Generality and Simple Hands

Matthew T. Mason, Siddhartha S. Srinivasa, Andres S. Vazquez and Alberto Rodriguez

November 24, 2010

CMU-RI-TR-10-40

November 2010

Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

© Carnegie Mellon University

Abstract

While complex hands seem to offer generality, simple hands are often more practical. This raises the question: how do generality and simplicity trade off in the design of robot hands? This paper explores the tension between simplicity in hand design and generality in hand function. It raises arguments both for and against simple hands; it considers several familiar examples; and it proposes a concept for a simple hand design with associated strategies for grasping, recognition, and localization. The central idea is to use learned knowledge of stable grasp poses as a cue for object recognition and localization. This leads to some novel design criteria, such as minimizing the number of stable grasp poses. Finally, we present some experimental results for a prototype simple hand in a bin-picking task.

Contents

1	Introduction	1
2	Previous work	3
2.1	Simple hands	5
2.2	Object recognition and localization	6
3	Design Goals	7
3.1	Stability	8
3.2	Capture	8
3.3	In-hand manipulation	9
3.4	Clutter	9
3.5	Object shape variation	10
3.6	Multiple and deformable objects	11
3.7	Recognition, Localization	11
3.8	Placing.	11
3.9	Other issues and requirements	11
4	Approach: Let the fingers fall where they may	11
5	Stable pose field	13
6	Experimental prototype	14
7	Experimental results	18
7.1	Grasp Recognition: Object singulation	19
7.2	In-hand Localization: Pose estimation	19
8	Conclusion	21

1 Introduction

Complex hands surely offer greater generality than simple hands. Yet simple hands surely offer greater generality than any autonomous system has yet demonstrated. There are both practical and scientific reasons to explore the potential generality of simple hands. The practical reason is that simple hands are smaller, lighter, and less expensive. For many applications such as micro-manipulation or minimally invasive surgery, simple hands will be in use for the foreseeable future. The scientific reason is that simple systems are easier to study. Study of simple hands may yield insights applicable to both simple and complex hands.

How do we define simplicity and generality? By a *simple hand* we mean a hand with a few actuators, a few simple sensors, and without complicated mechanisms, so that the whole hand can be small, light, and inexpensive. Figure 1 shows two examples: a simple hardware store gripper, and P1, a prototype simple gripper developed in this work for a bin-picking task. By *generality* we mean that the hand should address a broad range of tasks and task environments.

What is the nature of the tradeoff between simplicity and generality in a robot hand? Some arguments in favor of complexity are:

- Grippers for manufacturing automation are often simple and highly specialized, perhaps designed to grasp a single part. Human hands, in contrast, are complex and general;
- Hands grasp by conforming to the object shape. Motion freedoms are a direct measure of a hand's possible shape variations.
- Beyond grasping, many tasks benefit from more complexity. Manipulation in the hand, and haptic sensing of shape, to mention two important capabilities, benefit from more fingers, more controlled freedoms, and more sensors.
- The most general argument is: Design constraints have consequences. Restricting actuators, sensors, and fingers to low numbers eliminates most of the hand design space.

However, there *are* simple but general grippers, for example prosthetic hooks, teleoperated systems with simple pincer grippers, and the simple pickup tool shown in Figure 1a. With a human in the loop, unhindered by the limitations of current autonomous systems, we see the true generality offered by simple hands.

We conclude that while there is a tradeoff between simplicity and generality, the details of that tradeoff are important and poorly understood. Simple grippers can achieve a level of generality that is yet untapped in autonomous robotic systems. And whatever the limitations of any particular simple gripper, the most general robots will surely want a toolbox full of simple grippers for some tasks, just as humans do.

Our approach to grasping was inspired by the simple pickup tool of Figure 1a. The pickup tool is designed to grasp various shapes without careful positioning of fingers relative to the object. Our first prototype P1 in Figure 1b mimics the pickup tool but also addresses some of its limitations. Sometimes a robot needs not only to grasp, but also to know what it has grasped (recognition), and to know the object pose within the hand (localization). Our approach then is to apply the blind or near-blind grasping strategy of the pickup tool, augmented by simple sensors and offline training to address recognition and localization.

