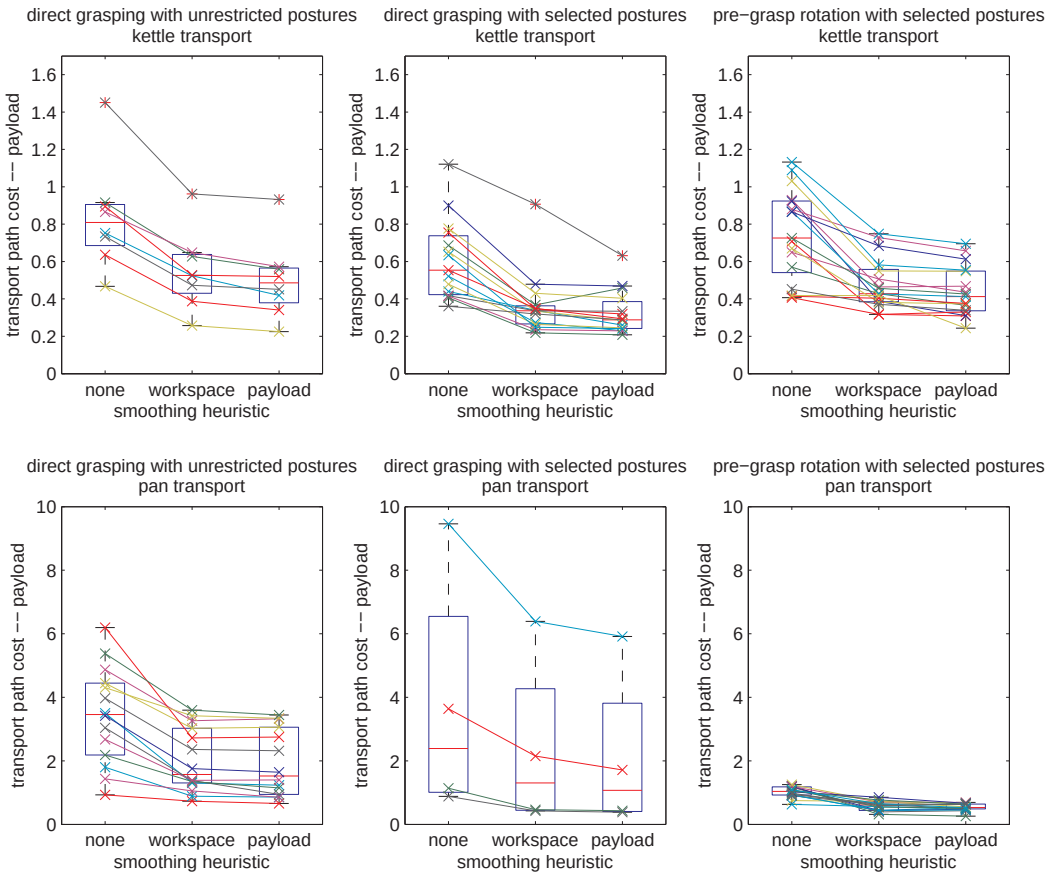

Figure 9.10: Effect of optimization during shortcutting. Transport path costs are shown for the original unrefined path and the two refined paths based on the workspace and payload metrics. The connecting lines within each plot indicate the costs for the paths modified from the same original path.

for the grasping postures were restricted to the top poses determined from the pre-computation sampling stage. (Note that this restriction differs from the method in the previous chapter, where the direct grasping target postures were unlimited within the allowable grasp regions.) Similarly, for the pre-grasp rotation strategy, there were $n_g = 5$ grasp plan attempts for the rotated object position and $n_t = 5$ transport plan attempts per successful grasp plan. The object configuration at the grasping transition point was always the optimal target orientation θ_t determined from the pre-sampling phase. That is, there was only a single pivoting action plan considered per initial object orientation, but there were multiple grasp plans and multiple transport plans as choices for the remainder of the task.

For each scenario, the multiple path re-initialization method was tested for 36 initial object orientations θ evenly spaced at 10-degree intervals. All initial feasible paths were smoothed according to the two cost metrics as in the previous section.

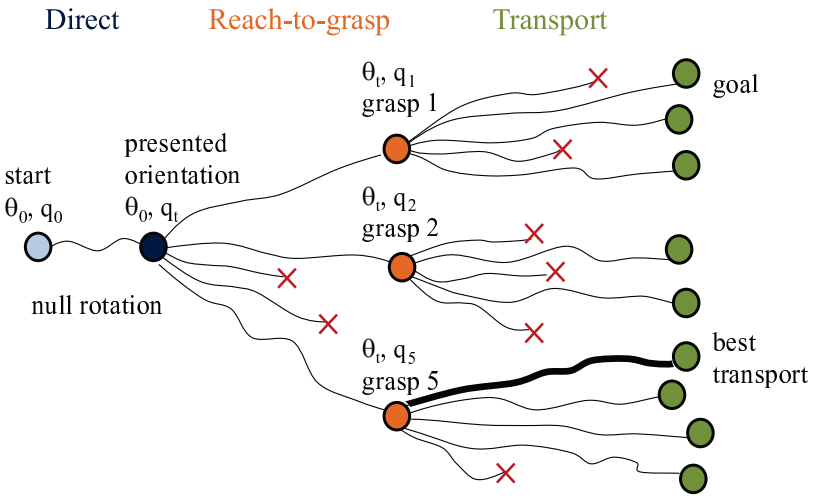

(a) Direct grasping

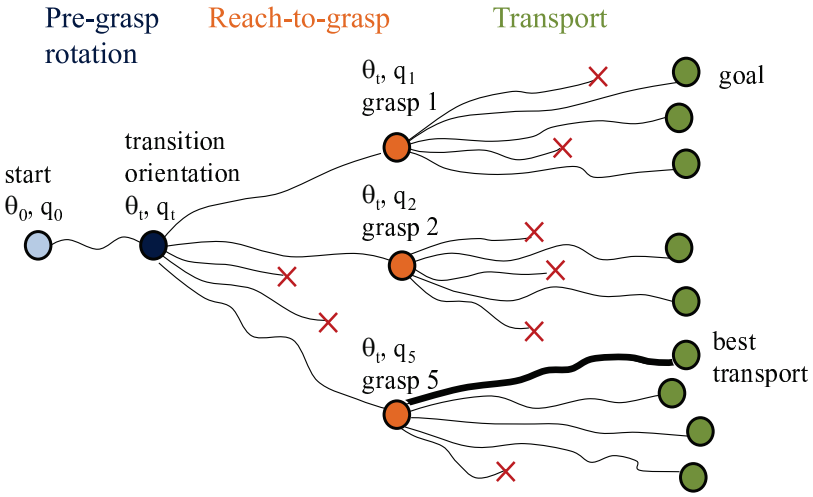

(b) Pre-grasp rotation

Figure 9.11: Planning with multiple re-initializations. In these diagrams there are $n_g = 5$ grasp attempts and $n_t = 5$ transport attempts. (a) Direct grasping. (b) Pre-grasp rotation. Note that for pre-grasp rotation there is only a single rotation plan attempt to the target object orientation angle θ_t. The only differences between the direct grasping method and the pre-grasp rotation method are the amount of object rotation from the starting object orientation θ_0 to the transition orientation θ_t and the manipulator configuration $\mathbf{q}_t$ at the transition.

All paths were evaluated according to the payload cost metric as before. The final selected plan was that with the lowest payload cost for the transport path.

The largest difference of the multiple attempt optimization in Fig. 9.12 from the single plan attempt method (Fig. 9.10) was the decreased maximum transport cost for each case. The comparison between the shortcut cost metrics (within each

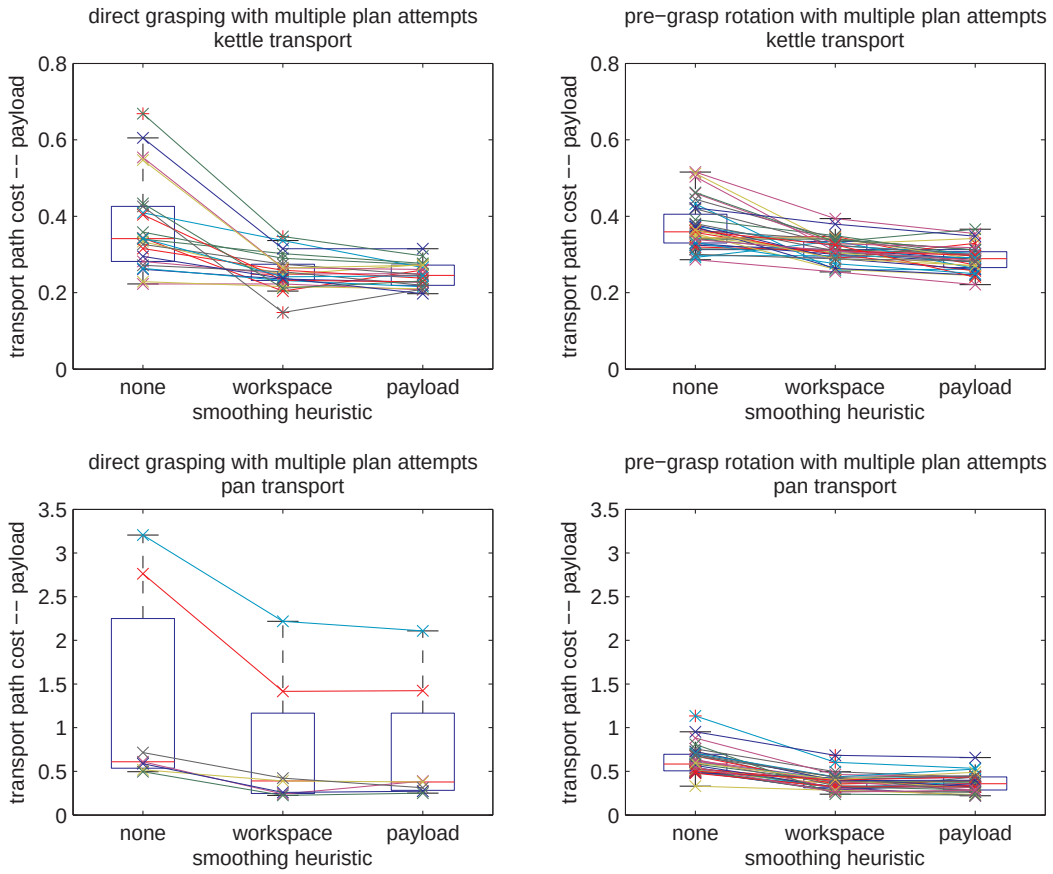

Figure 9.12: Effect of multiple plan re-initializations. Transport path costs are shown for the best unrefined path out of multiple paths, the best refined path after shortcutting with respect to the workspace distance, and the best refined path after shortcutting with respect to the payload margin metric. The connecting lines within each plot indicate the costs for the paths corresponding to the same task scenario with a particular initial object orientation.

plot of Fig. 9.12) reproduces the previous results, where the two refined paths have nearly equivalent transport costs. This similarity suggests that the variety of grasp postures as starting points for the transport paths is more influential than internal path segment modification. This result was partially expected due to the fact that path segment modifications for a fixed pair of end points represent a subset of the paths achievable from multiple path attempts. However, it demonstrates, for these example scenes, that the payload cost metric varies less for refinements of a single feasible plan than between plans with different end points.