To test the generality of this approach, we plan to explore a variety of objects in two different task domains: bin-picking for industrial automation, and object retrieval for a domestic service robot. This

(a)

(b)

Figure 1: (a) The common “pickup tool” is very simple, but also very effective in achieving stable grasps over a broad class of shapes. Four fingers of spring steel are driven by a single actuator. (b) Our first prototype gripper P1 was inspired by the pickup tool, but has three fingers rather than four, and has angle encoders for object recognition and localization.

paper focuses on the first application: bin-picking. We have barely begun the work on object retrieval and hope to report it in a later publication. Object retrieval does have some similarities to bin-picking, especially the problem of retrieving a particular object from a cluttered drawer.

The paper is organized as a design study. Following a discussion of previous work (Section 2), we discuss the set of goals or requirements for a general-purpose gripper (Section 3). Then we more fully describe our approach to grasping (Section 4). Section 5 describes numerical and analytical results with a simple idealized model. Section 6 describes the prototype P1 in greater detail. Section 7 describes experimental results with the prototype.

2 Previous work

This section addresses previous work on the design of hands and the tradeoffs between simplicity and generality. Some discussion of previous work appears elsewhere in the paper: Section 3 addresses previous work on specific topics such as stability, capture, clutter, object recognition, and localization. Section 7 mentions some previous work on bin picking.

Our work touches on many different areas of robotic manipulation, including topics that predate robotics. Our long-range goal, attaining general-purpose manipulation with simple hands, is of increasing interest, in part due to applications in service robotics, including for example teleoperated mobile manipulators for bomb disposal [94, 85], and domestic service robots for object retrieval [77, 50, 20, 78]. Generality is required because of the wide range of objects encountered, and the unpredictable variation from one task to the next. Simplicity is required because the manipulator is often mounted on a mobile platform, which constrains weight, size, and power. Additional considerations, including cost and an immediate need, favor the use of very simple grippers.

Most relevant to this study is the recent work by Dollar, Howe and colleagues. They have explored a simple planar two-fingered gripper, with two joints per finger [29], to attain grasp generality over object shape and pose. In subsequent work a three-dimensional version employs four two-jointed fingers, all compliantly coupled to a single actuator [30].

Ciocarlie and Allen [21] adopted a similar design, with two fingers, three joints per finger, all compliantly coupled to a single actuator, and tuned the compliant coupling to optimize the grasping outcomes across a broad range of object shapes.

Another closely related example is the end effector developed by Xu, Deyle and Kemp [93], resembling a dustpan combined with a multi-jointed compliantly coupled finger driven by a single actuator. In thorough empirical testing they demonstrated its effectiveness in capturing a wide range of objects with substantial variation in pose. Of special interest is the list of objects they work with, compiled after surveying motor impaired patients to determine the relative importance of household objects for automatic retrieval [19].

Although interest in generality and simple hands is high today, the tradeoff between simplicity and generality was discussed even as the first robotic hands were being developed 50 years ago. Tomovic and Boni [84] noted that for some designs additional hand movements would require additional parts, leading to “unreasonably complex mechanical devices”. The Tomovic/Boni hand, commonly called the “Belgrade hand”, was designed for prosthetic use, but with reference to possible implications for “automatic material handling equipment”.

Jacobsen and colleagues [47] likewise raised the issue in the context of the Utah/MIT Dextrous Hand, over 25 years ago:

“A very interesting issue, frequently debated, is the question ‘How much increased

function versus cost can be achieved as a result of adding complexity to an end-effector?’ . . . In short then the question (in the limit) becomes ‘Would you prefer N each, one-degree-of-freedom grippers or one each N-degree-of-freedom grippers?’ We believe that a series of well-designed experiments with the DH [dexterous hand] presented here could provide valuable insights about the trade-offs between complexity, cost, and functionality.”