The improvement in the path quality for the direct grasping method indicates that advantage of pre-grasp rotation is mitigated when multiple plan initializations are considered. Thus, if the computation required for re-planning a transport path

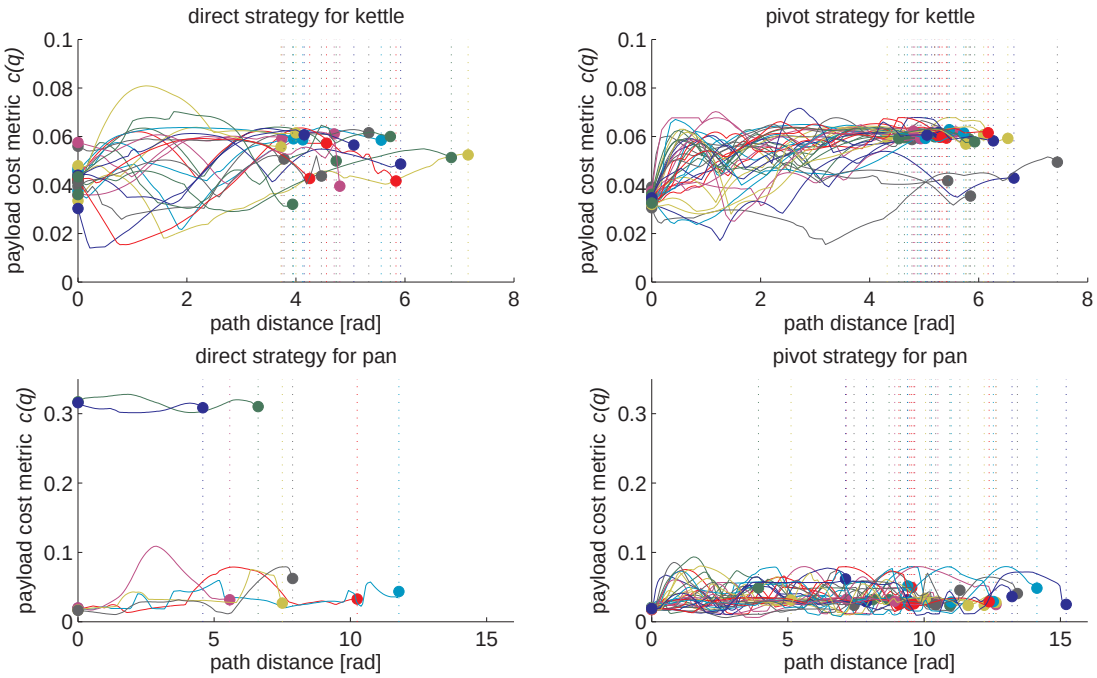

Figure 9.13: Cost evolution for the lowest cost path per initial angle. All paths were refined according to the payload metric. The vertical lines indicate the total path length in configuration space.

from direct grasping is favorable to that required for planning pre-grasp rotation, it may be preferable to choose a direct grasping strategy when it is feasible.

We also investigated the distribution of the transport paths according to the payload cost metric over the path length. Fig. 9.13 shows the cost evolution of the selected transport paths from the multiple plan attempts. While the initial configuration costs were more restricted for pre-grasp rotation, the cost variation per path after the starting point were similar in magnitude.

Some of the paths with shortest configuration space path length occurred for direct grasping. This result could be an effect of the limitation of pre-grasp rotation to the optimized transition orientation of the object angle, which was the same regardless of the starting orientation. In this sense, direct grasping had greater variety in the initial path points, some of which may have resulted in shorter paths for the scene. This feature is the main limitation of evaluating candidate grasp postures by an absolute state cost without consideration of the relation to the goal region or initial manipulator conditions.

9.6 Whole path metrics for manipulation quality

The transition selection method discussed previously is only relevant to the manipulation quality if the cost metric in question is computed from the cost of the individual configurations or points in the path. There are several other potential cost metrics relevant to robotics applications that depend on the relation between points and/or the timing assigned to the path execution. In these cases, there does not exist a cost metric for evaluating a single point independent of the path. Instead, the paths must be evaluated in their entirety in order to compare path costs, and there is no meaningful concept of a single point's contribution to the total path cost.

In this section, we examine one such metric, the workspace jerk of the path trajectory. We first present a method for optimizing and evaluating the jerk for a path. Evaluation of jerk requires a timing assignment to the path nodes so that the jerk is evaluated for a specific trajectory. We evaluated an example task scenario with pre-grasp rotation to investigate whether pre-grasp rotation could result in a manipulation action plan with lower jerk when direct grasping is feasible.

9.6.1 Minimum jerk timing for path trajectory

The total jerk over a trajectory measures the smoothness of the action as the cumulative first derivative of acceleration. For a trajectory $p(t)$, we evaluate the integral of squared jerk over time as the path cost:

$$J = \int (\dddot{p}(t))^2 dt. \tag{9.1}$$

Changes in path direction over a short time interval will result in large contributions to the cumulative jerk. Thus we expect a minimum jerk trajectory to have lower speed at sharp path turns and higher speed at smooth path segments, as illustrated by the examples in Fig. 9.14.

The jerk metric requires time derivatives of the motion, and thus the geometric path information is insufficient. In this section we present our method for assigning an explicit timing to the path nodes to yield a trajectory. This trajectory can then be evaluated for its total jerk for comparison to the trajectories resulting from alternative paths.

There are infinite choices for the timing assignment to a geometric path. We wish to assign a timing that achieves the minimum jerk for a specified path, such that the subsequent comparison with alternative paths is as favorable as possible. One limitation is that increasing the total time for path completion will always result in lower jerk, leading to an unbounded optimization problem. One approach is to fix a constant total time for all trajectories to constrain the optimization. We first optimize jerk for a unit time duration to find the optimal relative timing between path nodes. We then scale the unit duration timing to normalize the trajectory with respect to the maximum speed.

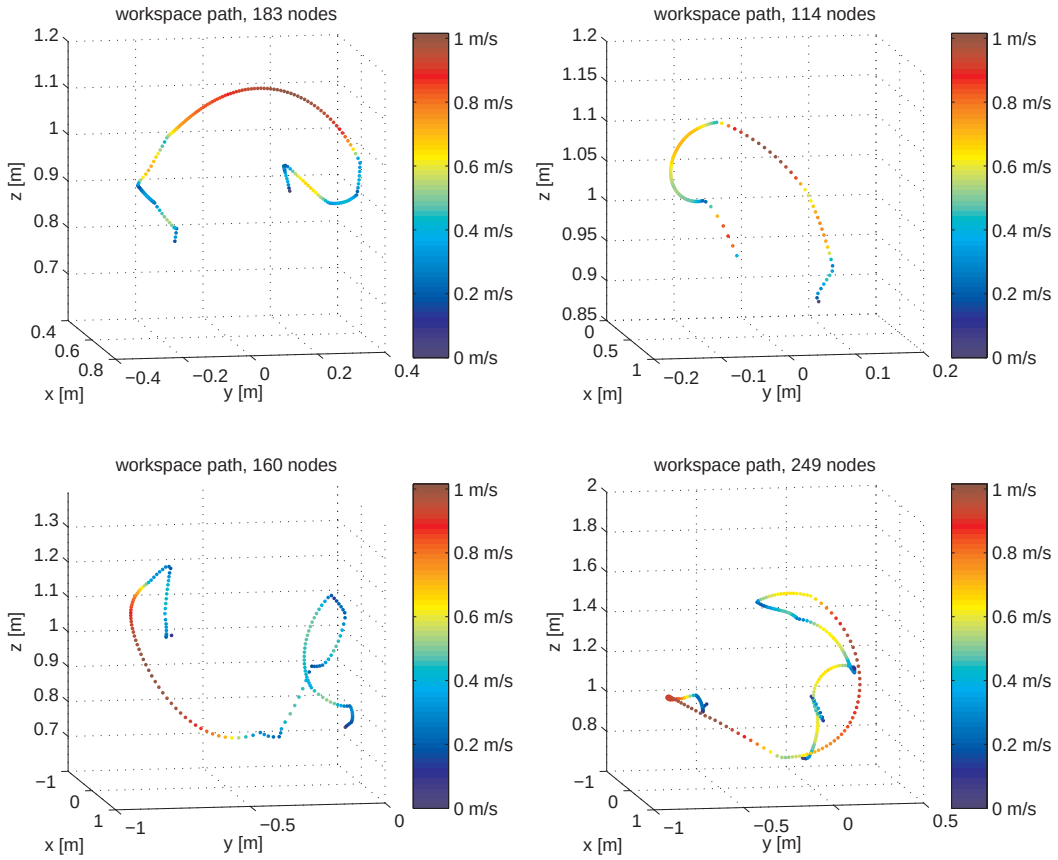

Figure 9.14: Example timings for grasp and transport paths. Speed along the path is lower when there are larger changes in direction in order to minimize jerk. The color of the path node corresponds to the workspace speed resulting from the retiming.

Spline representation for timing optimization

In the optimization of the trajectory with unit time duration, we seek to find the best relative timing between path nodes. For a path with N nodes, our search parameters are the $N-1$ incremental time durations between two nodes. Given a candidate set of $N-1$ internode durations, we scale the set to a unit sum. This normalization provides a candidate timing assignment $(t_1, t_2, \ldots, t_N)$ for the N path nodes, where $t_1 = 0$ and $t_N = 1$ to meet the unit total time constraint.