We share the interest in the tradeoff between simplicity and generality, but we depart from Jacobsen et al.’s thinking in two ways. First, we focus on the generality achievable with a single simple hand, rather than a toolbox of highly specialized grippers. Second, while we agree that complex hands can be used to emulate simpler devices, we work directly with simpler effectors designed specifically to explore the issues.

Three approaches to generality have dominated hand design research: anthropomorphism, grasp taxonomies, and in-hand manipulation.

- *Anthropomorphism.* Several reasons have been cited for directly emulating the form of the human hand: anthropomorphic designs interface well with anthropic environments; they are the most natural teleoperator device; they facilitate comparisons and interchange between biomechanical and robotic studies; they have purely esthetic advantages in some assistive, prosthetic, and entertainment applications; they are well-suited for communication by gesture; and finally and most simply, the human hand seems to be a good design [7, 27, 47, 66, 95].
- *Grasp taxonomies.* Rather than directly emulating the human hand, several hand designs address generality by emulating the most common grasp poses of the human hand. Grasp taxonomies [26, 73, 62] are the perfect source to evaluate the generality of hand designs over grasp pose.
- *In-hand manipulation.* The ability to use the fingers to control the six rigid-body degrees of freedom of an object in the hand. That skill is often called “dexterous manipulation” or “internal manipulation” [58].

Two of these three elements appear as early as 1962 in the early work of Tomovic and Boni on the Belgrade hand [84]. They demonstrate the hand’s generality by emulating six of the seven grasp poses in a taxonomy developed by Schlesinger [73], including palmar, tip, pencil, cylindrical, ball and hook grasps, but not the lateral pinch. And while they do not explicitly appeal to anthropomorphism, the Belgrade Hand appears even to have been cast from a human hand, with veins visible on the back.

Okada’s work [65] almost twenty years later explicitly appealed to the example of the human hand, but Okada’s motivation and examples are more focused on the third element: *in-hand manipulation*. Salisbury [71] similarly focuses on in-hand manipulation. In Salisbury’s work, and in numerous subsequent papers, in-hand manipulation is accomplished by a fingertip grasp, using three fingers to control the motion of three point contacts. For full mobility of the grasped object, each finger requires three actuated degrees of freedom, imposing a minimum of nine actuators. While it has been noted that in-hand manipulation can be accomplished in other ways involving fewer actuators [11, 9], Salisbury’s approach is amenable to analysis, requiring only the resources of the hand itself, without depending on outside contact, gravity or controlled slip.

Grasp taxonomies and in-hand manipulation are coupled. Grasp taxonomies often identify two broad classes of grasps: power grasps (also called enveloping grasps), and precision grasps (also called fingertip grasps) [62]. In-hand manipulation is usually performed with precision grasps. Salisbury’s

design [71] was optimized in part for in-hand manipulation, using a specific fingertip grasp of a one inch diameter sphere. Both the DLR-Hand II [16] and the UPenn Hand [89] are designed with an additional freedom, allowing the hand to switch from a configuration with parallel fingers, better suited to enveloping grasps, to a configuration where the fingers converge and maximize their shared fingertip workspace, better suited to fingertip grasps.

Taken together, these three elements: anthropomorphism, grasp taxonomies, and in-hand manipulation, have enabled the development of arguably quite general hands. At the same time, they have driven us towards greater complexity.

One interesting departure is the work of Ulrich and colleagues [88, 86, 89, 87] leading to the UPenn hand and ultimately to the Barrett Hand [80]. Ulrich explicitly attacked the problem of trading off generality for simplicity. He defined simple hands as having one or two actuators; and complex hands as having nine or more actuators. He then defined a new class: medium-complexity hands, with three to five actuators. He achieved the reduction by explicitly eschewing in-hand manipulation, and focusing on a smaller set of stereotypical grasp poses, favoring enveloping grasps over fingertip grasps.

Ulrich's approach also employed underactuation to reduce the number of actuators while retaining generality. Each finger has two joints driven by a single actuator. The two joints are coupled to close at the same rate, but a clutch mechanism decouples the proximal joint when a certain torque is exceeded.

Another interesting departure is the work by Hirose and Umetani on the Soft Gripper [43, 44], which achieves a surprising 18 articulated coupled finger joints. Each finger is an open chain of ten links, coupled so that the entire chain is driven by just two actuators, and conforms to arbitrary shapes, even concavities.

Some more recent work has taken a different approach to generality, by directly testing and optimizing systems to a variety of environments, rather than referencing human hands or grasp taxonomies. Xu, Deyle and Kemp [93] directly address the requirements of the application (domestic object retrieval). Their design is based on the observation that the task often involves an isolated object on a flat surface, and is validated using a prioritized test suite of household objects [19]. Similarly, Ciocarlie and Allen's work [21] is tuned to obtain a collective optimum over a test suite of 75 grasps applied to 15 objects. Saxena, Driemeyer and Ng [72] likewise use a suite of common household objects. Their work was focused on vision-guided grasping of unfamiliar objects, but their success is also testimony to generality of the underlying hardware. Dollar and Howe [29] adopt a single generic shape (a disk) and explore variations in pose.

There is more to simplicity than just counting actuators. As actuators, transmissions, sensors and electronics improve, and with innovative design of hand structures and component integration, it has become possible to build more compact and more modular hands. For example, where the Utah / MIT Dextrous Hand [47] employed 32 remotely located actuators to control its 16 finger, thumb and wrist joints, the DLR-Hand II [16] has 13 actuators and associated electronics fully integrated into the hand itself.

2.1 Simple hands

Thus far we have looked primarily at complex hands, but simple hands also have a long history. We have noted that simple hands in industrial automation applications are often designed for a single shape, but there are numerous examples of designs for multiple shapes, or for multiple grasps of a single shape [32, 61].

Some of the earliest work in robotic manipulation exhibited surprisingly general manipulation with simple effectors. Freddy II, the Edinburgh manipulator, used a parallel-jaw gripper to grasp a variety

of shapes [3]. Freddy II is also one of the rare robotics research projects to address grasping in clutter. It used a camera to capture silhouettes of objects. If the objects were in a heap it would first look for a possibly graspable protrusion, then it would try to just pick up the whole heap, and if all else failed it would simply plow through the heap at various levels, hoping to break it into more manageable parts.

The Handey system also used a simple parallel-jaw gripper exhibiting impressive generality [53, 54]. It grasped a range of plane-faced polyhedra. Handey addressed clutter by planning grasp poses that avoided the clutter at both the start and goal, and planning paths that avoided the clutter in between. It also planned re-grasping procedures when necessary.

Theobald et al. [83] developed a simple gripper called “Talon” for grasping rocks of varying size and shape in a planetary exploration scenario. Talon’s design involved a single actuator operating a squeeze motion with three fingers on one side and two fingers on the other. The shape of the fingers, including serrations and overall curvature, as well as their compliant coupling, were refined to grasp a wide variety of rocks, even partially buried in soil.

Our gripper concept is similar to Hanafusa and Asada’s [41], who analyzed the stability of planar grasps using three frictionless compliant fingers. Our work is also close to theirs in our analysis of stability: stable poses correspond to local minima in potential energy.

Another area of relevant research is the design of modular work-holding fixtures, which was sometimes also applied to the design of simple hands for particular shapes, and to designs that accommodate several shapes, or several grasps of a given shape [13, 14, 90, 91, 79]. Broad arguments in favor of simplicity in industrial automation algorithms and hardware are advanced under the name “RISC Robotics” [17, 18].

2.2 Object recognition and localization

Others have explored haptic object recognition and localization. Lederman and Klatzky [52] survey work on human haptic perception, including object recognition and localization. Here we will borrow the biological terminology distinguishing between *kinesthetic* sensors such as joint angle, versus *cutaneous* sensors such as pressure or contact. The present work employs kinesthetic sensing along with knowledge of stable poses, both for recognition and localization.