The path points combined with a candidate timing $(t_1, t_2, \ldots, t_N)$ specifies a trajectory that can be evaluated with respect to the jerk or other kinematic metrics. To evaluate the total jerk over the trajectory, we compute a spline representation from the path points and their assigned times. A spline representation with piecewise polynomials provides a smooth path interpolation between path nodes. (An alternative is to compute total jerk using a finite difference scheme, but this differencing would tend to amplify the derivative discontinuities in the linear piecewise

representation of the path.) In this section, we present spline fitting with respect to a univariate curve. The method can be used independently for each geometric dimension of the path.

For a cubic spline, each internode segment has a cubic polynomial form

$$p(t) = c_0 + c_1 t + c_2 t^2 + c_3 t^3. \tag{9.2}$$

The $4(N-1)$ coefficients for the entire cubic spline can be determined exactly from linear constraints. First, there are $2(N-1)$ position constraints of the form

$$p(t_i) = c_{0_i} + c_{1_i} t_i + c_{2_i} t_i^2 + c_{3_i} t_i^3, \tag{9.3}$$

which ensure the spline passes through the N path nodes at assigned times. There are $N-2$ velocity constraints of the form

$$c_{1_i} + 2c_{2_i} t_i + 3c_{3_i} t_i^2 = c_{1_{i+1}} + 2c_{2_{i+1}} t_i + 3c_{3_{i+1}} t_i^2, \tag{9.4}$$

which ensure velocity continuity on the interior node points. Another $N-2$ acceleration constraints of the form

$$2c_{2_i} + 6c_{3_i} t_i = 2c_{2_{i+1}} + 6c_{3_{i+1}} t_i \tag{9.5}$$

ensure acceleration continuity on the internode points. Finally two stationary boundary constraints

$$0 = c_{1_i} + 2c_{2_i} t_i + 3c_{3_i} t_i^2 \tag{9.6}$$

for $i = 1$ and $i = N-1$ enforce zero velocity at the spline end points.

The coefficients of the cubic spline are used to directly compute the jerk for the overall trajectory. Jerk is constant for each cubic spline segment such that the cost metric simplifies to

$$J_{\text{cubicspline}} = \int (\dddot{p}(t))^2 dt = \sum_{i=1}^{N-1} (6c_{3_i})^2 (t_{i+1} - t_i). \tag{9.7}$$

A quintic spline may also be used to represent the trajectory. The spline coefficients are determined by a similar linear system as that described above for the cubic spline. The additional $2(N-1)$ constraints enforce zero acceleration at the end points (2 constraints), continuous jerk at the interior points ($N-2$ constraints), and continuous snap – the fourth-order position derivative – at the interior points ($N-2$ constraints). The jerk cost function is then computed from the spline coefficients by

$$J_{\text{quinticspline}} = \int (\dddot{p}(t))^2 dt = \sum_{i=1}^{N-1} \int_{t_i}^{t_{i+1}} (6c_{3_i} + 24c_{4_i} t + 60c_{5_i} t^2)^2 dt. \tag{9.8}$$

Using the spline jerk cost functions above, we select a candidate unit timing $(t_1, t_2, \ldots, t_N)$ for the path by numerical optimization in the $(N-1)$-dimensional search space of internode time durations. Because the search space is of high dimension for a path with several nodes, the initialization is important to find a low-jerk timing solution. We address this dimensionality sensitivity with a hierarchical optimization scheme. We first optimize the timing for a coarse path consisting of the path end points and a small number of interior points. Each optimization level increases the number of included interior nodes, which increases the search space dimension. Between levels, the next timing initialization is determined by up-sampling the optimized time course according to the path distance. We verified that this method produced reasonable timing results for a simple straight line path progressively sampled at smaller intervals.

This method provides the relative timing between path nodes that minimizes jerk, within the limits of the numerical optimization search. Next we describe how the unit duration timing is used as a reference to find the minimum jerk for non-unit time durations.

Maximum speed normalization

Given a reference timing for a path, we can rescale it to meet constraints such as a new total time, maximum speed, or maximum acceleration. Here we show the relationship between a reference timing $s(t)$ and its scaling $s(r)$ by a constant c, where

$$r(t) = \frac{1}{c}t. \tag{9.9}$$

The path geometry is the same for both trajectories, and the inter-node times have the same proportions.

Fig. 9.15 shows an example for $c > 1$, where the retimed trajectory $s(r)$ is c times faster than the original $s(t)$. The velocity derivatives of the trajectories relative to the timing and retiming, $\dot{s}(t) = ds/dt$ and $s'(r) = ds/dr$ respectively, are related as:

$$\frac{ds}{dt} = \frac{dr}{dt} \cdot \frac{ds}{dr}, \tag{9.10}$$

$$\dot{s} = \frac{1}{c}s'. \tag{9.11}$$

The higher order derivatives are similarly related:

$$\frac{d^n s}{dt^n} = \frac{1}{c^n} \cdot \frac{d^n s}{dr^n}, \tag{9.12}$$

$$\tag{9.13}$$

Then we can write the jerk of the scaled trajectory J_c in terms of the jerk for the original reference trajectory J_0. This relationship is true for any two timings related

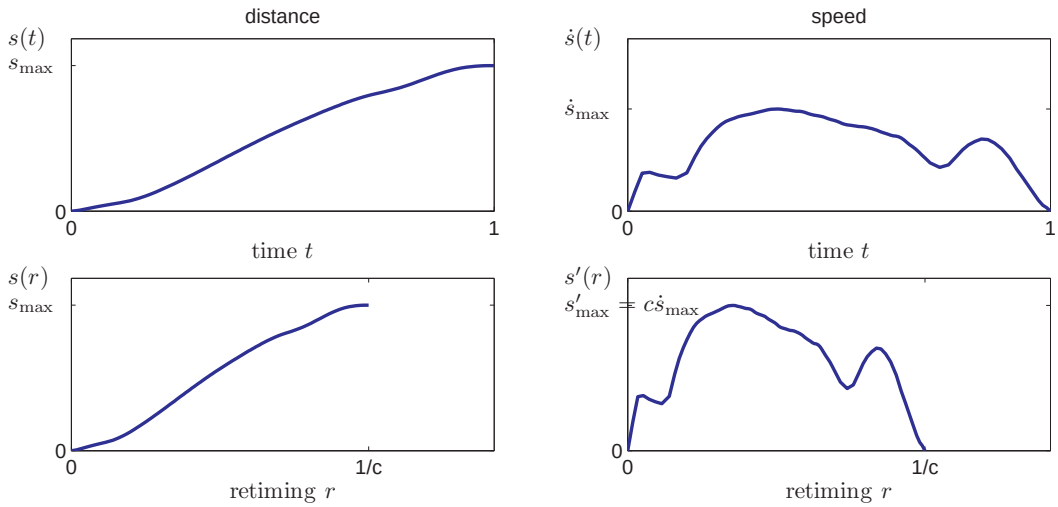

Figure 9.15: Retiming scaling relationships for distance and speed. The reference trajectory $s(t)$ travels the path length in unit time duration, and the retimed trajectory $s(r)$ travels the same path length in time duration $1/c$. The retiming corresponds to an increase in the speed along the path by factor c.

by a constant scaling, without requiring that one trajectory have unit time duration T for reference.

$$J_0 = \int_{t=0}^{t=T} \left(\frac{d^3 s}{dt^3} \right)^2 dt \tag{9.14}$$

$$J_c = \int_{r=0}^{r=\frac{T}{c}} \left(\frac{d^3 s}{dr^3} \right)^2 dr \qquad = \int_{t=0}^{t=T} \left(c^3 \frac{d^3 s}{dt^3} \right)^2 \frac{1}{c} dt = c^5 J_0 \tag{9.15}$$

This scaling relationship illustrates that the relative timing optimization in the previous section could be done with respect to any time duration, as long as it is fixed as an optimization constraint. Note that for a multi-dimensional path, each dimension must be scaled by the same constant to preserve the correct path geometry at each node. Thus for scaling based on constraints on the motion derivatives such as velocity (instead of total time constraints), it may be more practical and meaningful to choose a scale based on the path arc length to evaluate the contribution of movement from each dimension.

In our evaluation of jerk for the pre-grasp rotation plans, we rescaled the reference timing such that the retimed result achieved a given maximum speed constraint. Fig. 9.14 shows examples of the timing assignments for sample reach-to-grasp and transport paths. The timing of a given path was assigned by first minimizing the jerk of the workspace trajectory in normalized coordinates for unit time duration and then scaling the timing to achieve a maximum end-effector speed of 1 m/s.

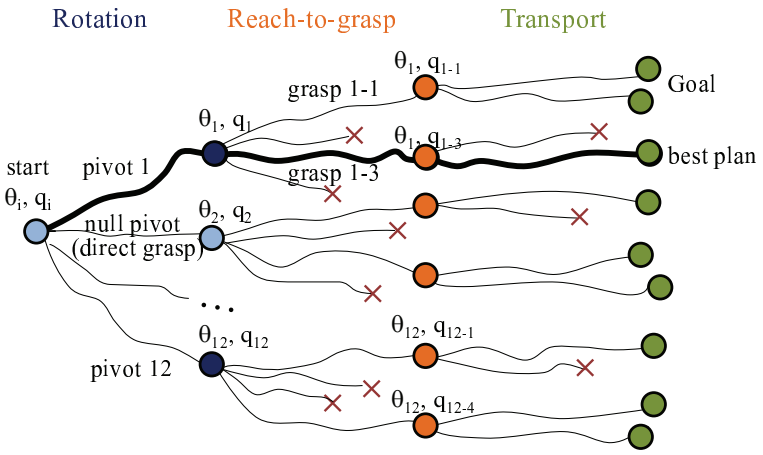

Figure 9.16: Multiple candidate paths for trajectory optimization experiments. When there is no point-based cost to optimize the grasp posture configuration of the cost metric of interest, multiple pre-grasp rotation plans are considered to different target object orientations θ. There is a null pre-grasp rotation plan for the initial object orientation if it is already reachable by direct grasping, . In addition, multiple grasp plans and transport plans are attempted for each successful pre-grasp rotation.

9.6.2 Simulation experiments

We evaluated the workspace jerk on multiple pre-grasp interaction plans (Fig. 9.16) for the pan transport task described earlier in §9.5.1. As before, task feasibility was prioritized over optimality. Thus we first searched for multiple path alternatives that satisfy path completion as in §9.5.3, evaluated the costs per path, and then selected the lowest cost option.

Since there was no direct heuristic for evaluating a grasp configuration's contribution to the total jerk, we evaluated plans to multiple object orientations θ instead a single target orientation θ_t. Note that in contrast to §9.5.3, the experiment in this section searched for multiple pre-grasp rotation plans, as illustrated in Fig. 9.16. For the pan scene, we sampled candidate object orientations θ at 10-degree intervals and grasp postures at each orientation, in the same manner as the transition point sampling described in §8.4. We searched for a pre-grasp rotation plan from the initial presented object orientation to each candidate orientation with at least one associated feasible grasp posture. Thus a pre-grasp rotation plan to an orientation was not attempted if the object could not be grasped at the new orientation (even though it could have been feasibly pivoted there).

For the pan transport task scene, we tested the pivoting from initial object orientation of 180 deg. The 7 candidate orientation angles with feasible grasp postures ranged from 60 deg to 120 deg. These are the orientations where the cooking pan handle is directed toward the robot base and overhangs the tabletop edge. For each candidate object orientation, there were $n_g = 5$ grasp plan attempts, and per suc-

cessful grasp there were $n_t = 5$ transport attempts. All paths were refined during the shortcut stage by the configuration path length.

Each successful plan within the grasp was assigned a timing using the minimum jerk optimization described above. The jerk was minimized for the end-effector workspace path separately for the rotation plan, grasp plan, and transport plan. The pre-grasp rotation plan timing was rescaled to maximum speed of 0.05 m/s to meet the quasi-static assumptions. The grasp plans and transport plans were rescaled to maximum speed of 1 m/s, which is similar to the order of magnitude of human hand speeds.

Fig. 9.17 shows the resulting jerk costs grouped by the object orientation angle θ_t at the lifting grasp transition. The initial object orientation was 180 deg. Note that the jerk costs tended to be higher for some candidate orientations, such as $\theta = 120$ deg. This example demonstrates the potential benefit of pre-grasp rotation for decreasing path jerk. A lift angle of 120 deg required the least amount of object rotation, 60 deg, from the initial object orientation of 180 deg, but the corresponding motion plans exhibited high squared jerk costs. The paths with lowest jerk costs required a larger object rotation of 120 degrees to achieve lift angle $\theta_t = 60$ deg. Thus the amount of object adjustment did not necessarily correlate with whole path cost of minimum jerk.

The plot of total path cost in Fig. 9.17 is more asymmetric than the plot of grasp posture costs in Fig. 9.1(b). We hypothesize that this asymmetry is likely due to the difference in path length for the pushing actions to achieve the grasping postures in the orientation range between 60 deg and 120 deg from the initial angle of 180 deg. In addition, the optimal object orientation for the payload safety metric cost in Fig. 9.1(b) was $\theta = 80$ deg while the lowest cost paths for the workspace squared jerk cost occurred for $\theta_t = 60$ deg .

9.7 Summary of action optimization experiments

In this chapter, we examined how modified cost criteria change the target object angle after pre-grasp rotation for the object transport tasks.

First we examined variations on the cost function for evaluating single grasping postures at candidate object angles. Changes in the manipulator torque limits shifted the limiting joint from the distal wrist joints to the proximal shoulder joints in two example tasks. However, even though the distal joints closest to the end effector were not the limiting joint, the range of cost values over the object angles did not always decrease. Thus the interaction of end-effector pose constraints and weak distal joints did not alone predict the magnitude of cost variation that measures the relative benefit of pre-grasp rotation to direct grasping.

We also evaluated whether the torque at a single joint could serve as a proxy for the multi-joint payload margin metric. The minimization of a selected joint torque

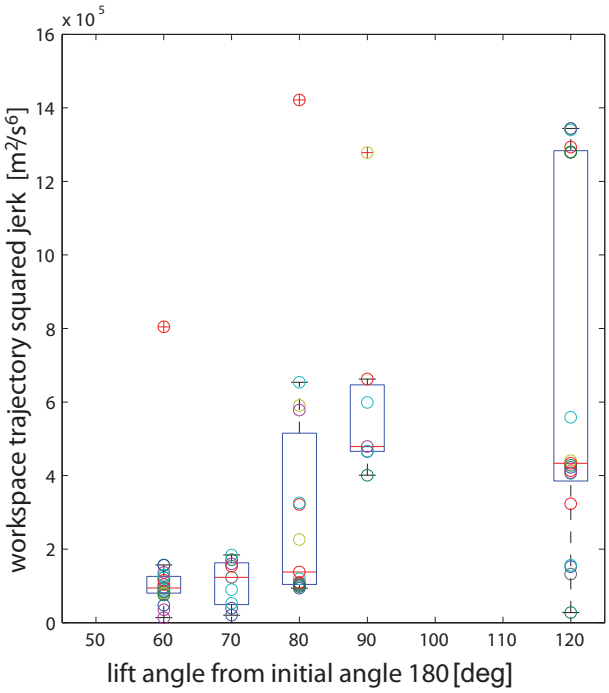

Figure 9.17: Trajectory timing optimization for minimum jerk in the example pan transport task. The squared jerk cost values of individual paths are shown as individual marks and grouped along the horizontal axis according to the object lift angle for each plan. The cost over the plan includes the jerk from the combination of the rotation, grasp, and transport actions. In this example, the initial object angle was 180 degrees. Thus a plan where the object is grasped at angle 60 degrees requires 120 degrees of object rotation.

did yield similar choices for the desired object angle and grasping postures if the joint was the same as the limiting joint of the multi-joint payload cost. However, such a proxy required prior sampling of feasible grasping postures to determine the most frequent limiting joint. The optimization method for adapting a feasible grasping posture also provided a method for refining the target object angle between the sample discretization points.

The latter two sets of experiments examined how the transport action quality could be further improved given addition computation allowance for evaluating more than only the grasping posture. In the evaluation of payload margin costs over the transport path, the high-level sampling of multiple alternative paths had larger impact on the resulting cost than low level modifications of the path segments. With the high level sampling, the direct grasping strategy could yield similar cost plans as the pre-grasp rotation strategy. This result is consistent with the result in the previous chapter §8 where the transition grasping posture selection had large impact on the transport path quality even though the transport path planner was agnostic to the payload cost metric. For the particular lifted tasks we evaluated, the path

starting point was an influential constraint on the remainder of the transport path.

We also examined optimization of path quality with respect to trajectory-based, rather than purely path-based metrics. We provided a method for optimizing the path jerk over a set of feasible pre-grasp rotation plans. Our results illustrate an example where the optimal target orientation for pre-grasp rotation is different from that predicted from evaluating the feasible grasp postures on a point-based payload safety margin metric.

Overall, our experiments showed that the advantage of pre-grasp rotation and the desired object lifting orientation will be dependent on the quality metric as well as the specific scene parameters. The optimal handle directions for lifting after pre-grasp rotation were sometimes to the sides of and not necessarily oriented toward the robot base or "torso" as observed in the human experiments. The variation in the resulting pre-grasp rotation plans and cost benefits suggests that an initial sampling to assess the possible grasping postures in a scene may be a useful pre-planning step for deciding when to attempt pre-grasp interaction.

10 Summary and future work

This thesis has investigated pre-grasp interaction as a manipulation strategy for successfully grasping objects during a difficult task. In our survey of human manipulation actions in natural settings (§3), we found several examples where humans do not or cannot grasp an object at the desired surface contacts without first re-configuring the object in the workspace. Regrasping is one method of reconfiguration, but pre-grasp interaction also occurs in the form of non-prehensile manipulation such as pushing or tumbling. We suggest that pre-grasp interaction may be especially useful for difficult tasks involving heavy objects and/or cluttered environments where the object is not conveniently presented for reaching the desired grasp.

In particular, we examined in detail the example of pre-grasp object rotation, where objects are pivoted in the plane of the horizontal support surface to achieve feasible task completion or increased action quality. The first part of the thesis in §4 and §5 investigated how humans exhibit the pre-grasp rotation strategy and possible selection criteria underlying this movement strategy. The latter portion in §6–§9 presented an analysis of how this pre-grasp rotation strategy can improve robotic manipulation by increasing task completion over a range of initial task conditions or increasing the quality of the primary manipulation action.

Overall, we promote a broad view of the richness of manipulation skills for object interaction. We have shown that the seemingly simple skill of object acquisition can involve a sophisticated process of pre-grasp manipulation to achieve the desired grasp for the subsequent task. Here we review the highlights and limitations of our findings. We also suggest avenues for future research on human manipulation, robot manipulation, and the interaction between the two.

10.1 Human manipulation strategies

Pre-grasp object interaction is a broad description of one approach humans use to complete manipulation tasks robustly. Our experiments on human pre-grasp rotation represent an initial investigation of how object manipulation actions such as pushing or pivoting interact with the grasp acquisition.

Within the pre-grasp rotation category, we investigated the effect of particular task parameters such as initial object orientation, object weight, and lifting precision

on the choice of object rotation. We also found that aspects of the manipulator, such as an individual's lifting capability or a robot's payload capacity, may drive the selection of a preferred grasping posture enabled by pre-grasp rotation. Other factors that plausibly influence the choice of preferred object orientation include the direction of transport to the goal location, the surface friction between the object and support surface, and obstacles in the environment. Learning how these factors change the performance of pre-grasp rotation may contribute to the understanding of why and when it is chosen as a motor action.

At a higher level, there remain several interesting questions about how humans choose among multiple strategies rather than the parameters for a single strategy. The taxonomy of pre-grasp interaction strategies presented in §3 categorizes interaction examples based on the degrees of freedom in the object reconfiguration. The categorization provides a guide for classes of pre-grasp strategies, but it would benefit from a model of how the task specifications determine the choice between two possible strategies. For example, for articulated pliers, how does the primary manipulation task (for example, either transport, hand-off to another person, or bending a wire) determine whether the tool is pivoted, tumbled, or folded during the pre-grasp interaction?

We have shown for pre-grasp rotation that the choice of whether and how much to rotate the object is in part a response to the task difficulty factors of object weight and upright precision for our canister lifting task experiment (§4, §5). Along these lines, it would be useful to develop a formal, generalizable representation of the primary manipulation task difficulty. Descriptors of task difficulty in such a representation may provide a mapping for predicting the probable resulting pre-grasp interaction. Example factors that we speculate contribute to difficulty and thus more sophisticated strategies include the number and occlusion of possible object surface contacts for grasps, the object inertial properties, and the friction in or resistance to the object or environment to the possible reconfiguration freedoms. Manipulator constraints of the human body – e.g., strength, bimanual availability, or position mobility – are additional considerations.