Most previous work in robotic haptic object recognition and localization assumes contact location data—point data sampled from the object surface, perhaps along with contact normal [36, 39, 2]. While cutaneous sensing is the most obvious technique for obtaining contact location, it is also possible to obtain contact location from kinesthetic sensors using a technique called *intrinsic contact sensing*. If you know a finger’s shape and location, and the total applied wrench, and if you assume a single contact, then you can solve for the contact location. Bicchi, Salisbury and Brock [9] explored and developed the technique in detail. From a very general perspective, our approach is similar to intrinsic contact sensing. Both approaches use the sensed deformation of elastic structures in the hand, but our learning approach transforms that information directly to object recognition and localization, without the intermediate representation in terms of contact locations.

Our work fuses kinesthetic information with information from the system dynamics and controls, specifically the expectation of a stable grasp pose. Siegel [76] localized a planar polygon using kinesthetic data (joint angles) along with the knowledge that each joint was driven to some torque threshold. Our work could be viewed as a machine learning approach to the same problem, extended to arbitrary three dimensional shapes, using a simpler hand. Jia and Erdmann [49] fused contact data with system dynamics to estimate object pose and motion, and Moll and Erdmann [60] also fused haptic data with system dynamics and controls to estimate both the shape and the pose of a body in three

Table 1: Characterizing manipulation tasks and systems. A check broadly indicates either a task requirement or a hand capability. While P1 does produce stable grasps across a range of object shapes, we have not yet tested shape variation in combination with all the other requirements of bin-picking.

	Stability	Capture	In-hand manipulation	Clutter	Object shape variation	Multiple and deformable objects	Recognition, localization	Placing
Pickup Tool	✓	✓			✓	✓		
Bin picking	✓	✓		✓	✓		✓	✓
P1	✓	✓		✓	✓		✓	

dimensions.

Another related body of work is called “shape from probing” [24]. In its most common form, the problem is to determine the shape of a planar polygon by choosing a sequence of probes. For each probe, you aim a line at the polygon, and the probe returns some data, specifically the first intersection of the line with the polygon. Among the many variations on this problem, work by Wallack and Canny [92] and by Jia and Erdmann [48] are closest to our application.

Our approach is related to parts orienting research [40, 51, 37, 18, 33], where the ideas of using knowledge of stable poses and iterative randomized blind manipulation strategies are well known, and have even been used in the context of simple hands [56, 37].

3 Design Goals

This section tries to identify a list of characteristics of general-purpose grasping. We have tried to incorporate design criteria and task requirements most commonly used or referenced in the literature of hand design and grasping. The eight characteristics we identify—stability, capture, in-hand manipulation, clutter, object shape variation, multiple/deformable objects, recognition/localization and placing—are discussed, one by one, in the following subsections.

These eight characteristics are broadly defined properties that might be used to characterize either the requirements of an application, or the capabilities of a grasping hand. As an example, Table 1 characterizes the bin-picking task, the hardware store gripper, and our prototype gripper P1. By prioritizing the characteristics we would obtain a taxonomy of tasks and robotic systems. All eight characteristics have previously been studied, some more than others, but we are unaware of earlier attempts to organize them as a way of characterizing either tasks or robots.

Ultimately our aim is to refine these characteristics and to associate metrics with them, so that we can compare robot systems and tasks more precisely, and measure progress of the field in its pursuit of generality. As it stands, our list of characteristics is imprecise and incomplete, but still it facilitates discussion of general-purpose grasping.

The set of characteristics has been useful in framing our design goals. Following a similar process to Ulrich’s design of the UPenn Hand [89], we have identified certain capabilities we are willing to sacrifice for the sake of simplicity. Unlike Ulrich, we cannot describe our tradeoffs in terms of the well

known grasp taxonomies, or in terms of in-hand manipulation or even anthropomorphism. Bin picking has its own requirements and demands for specific tradeoffs. As a consequence, P1 cannot convincingly emulate any of the usual stereotype grasp poses, and cannot perform in-hand manipulation using just the fingers. Our grasps are neither power nor precision. The role of the fingers is tactile sensing and exploration as much as manipulation. Thus the requirements of the application do not align with previously listed design goals for hand design.