The uncertainty in and/or the human perception of these difficulty factors may be more critical than the actual values, such that cognitive load in planning a difficult maneuver is a driving factor. For example, the strategy choice will be limited by the perception of allowable object affordances, such as whether an object can be tumbled or whether a set of objects are stackable.

Finally, we have implicitly referred to pre-grasp interaction as an intentional, planned strategy that may optimize some aspect of the primary task. It is possible that some pre-grasp interactions arise from exploratory interactions with an object, such as in the haptic exploration of an unfamiliar surface for a grasp point, or from recovery responses to unintentional disturbances, such as when an object slips during an imprecise grasp action.

10.2 Robotic pre-grasp interaction

To evaluate the pre-grasp rotation strategy for a robot manipulator, we first presented a workspace analysis (§6) and a manually-programmed manipulation action that demonstrated the increased task feasibility over workspace conditions (§7). We have also contributed a method for incorporating pre-grasp object rotation automatically in an object transport plan for a robotic manipulator. An increase in the success of finding feasible task plans was also achieved by our automated planning method (§8). Furthermore, our evaluation of this method and the related optimization variations demonstrates that pre-grasp rotation can also improve plan quality in certain scenarios (§9).

There are two main limitations of the work as a step toward intelligent, robust robot manipulation. First, the method models a specific type of pre-grasp interaction in the context of an object transport task. Second, while physical demonstrations illustrate the plausibility of the manipulation plans, the advantages of increased task success and plan quality are evaluated with respect to the simulated motion plans rather than the execution.

Our initial investigation of autonomous pre-grasp rotation planning suggests a number of interesting challenges to achieve more general pre-grasp interaction capabilities for a robot. Typically the configuration space for the manipulator arm is already high-dimensional, and pre-grasp object interaction adds additional complexity of the object configuration freedoms. We showed in our examples of pre-grasp rotation that sampling key transition points in the overall manipulation plan for both the object configuration and manipulator grasp postures can focus the search and make it tractable. Even with only the single degree of freedom of object orientation, there were gains in task success rate and sometimes plan quality.

To add further pre-grasp interaction strategies to a robot's repertoire, similar reduction of the search complexity will be necessary for search tractability and high quality plans. The dimensionality quickly increases once allowable motions include more complex object reconfiguration such as 6-DoF tumbling or more manipulator freedoms such as a mobile base or bimanual manipulation. Sampling transition points within the manipulation plan is unlikely to sufficiently narrow the search without other strategic choices. For both the object and the robot motions, the space may be reduced by considering only a small set of action primitives such as bounded displacements or canonical arm trajectories. Data-driven or learning techniques comprise one approach to determining and reusing promising canonical actions.

Further investigation of robust action execution and error recovery is also necessary for bringing robot pre-grasp interaction to fruition on physical platforms. Motion plans that are feasible or even promising in simulation can fail in the real world due to uncertainty in perception or action execution. A model of the uncertainty can also be incorporated into the plan simulation itself as a metric for predicting robust-

ness of the candidate actions. Another extension could be a more sophisticated model of the hand grasp contacts on the object surface. This thesis used pre-defined hand grasp shapes, but following work by Kappler *et al.* [2010] presented an example-based approach for learning and refining hand postures for pre-grasp pushing. Overall, the features describing object localization uncertainty and the grasp contacts could be used to inform the high level decision of what type of pre-grasp interaction to use when there are multiple possible behaviors. For example, the predicted grasp success could be used as a heuristic to select from among available choices of sliding, tumbling, or pivoting.

10.3 Interactive manipulation scenarios

A common theme between human and robotic manipulation is the selection of a strategy and furthermore a specific action and the underlying criteria of either feasibility or optimality. This connection may play a role in understanding and improving the interaction between a human and a robot in a cooperative manipulation task.

In a home environment, robots may not complete tasks in isolation but instead may be observed by and interact with humans. In these situations, predictable and natural-looking motion may facilitate efficient cooperation with an end user or high-level control by a human operator. The human may perceive a robot to be more capable or easier to control if it uses similar manipulation strategies or optimizes similar task criteria that influence human actions. Future research could evaluate whether the mimicking of a particular manipulation strategy consistent with the human choice in the same task conditions would improve a human's perception of a robot manipulator's intelligence or dexterity.

Another interesting research question is what the manipulation strategy could communicate to an observer about the object or task parameters in a novel environment. This communication can work in both directions, from a human observing a robot's object manipulation before receiving the same object, to a robot estimating object properties based on tracking a human manipulation example.

10.4 Generalization of manipulation planning

10.4.1 Promising human strategies for robot repertoire

In this thesis, we presented a high-level extraction of the human pre-grasp interaction strategy for transfer to a robot manipulation technique. While grasping for object acquisition is a key part of fetching tasks, there are additional manipulation behaviors where human strategies may offer useful insight on methods for robot manipulation.

One area for future work is learning gathering strategies for carrying multiple objects at once in whole-body grasps. During our observation of human manipula-

tion behaviors presented in §3, we observed examples of stacking or piling objects in addition to pre-grasp interactions. For example, when tidying up a cluttered table or preparing ingredients for cooking, a person may grasp multiple small and/or lightweight objects at one location before walking to a new location to release or place the objects. In order to free one hand for grasping a new object, the acquired objects are supported by the upper limbs, often cradled between the elbow and torso or stacked upon other objects (Fig. 10.1(a)). In addition, multiple objects may be grasped in one hand, especially for similar objects such as jars (Fig. 10.1(b)) or even non-rigid items such as food scraps (Fig. 3.3(c)). A robot may benefit from learning patterns for groupable or stackable objects based on human examples.

The opposite maneuver of object separation from clutter is another behavior where human examples may motivate new robotic manipulation techniques. Several pre-grasp interaction examples may be motivated by the need to separate the target object from surrounding clutter in order to achieve the desired grasp. This interaction type includes the example of pivoting a book out from a bookshelf (Fig. 1.1(a)) or tumbling a large piece of furniture to find a handhold (Fig. 3.2(b)). Recent work by Dogar and Srinivasa [2010] showed how pre-grasp pushing can be used to separate an object from surrounding clutter with a simple pushing primitive in order to achieve additional clearance for the fingers to complete a grasp. However, in addition to object acquisition, separation strategies also include wedging techniques for object placement tasks. For example, when inserting a book back onto a bookshelf or a food container onto a fridge shelf, humans sometimes use the grasped object itself to push the surrounding movable objects to create sufficient support space for setting the object down. There has been some initial work on planning how a robot can individually move obstacles in a cluttered environment [Stilman and Kuffner, 2005]. The human examples of clearing a space indirectly with a grasped object as a pushing tool suggests research on different aspects of the task: the identification of pushable surfaces on the grasped object as a tool, and the regulation of the indirect pushing via force or visual feedback.

Bracing is another potentially useful behavior for a robot's repertoire. As pre-grasp interaction takes advantage of shared object load with a support surface, bracing actions use environmental structures for additional stability of the manipulator itself during a high-load action. Bracing has been used previously as a method for improving precision and reducing vibrations in fine manipulation tasks within structured environments [Hollis and Hammer, 1992; Lew and Book, 1994]. For a mobile manipulator operating amidst dynamic obstacles, perceiving and classifying objects as either load-supporting structures or non-supporting movable items is necessary for action planning with bracing.

(a) (b)

Figure 10.1: Gathering multiple objects in whole-body grasps. (a) Cradling multiple books and papers against the torso. (b) Stacking similar sized jars on top of each other to grasp four jars with two hands.

10.4.2 Motor planning optimization functions

While one approach to achieving robust and dexterous manipulation is to build a repertoire of task-specific behaviors, another avenue toward general manipulation capabilities is the identification of an explanatory optimization function for planning any action. In addition to providing insight into human cognition and motor skills, a unifying model for motor actions would avoid the need for manual observation and extraction of promising strategies for mimicry. The behaviors of gathering, separation, and bracing described previously are similar to pre-grasp interaction in that they involve some preparatory action to "set-up" a task.

Below is a non-exhaustive list of possible criteria from whose optimization could emerge the natural motor strategies of pre-grasp interaction and other behaviors:

- Time or speed
  - Time for task completion
  - Planning or decision time
  - Cognitive retrieval time or load for new versus repeated/similar actions
- Torque or load or effort
  - Total joint torque magnitudes
  - Torque of weakest joint
  - Margin from multiple joint torque limits
  - Margin from weakest joint torque limit (payload margin, lifting capability, or available strength)
  - Energy or perceived exertion

- Posture or form
  - Margin from joint limits (available travel)
  - Manipulability or margin from singularities
  - Balance or stance stability
  - Grasp stability
- Confidence, robustness, or flexibility
  - Reachability of target from multiple postures
  - Reliability of sensory information: visibility, tactile redundancy
  - Reliability of action execution or remaining within actuation capabilities
  - Availability of alternative exit strategies: setting object down, or changing grip mid-task
  - Collision avoidance or distance to obstacles
- Comfort in action execution, or avoidance of pain
- Interaction predictability or social acceptance: avoidance of awkward or unlikely postures

In this thesis we have examined a small set of metrics in the context of pre-grasp interaction. In the human motion studies described in §4 and §5, we found that time for task completion was not an explanatory metric for the pre-grasp rotation, since task completion took longer for pre-grasp rotation compared to direct grasping. However, time for planning initial movement was shorter and suggests that decision making or cognitive load could be a driving factor. We also tested metrics based on physical load at the time of object lift-off. These included the magnitude of joint torques over multiple joints (§4) and available strength for a given grasping posture (§5). While the joint loads at the lifting posture are a specific metric for pre-grasp rotation in lifting tasks, they can be a heuristic that is representative of a more general cost function of payload safety margin or joint torques over an entire action, as evaluated in §9. Furthermore we surveyed subject's self report of comfort (§4) but did not find a clear preference for the pre-grasp rotation method. An objective measure of discomfort, such as local muscle fatigue [Chaffin, 1969, 1973], may be a more reliable metric if it could be practically recorded.

Future research is required to determine if there is indeed an underlying optimization criteria driving general human manipulation actions. A holistic criteria such as robustness or success rate that includes multiple factors may be necessary for modeling complex human motion over a wide range of behaviors. In addition, human motion may not necessarily be optimal but instead merely sufficient to complete the task objective using any one of multiple similar strategies. It also remains to be seen whether identification and implementation of such an optimization criteria on a robotic system would result in the emergence of similar manipulation behaviors observed in humans, despite the differences in kinematic, sensory, actuation, and

computation capabilities.

10.5 Closing remarks

Manual dexterity is one of the unique hallmarks of human intelligence. Neuroscientists have long studied manipulation skills as a window to the brain, and I believe that *dexterity on a machine demonstrates artificial intelligence*. The dual inquiries of understanding human actions and developing robot algorithms both contribute to the progress in dexterous interaction with physical or even virtual environments. This thesis has presented several investigations on pre-grasp interaction as a particular class of manipulation strategies for movable objects. Ultimately, these are steps toward realizing the promise of fast, robust autonomous manipulation in real world environments and to enhance the quality of human interaction with the physical environment.

References

Aboaf, E. W. and Paul, R. P., 1987. Living with the singularity of robot wrists. In *IEEE International Conference on Robotics and Automation (ICRA)*, volume 4, pages 1713–1717. [§9.1]

Ahuactzin, J. M., Gupta, K., and Mazer, E., 1998. Manipulation planning for redundant robots: A practical approach. *The International Journal of Robotics Research*, 17(7):731–747. [§8.2]

Balasubramanian, R. and Matsuoka, Y., 2008. Biological stiffness control strategies for the Anatomically Correct Testbed (ACT) hand. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 737–742. [§2.2]

Barraquand, J., Langlois, B., and Latombe, J.-C., 1992. Numerical potential field techniques for robot path planning. *IEEE Transactions on Systems, Man and Cybernetics.*, 22(2):224–241. [§9.1]

Berenson, D., Chestnutt, J., Srinivasa, S. S., Kuffner, J. J., and Kagami, S., 2009a. Pose-constrained whole-body planning using task space region chains. In *IEEE/RAS International Conference on Humanoid Robots (Humanoids)*, pages 181–187. [§8.3.2, 8.3.3, 8.5.1, 9.5.1, 9.5.2]

Berenson, D., Diankov, R., Nishiwaki, K., Kagami, S., and Kuffner, J., 2007. Grasp planning in complex scenes. In *IEEE/RAS International Conference on Humanoid Robots (Humanoids)*, pages 42–48. [§2.3.2, 8.2]

Berenson, D., Srinivasa, S. S., Ferguson, D., and Kuffner, J. J., 2009b. Manipulation planning on constraint manifolds. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 625–632. [§2.3.2, 8.2]

Berenson, D., Srinivasa, S. S., Ferguson, D., Romea, A. C., and Kuffner, J., 2009c. Manipulation planning with workspace goal regions. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 618–624. [§8.1, 8.3.2, 8.4.2]

Bernardin, K., Ogawara, K., Ikeuchi, K., and Dillmann, R., 2005. A sensor fusion approach for recognizing continuous human grasping sequences using hidden Markov models. *IEEE Transactions on Robotics*, 21(1):47–57. [§2.3.1]

Bertram, D., Kuffner, J., Dillmann, R., and Asfour, T., 2006. An integrated approach to inverse kinematics and path planning for redundant manipulators. In *IEEE*

International Conference on Robotics and Automation (ICRA), pages 1874–1879. [§2.3.2, 8.2]

Billard, A. and Siegwart, R., 2004. Robot learning from demonstration. *Robotics and Autonomous Systems*, 47(2-3):65–67. [§2.4]

Bongers, R. M., Michaels, C. F., and Smitsman, A. W., 2004. Variations of tool and task characteristics reveal that tool-use postures are anticipated. *Journal of Motor Behavior*, 36(3):305–315. [§4.1]

Brock, D., 1988. Enhancing the dexterity of a robot hand using controlled slip. In *IEEE International Conference on Robotics and Automation (ICRA)*, volume 1, pages 249–251. [§2.3.3]

Brock, O. and Khatib, O., 2002. Elastic strips: A framework for motion generation in human environments. *The International Journal of Robotics Research*, 21(12):1031–1052. [§9.1]

Burgess-Limerick, R. and Abernethy, B., 1997. Qualitatively different modes of lifting. *International Journal of Industrial Ergonomics*, 19:413–417. [§2.2, 4.1, 5.1, 9.1]

Calinon, S. and Billard, A., 2008. A probabilistic programming by demonstration framework handling constraints in joint space and task space. In *IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*, pages 367–372. [§2.4]

Chaffin, D. B., 1969. Surface electromyography frequency analysis as a diagnostic tool. *Journal of Occupational Medicine*, 11(3):109–115. [§10.4.2]

—, 1973. Localized muscle fatigue–definiton and measurement. *Journal of Occupational Medicine*, 15(4):346–354. [§10.4.2]

Chang, C. C., Brown, D. R., Bloswick, D. S., and Hsiang, S. M., 2001. Biomechanical simulation of manual lifting using spacetime optimization. *Journal of Biomechanics*, 34(4):527–532. [§2.2, 4.1, 9.1]

Chang, L. Y., Klatzky, R. L., and Pollard, N. S., 2010a. Selection criteria for preparatory object rotation in manual lifting actions. *Journal of Motor Behavior*, 42(1):11–27. [§1.5, 5]

Chang, L. Y. and Pollard, N. S., 2008. On preparatory object rotation to adjust handle orientation for grasping. Technical Report CMU-RI-TR-08-10, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA. [§1.5, 4, 4.4.1, 4.6.2]

—, 2009. Video survey of pre-grasp interactions in natural hand activities. In *Robotics: Science and Systems (RSS) 2009 Workshop: Understanding the Human Hand for Advancing Robotic Manipulation*, University of Washington, Seattle, USA. [§1.5, 3, 8.1, 8.2, 8.3.1]

Chang, L. Y., Srinivasa, S. S., and Pollard., N. S., 2010b. Planning pre-grasp manipulation for transport tasks. In *IEEE International Conference on Robotics and*

Automation (ICRA), IEEE. [§1.5, 8]

Chang, L. Y., Zeglin, G. J., and Pollard, N. S., 2008. Preparatory object rotation as a human-inspired grasping strategy. In *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, Daejeon, Korea, pages 527–534. [§1.5, 4, 5.1, 5.10.1, 7, 8.1, 8.3.1]

Chen, P. C. and Hwang, Y. K., 1998. Sandros: a dynamic graph search algorithm for motion planning. *IEEE Journal of Robotics and Automation*, 14(3):390–403. [§9.1]

Chow, C. K. and Jacobson, D. H., 1971. Studies of human locomotion via optimal programming. *Mathematical Biosciences*, 10(3-4):239–306. [§4.1]

Ciocarlie, M. and Allen, P. K., 2009. Hand posture subspaces for dexterous robotic grasping. *International Journal of Robotics Research*, 28:851–867. [§2.2]

Clauser, C. E., McConville, J. T., and Young, J., 1969. Weight, volume and center of mass of segments of the human body. Technical Report AMRL-TR-69-70, Aerospace Medical Research Laboratory, Wright-Patterson Air Force Base, Ohio, Antioch College, Yellow Springs, OH. [§4.4.1, 6.2]

Cole, A., Hauser, J., and Sastry, S., 1989. Kinematics and control of multifingered hands with rolling contact. *IEEE Transactions on Automatic Control*, 34(4):398–404. [§2.3.3]

Cole, A., Hsu, P., and Sastry, S., 1992. Dynamic control of sliding by robot hands for regrasping. *IEEE Transactions on Robotics and Automation*, 8(1):42–52. [§2.3.3]

Cooper, L. A. and Shepard, R. N., 1973. The time required to prepare for a rotated stimulus. *Memory & Cognition*, 1:246–250. [§5.10.3]

Corke, P., 1996. A robotics toolbox for MATLAB. *IEEE Robotics and Automation Magazine*, 3(1):24–32. [§6.4.2, 7.5]

Cutkosky, M. R., 1989. On grasp choice, grasp models, and the design of hands for manufacturing tasks. *IEEE Journal of Robotics and Automation*, 5(3):269–279. [§2.1, 2.3.1]

Desmurget, M., Prablanc, C., Arzi, M., Rossetti, Y., Paulignan, Y., and Urquizar, C., 1996. Integrated control of hand transport and orientation during prehension movements. *Experimental Brain Research*, 110(2):265–278. [§4.1]

Diankov, R., Kanade, T., and Kuffner, J., 2009. Integrating grasp planning and visual feedback for reliable manipulation. In *IEEE/RAS International Conference on Humanoid Robots (Humanoids)*, pages 646–652. [§9.1]

Diankov, R. and Kuffner, J., 2008. OpenRAVE: A planning architecture for autonomous robotics. Technical Report CMU-RI-TR-08-34, Robotics Institute, Pittsburgh, PA. [§8.4.2]

Diankov, R., Ratliff, N., Ferguson, D., Srinivasa, S., and Kuffner, J., 2008a. Bispace

planning: Concurrent multi-space exploration. In *Robotics: Science and Systems*. [§2.3.2]

Diankov, R., Srinivasa, S. S., Ferguson, D., and Kuffner, J., 2008b. Manipulation planning with caging grasps. In *IEEE/RAS International Conference on Humanoid Robots (Humanoids)*, pages 285–292. [§8.2, 8.3.1, 8.5.1, 9.5.1]

Dogar, M. and Srinivasa, S., 2010. Push-grasping with dexterous hands: Mechanics and a method. In *IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*. [§10.4.1]

Doriot, N. and Wang, X., 2006. Effects of age and gender on maximum voluntary range of motion of the upper body joints. *Ergonomics*, 49(3):269–281. [§6.2]

Dysart, M. J. and Woldstad, J. C., 1996. Posture prediction for static sagittal-plane lifting. *Journal of Biomechanics*, 29(10):1393–1397. [§2.2, 4.1, 9.1]

Edwards, S. J., Buckland, D. J., and McCoy-Powlen, J. D., 2002. *Developmental & Functional Hand Grasps*. Slack Incorporated, Thorofare, New Jersey. [§2.1]

Ekvall, S. and Kragic, D., 2005. Grasp recognition for programming by demonstration. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 748–753. [§2.3.1]

Fischman, M. G., 1998. Constraints on grip-selection: minimizing awkwardness. *Perceptual and Motor Skills*, 86(1):328–330. [§4.1, 5.1, 9.1]

Flash, T. and Hogan, N., 1985. The coordination of arm movements: an experimentally confirmed mathematical model. *Journal of Neuroscience*, 5(7):1688–1703. [§2.2, 9.1]

Fu, J. L. and Pollard, N. S., 2006. On the importance of asymmetries in grasp quality metrics for tendon driven hands. In *IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*, pages 1068–1075. [§2.3.1]

Gienger, M., Toussaint, M., and Goerick, C., 2008. Task maps in humanoid robot manipulation. In *IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*, pages 2758–2764. [§2.3.2, 8.1]

Hauser, K., Ng-Thow-Hing, V., and Gonzalez-Banos, H., 2007. Multi-modal motion planning for a humanoid manipulation task. In *Proceedings of the International Symposium on Robotics Research (ISRR)*. [§2.3.2, 7.1, 8.2]

Hollerbach, J. and Suh, K., 1985. Redundancy resolution of manipulators through torque optimization. In *IEEE International Conference on Robotics and Automation (ICRA)*, volume 2, pages 1016–1021. [§9.1]

Hollis, R. L. and Hammer, R., 1992. Real and virtual coarse-fine robot bracing strategies for precision assembly. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 767–774. [§10.4.1]

Hsiao, K. and Lozano-Perez, T., 2006. Imitation learning of whole-body grasps. In

IEEE/RSJ Conference on Intelligent Robots and Systems (IROS), pages 5657–5662. [§2.4]

Huber, M., Rickert, M., Knoll, A., Brandt, T., and Glasauer, S., 2008. Human-robot interaction in handing-over tasks. In *IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 107–112. [§2.5]

Ikeuchi, K. and Suehiro, T., 1994. Toward an assembly plan from observation Part I. Task recognition with polyhedral objects. *IEEE Journal of Robotics and Automation*, 10(3):368–385. [§2.3.1, 2.4]

Jarvis, R. A., 1985. Collision-free trajectory planning using distance transforms. In *Mechanical Engineering Transaction of The Institution of Engineers*, volume ME 10, pages 187–191. [§9.1]

—, 1989. Collision-free path planning in time-varying environments. In *IEEE/RSJ International Workshop on Intelligent Robots and Systems*, pages 99–106. [§9.1]

Jeannerod, M., 1981. Intersegmental coordination during reaching at natural visual objects. In J. Long and A. Baddeley, editors, *Attention and Performance IX*. Lawrence Erlbaum Associates, pages 153–169. [§2.2, 4.1]

Johansson, R., 1996. Sensory control of dexterous manipulation in humans. In A. Wing, P. Haggard, and J. J.R. Flangan, editors, *Hand and brain: The neurophysiology and psychology of hand movements*. New York: Academic Press, pages 381–414. [§2.2, 4.1]

Kamakura, N., Matsuo, M., Ishii, H., Mitsuboshi, F., and Miura, Y., 1980. Patterns of static prehension in normal hands. *American Journal of Occupational Therapy*, 34(7):437–445. [§2.1]

Kaneko, M., Shirai, T., and Tsuji, T., 2000. Scale-dependent grasp. *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans*, 30(6):806–816. [§2.3.3]

Kang, S. B., 1994. *Robot Instruction by Human Demonstration*. Ph.D. thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA. [§2.4]

Kang, S. B. and Ikeuchi, K., 1997. Toward automatic robot instruction from perception-mapping human grasps to manipulator grasps. *IEEE Journal of Robotics and Automation*, 13(1):81–95. [§2.3.1, 2.4]

Kappler, D., Chang, L. Y., Przybylski, M., Pollard, N., Asfour, T., and Dillmann, R., 2010. Representation of pre-grasp strategies for object manipulation. In *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*. [§10.2]

Kavraki, L. and Latombe, J. C., 1994. Randomized preprocessing of configuration space for path planning: articulated robots. In *IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*, volume 3, pages 1764–1771. [§8.2]

Kavraki, L. E. and Latombe, J.-C., 1998. *Practical Motion Planning in Robotics:*

Current Approaches and Future Directions, chapter Probabilistic Roadmaps for Robot Path Planning. John Wiley, pages 33–53. [§9.1]

Kavraki, L. E., Svestka, P., Latombe, J. C., and Overmars, M. H., 1996. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Journal of Robotics and Automation*, 12(4):566–580. [§8.2]

Kilner, J., de C Hamilton, A. F., and Blakemore, S.-J., 2007. Interference effect of observed human movement on action is due to velocity profile of biological motion. *Soc Neurosci*, 2(3-4):158–166. [§2.5]

Kleinmann, K., Hennig, J.-O., Ruhm, C., and Tolle, H., 1996. Object manipulation by a multifingered gripper: on the transition from precision to power grasp. In *IEEE International Conference on Robotics and Automation (ICRA)*, volume 3, pages 2761–2766. [§2.3.3]

Kry, P. G., 2005. *Interaction Capture and Synthesis of Human Hands*. Ph.D. thesis, University of British Columbia. [§2.4]

Lacquaniti, F. and Soechting, J. F., 1982. Coordination of arm and wrist motion during a reaching task. *Journal of Neuroscience*, 2(4):399–408. [§2.2]

Latash, M. L., Gelfand, I. M., Li, Z. M., and Zatsiorsky, V. M., 1998. Changes in the force-sharing pattern induced by modifications of visual feedback during force production by a set of fingers. *Experimental Brain Research*, 123(3):255–262. [§2.2, 4.1]

Latash, M. L. and Jaric, S., 2002. Organization of drinking: postural characteristics of arm-head coordination. *Journal of Motor Behavior*, 34(2):139–150. [§5.1, 5.10.2]

Latombe, J.-C., 1991. *Robot Motion Planning*. Kluwer Academic Publishers, Norwell, MA, USA. [§8.2]

LaValle, S. M., Gonzalez-Banos, H. H., Becker, C., and Latombe, J.-C., 1997. Motion strategies for maintaining visibility of a moving target. In *IEEE International Conference on Robotics and Automation (ICRA)*, volume 1, pages 731–736. [§9.1]

LaValle, S. M. and Kuffner, J. J., 2000. Rapidly-exploring random trees: Progress and prospects. In A. K. Peters, editor, *Robotics: The Algorithmic Perspective. 4th Int'l Workshop on the Algorithmic Foundations of Robotics*. Hanover, NH. [§2.3.2, 8.2]

Lederman, S. J. and Klatzky, R. L., 1993. Extracting object properties through haptic exploration. *Acta Psychologica*, 84(1):29–40. [§2.2]

Lee, P., Wei, S., Zhao, J., and Badler, N. I., 1990. Strength guided motion. In *SIGGRAPH '90: Proceedings of the 17th Annual Conference on Computer Graphics and Interactive Techniques*, ACM, New York, NY, USA, pages 253–262. [§9.1]

Lew, J. Y. and Book, W. J., 1994. Bracing micro/macro manipulators control. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 2362–

2368. [§10.4.1]

Li, Y. and Pollard, N., 2005. A shape matching algorithm for synthesizing humanlike enveloping grasps. In *IEEE-RAS International Conference on Humanoid Robots (Humanoids 2005)*. [§2.3.1]

Li, Z., Hsu, P., and Sastry, S., 1989. Grasping and coordinated manipulation by a multifingered robot hand. *The International Journal of Robotics Research*, 8(4):33–50. [§2.3.1]

Li, Z. and Sastry, S., 1988. Task-oriented optimal grasping by multifingered robot hands. *IEEE Journal of Robotics and Automation*, 4(1):32–44. [§2.3.1]

Li, Z. M., Latash, M. L., and Zatsiorsky, V. M., 1998. Force sharing among fingers as a model of the redundancy problem. *Experimental Brain Research*, 119(3):276–286. [§2.2, 4.1]

Lin, Q., Burdick, J., and Rimon, E., 2000. A stiffness-based quality measure for compliant grasps and fixtures. *IEEE Transactions on Robotics and Automation*, 16(6):675–688. [§2.3.1]

Lozano-Perez, T., 1983. Spatial planning: A configuration space approach. *IEEE Transactions on Computers*, (2):108–120. [§8.2]

—, 1987. A simple motion-planning algorithm for general robot manipulators. *IEEE Journal of Robotics and Automation*, 3(3):224–238. [§8.2]

Lozano-Perez, T., Mason, M., and Taylor, R. H., 1984. Automatic synthesis of fine-motion strategies for robots. *International Journal of Robotics Research*, 3(1):3–24. [§2.3.2, 7.1, 8.2]

Lukos, J., Ansuini, C., and Santello, M., 2007. Choice of contact points during multidigit grasping: effect of predictability of object center of mass location. *Journal of Neuroscience*, 27(14):3894–3903. [§2.2]

Lukos, J. R., Ansuini, C., and Santello, M., 2008. Anticipatory control of grasping: independence of sensorimotor memories for kinematics and kinetics. *Journal of Neuroscience*, 28(48):12765–12774. [§2.2]

Lynch, K., 1999. Toppling manipulation. In *IEEE International Conference on Robotics and Automation (ICRA)*, volume 4, pages 2551–2557. [§2.3.2, 8.2]

Lynch, K. M., Maekawa, H., and Tanie, K., 1992. Manipulation and active sensing by pushing using tactile feedback. In *IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*, volume 1, pages 416–421. [§9.1]

Lynch, K. M. and Mason, M. T., 1995. Controllability of pushing. In *IEEE International Conference on Robotics and Automation (ICRA)*, Nagoya, Japan, pages 112–119. [§2.3.2, 7.1, 8.2]

Mason, C., Gomez, J., and Ebner, T., 2001. Hand synergies during reach-to-grasp. *Journal of Neurophysiology*, (86):2896–2910. [§2.1, 2.2]

Mason, M. T., 1986. Mechanics and planning of manipulator pushing operations. *International Journal of Robotics Research*, 5(3):53–71. [§2.3.2, 7.1, 8.2]

—, 2001. *Mechanics of Robotic Manipulation*. The MIT Press. [§2.3.1]

Mishra, B., Schwartz, J. T., and Sharir, M., 1987. On the existence and synthesis of multifinger positive grips. *Algorithmica*, 2:541–558. [§2.3.1]

Moussa, M. and Kamel, M., 1995. A connectionist model of human grasps and its application to robot grasping. In *IEEE International Conference on Neural Networks*, volume 5, pages 2555–2559. [§2.3.1]

Napier, J. R., 1993. *Hands*. Princeton University Press, Princeton, New Jersey. [§2.1]

Ng-Thow-Hing, V., Drumwright, E., Hauser, K., Wu, Q., and Wormer, J., 2007. Expanding task functionality in established humanoid robots. In *IEEE/RAS International Conference on Humanoid Robots (Humanoids 2007)*. [§2.3.2, 7.1, 8.2]

Park, W., Martin, B. J., Choe, S., Chaffin, D. B., and Reed, M. P., 2005. Representing and identifying alternative movement techniques for goal-directed manual tasks. *Journal of Biomechanics*, 38(3):519–527. [§2.2, 4.1]

Pinheiro, J. C. and Bates, D. M., 2000. *Mixed-effects models in S and S-PLUS*. Springer, New York. [§4.4.3, 5.6.3]

Pollard, N., 2004. Closure and quality equivalence for efficient synthesis of grasps from examples. *International Journal of Robotics Research*, 23(6):595–614. [§2.3.1, 2.4]

Pollard, N. and Zordan, V. B., 2005. Physically based grasping control from example. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. [§2.4]

Ponce, J., Sullivan, S., Sudsang, A., Boissonnat, J., and Merlet, J., 1997. On computing four-finger equilibrium and force-closure grasps of polyhedral objects. *International Journal of Robotics Research*, 16:11–35. [§2.3.1]

Quinlan, S. and Khatib, O., 1993. Elastic bands: connecting path planning and control. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 802–807. [§9.1]

R Development Core Team, 2008. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. http://www.R-project.org. [§4.4.3, 5.6.3]

Rand, M. K. and Stelmach, G. E., 2005. Effect of orienting the finger opposition space in the control of reach-to-grasp movements. *Journal of Motor Behavior*, 37(1):65–78. [§4.1]

Ratliff, N., Bagnell, J. A., and Srinivasa, S., 2007. Imitation learning for locomotion and manipulation. In *IEEE-RAS International Conference on Humanoid Robots*. [§2.4]

Ratliff, N., Zucker, M., Bagnell, J. A., and Srinivasa, S., 2009. Chomp: Gradient optimization techniques for efficient motion planning. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 489–494. [§9.1]

Rosenbaum, D. A., Engelbrecht, S. E., Bushe, M. M., and Loukopoulos, L. D., 1993a. A model for reaching control. *Acta Psychologica*, 82(1-3):237–250. [§5.1, 5.10.1]

Rosenbaum, D. A. and Gaydos, M. J., 2008. A method for obtaining psychophysical estimates of movement costs. *Journal of Motor Behavior*, 40(1):11–17. [§4.1]

Rosenbaum, D. A., Loukopoulos, L. D., Meulenbroek, R. G., Vaughan, J., and Engelbrecht, S. E., 1995. Planning reaches by evaluating stored postures. *Psychological Review*, 102(1):28–67. [§5.1, 5.10.1]

Rosenbaum, D. A., Marchak, F., Barnes, H. J., Vaughan, J., Slotta, J. D., and Jorgensen, M. J., 1990. Constraints for action selection: Overhand versus underhand grips. In M. Jeannerod, editor, *Attention and Performance XIII: Motor Representation and Control*. Lawrence Erlbaum Associates, Hillsdale, NJ, pages 321–342. [§4.1, 5.1, 5.10.1]

Rosenbaum, D. A., van Heugten, C. M., and Caldwell, G. E., 1996. From cognition to biomechanics and back: the end-state comfort effect and the middle-is-faster effect. *Acta Psychologica*, 94(1):59–85. [§5.1, 5.10.2, 9.1]

Rosenbaum, D. A., Vaughan, J., Barnes, H. J., and Jorgensen, M. J., 1992. Time course of movement planning: selection of handgrips for object manipulation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18(5):1058–1073. [§4.1, 5.1, 9.1]

Rosenbaum, D. A., Vaughan, J., Jorgensen, M. J., Barnes, H. J., and Stewart, E., 1993b. Plans for object manipulation. In D. E. Meyer and S. Kornblum, editors, *Attention and performance XIV - A silver jubilee: Synergies in experimental psychology, artificial intelligence and cognitive neuroscience*. MIT Press, Bradford Books, Cambridge, pages 803–820. [§4.1, 5.1, 5.10.1]

Safonova, A., Hodgins, J. K., and Pollard, N. S., 2004. Synthesizing physically realistic human motion in low-dimensional, behavior-specific spaces. *ACM Transactions on Graphics (SIGGRAPH 2004)*, 23(3). [§2.4]

Sahbani, A., Cortes, J., and Simeon, T., 2002. A probabilistic algorithm for manipulation planning under continuous grasps and placements. In *IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*, volume 2, pages 1560–1565. [§2.3.2, 8.2]

Santello, M., Flanders, M., and Soechting, J., 1998. Postural hand synergies for tool use. *Journal of Neuroscience*, (18):10105–15. [§2.1, 2.2]

Saxena, A., Driemeyer, J., and Ng, A. Y., 2008. Robotic Grasping of Novel Objects using Vision. *The International Journal of Robotics Research*, 27(2):157–173. [§2.4]

Shibata, S. and Inooka, H., 1998. Psychological evaluations of robot motions. *International Journal of Industrial Ergonomics*, 21(6):483–494. [§2.5]

Shiller, Z., 1994. Time-energy optimal control of articulated systems with geometric path constraints. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 2680–2685. [§9.1]

—, 2000. Online suboptimal obstacle avoidance. *The International Journal of Robotics Research*, 19(5):480–497. [§9.1]

Shiller, Z. and Dubowsky, S., 1987. Time optimal paths and acceleration lines of robotic manipulators. In *Proceedings of the 26th IEEE Conference on Decision and Control*, volume 26, pages 199–204. [§9.1]

—, 1991. On computing the global time-optimal motions of robotic manipulators in the presence of obstacles. *IEEE Journal of Robotics and Automation*, 7(6):785–797. [§9.1]

Short, M. W. and Cauraugh, J. H., 1997. Planning macroscopic aspects of manual control: end-state comfort and point-of-change effects. *Acta Psychologica*, 96(1-2):133–147. [§4.1, 5.1, 9.1]

Simeon, T., Cortes, J., Sahbani, A., and Laumond, J., 2002. A manipulation planner for pick and place operations under continuous grasps and placements. In *IEEE International Conference on Robotics and Automation (ICRA)*, volume 2, pages 2022–2027. [§2.3.2, 8.2]

Simeon, T., Laumond, J.-P., Cortes, J., and Sahbani, A., 2004. Manipulation planning with probabilistic roadmaps. *International Journal of Robotics Research*, 23(7-8):729–746. [§8.2]

Spong, M. W., Hutchinson, S., and Vidyasagar, M., 2006. *Robot Modeling and Control*, chapter 3.2. John Wiley & Sons, Inc., Hoboken, NJ. [§6.2]

Srinivasa, S. S., Ferguson, D., Helfrich, C., Berenson, D., Romea, A. C., Diankov, R., Gallagher, G., Hollinger, G., Kuffner, J. J., and VandeWeghe, M., 2010. HERB: a home exploring robotic butler. *Autonomous Robots*, 28(1):5–20. [§8.5.3]

Stanley, J., Gowen, E., and Miall, R. C., 2007. Effects of agency on movement interference during observation of a moving dot stimulus. *Journal of Experimental Psychology – Human Perception and Performance*, 33(4):915–926. [§2.5]

Stelmach, G. E., Castiello, U., and Jeannerod, M., 1994. Orienting the finger opposition space during prehension movements. *Journal of Motor Behavior*, 26(2):178–186. [§4.1]

Stilman, M. and Kuffner, J., 2005. Navigation Among Movable Obstacles: Real-Time Reasoning In Complex Environments. *International Journal of Humanoid Robotics*, 2(4):479–503. [§2.3.2, 7.1, 8.2, 10.4.1]

Stoytchev, A., 2005. Behavior-grounded representation of tool affordances. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 3060–3065. [§2.2]

Torabi, L., Kazemi, M., and Gupta, K., 2007. Configuration space based efficient view planning and exploration with occupancy grids. In *IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*, pages 2827–2832. [§9.1]

Tournassoud, P., Lozano-Perez, T., and Mazer, E., 1987. Regrasping. In *IEEE International Conference on Robotics and Automation (ICRA)*, volume 4, pages 1924–1928. [§8.2]

Turvey, M. T., Shockley, K., and Carello, C., 1999. Affordance, proper function, and the physical basis of perceived heaviness. *Cognition*, 73(2):B17–B26. [§4.1]

Uno, Y., Kawato, M., and Suzuki, R., 1989. Formation and control of optimal trajectory in human multijoint arm movement. minimum torque-change model. *Biological Cybernetics*, 61(2):89–101. [§2.2, 9.1]

Verbeke, G. and Molenberghs, G., 2000. *Linear mixed models for longitudinal data*. Springer, New York. [§4.4.3, 5.6.3]

Weghe, M. V., Ferguson, D., and Srinivasa, S., 2007. Randomized path planning for redundant manipulators without inverse kinematics. In *IEEE/RAS International Conference on Humanoid Robots (Humanoids)*. [§2.3.2, 8.2]

Yamane, K., Kuffner, J., and Hodgins, J. K., 2004. Synthesizing animations of human manipulation tasks. *ACM Transactions on Graphics (SIGGRAPH 2004)*. [§2.4, 8.2]

Yen, V. and Nagurka, M., 1987. Suboptimal trajectory planning of a five-link human locomotion model,. In J. L. Stein, editor, *Biomechanics of Normal and Prosthetic Gait*, volume BED-Vol. 4 and DSC Vol. 7, ASME Winter Annual Meeting, Boston, MA, pages 17–22. [§9.1]

Yoshida, E., Poirier, M., Laumond, J.-P., Alami, R., and Yokoi, K., 2007. Pivoting based manipulation by humanoids: a controllability analysis. In *IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*, pages 1130–1135. [§8.2]

Yoshida, E., Poirier, M., Laumond, J. P., Kanoun, O., Lamiraux, F., Alami, R., and Yokoi, K., 2008. Whole-body motion planning for pivoting based manipulation by humanoids. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 3181–3186. [§8.2]

Yoshida, E., Poirier, M., Laumond, J.-P., Kanoun, O., Lamiraux, F., Alami, R., and Yokoi, K., 2010. Pivoting based manipulation by a humanoid robot. *Autonomous Robots*, 28(1):77–88. [§8.2]

Zhang, W. and Rosenbaum, D. A., 2008. Planning for manual positioning: the end-state comfort effect for manual abduction-adduction. *Experimental Brain Research*, 184(3):383–389. [§4.1, 5.1, 9.1]

Zhu, X., Ding, H., and Li, H., 2001. A quantitative measure for multi-fingered grasps. In *Proceedings of the IEEE/ASME International Conference on Advanced Intelligent Mechatronics*, page 213219. [§2.3.1]