3.1 Stability

Stability can be broadly defined as holding the object without dropping it. Several more precise definitions are possible, but they depend on the process model and the nature of the disturbances. There is a substantial literature on stability, exploring several different aspects of grasp stability.

Our approach is similar to Hanafusa and Asada’s work [41], which analyzes the stability of planar grasps using three compliant fingers, where the fingers roll without friction on the object shape. Stable grasps correspond to minima in the potential field.

Some previous work assumes that the fingers are fixed, and make point contact with a rigid object, sometimes including friction. One can then either analyze the kinematic constraints to determine whether the object has any motion freedom, or analyze the system of contact forces to determine whether disturbance forces can be balanced. Many papers use the terms *force closure* and *form closure*, but with little agreement on their precise meaning [68, 59, 6, 8, 57].

Under the most liberal definitions of *force closure* and/or *form closure*, neither suffices for stability. However, Nguyen [64] showed that with a particular model of fingertip compliance, it is possible to obtain a stable grasp by assuming a high enough stiffness.

The grasps we explore in this paper do *not* satisfy the usual force/form closure tests. We use contact with the palm to yield a planar problem, nominally with three degrees of freedom, and then introduce frictionless point contacts with our three fingers. The analysis of [69, 70] would show that these grasps yield a second-order immobilization. Or, more simply, we can analyze stability by identifying local minima in the potential field.

The second-order nature of our grasps does have practical consequences: the effective stiffness is zero, meaning that disturbance forces including frictional forces may cause significant deviations from an expected stable pose, leading to localization noise. While the precision reported in Section 7 is adequate for some tasks, future designs might benefit from increasing the number of fingers, refining the finger shapes, or refining the palm shape for improved precision.

3.2 Capture

Capture means that the grasp process terminates with a stable grasp, despite variations in object pose. One way to characterize a grasp is to measure the set of all object poses from which the grasp succeeds, called the *capture region*. Despite the obvious importance of capture, there is relatively little work. Dollar and Howe [29] optimize their simple gripper design for capture. Monkman [61] notes the coupling between object pose uncertainty and gripper design in factory automation.

Such examples notwithstanding, capture has been more carefully studied in the context of planning and control. The capture region of a grasp defines the requirements of earlier operations, so it is fundamental to planning under uncertainty [55].

It is impossible to discuss capture without addressing some fundamental variations in modeling the grasp process. In the most general case, grasping is a very messy process, with one complex multi-body

system interacting with another, through contacts involving friction, impact, and material deformation. In order to model, analyze, control, and plan this process, simplifying assumptions are necessary.

The simplest set of assumptions is that the object is rigid, its shape and pose are known, and the robot has precise control of its hand. Then it is possible to choose contact locations, move the fingers there, and attain a stable grasp. The capture region is trivial in this case—just the singleton set comprising the object’s known pose.

More realistically, both the object pose and the robot control introduce uncertainty. It is still possible to simplify the analysis by proposing a reactive control strategy using tactile sensors to move the fingers to a stable grasp pose, without disturbing the object’s position [81].

So far we have only addressed *object priority* grasps: the object never moves; all the compliance is in the hand [35]. The other extreme is called a *hand priority* grasp: the hand’s motion is determined, and the compliance is all in the motion of the object.

Hand priority grasps arise in many situations where it may not even be possible to grasp an object without disturbing it. Dogar and Srinivasa [28] observe that it is often advantageous to move the object while grasping it, and develop techniques for predicting capture regions for a push-grasp. Brost and Christiansen [12] pursued a similar approach including an experimentally validated probability density function describing the likelihood of success as a function of object pose.

The present work employs grasps that are neither hand- nor object-priority. The object moves, but the fingers must move as well since it is their motion that gives us the information necessary for recognition and localization.

3.3 In-hand manipulation

In-hand manipulation means controlled motion of a grasped object relative to the hand. The origins of robotic in-hand manipulation, and its historical role in hand design, are addressed in Section 2.

Salisbury and Craig [71] motivated in-hand manipulation as a mechanism for obtaining *fine motion control*. The large motors used to move an arm, and the intervening masses, limit the arm’s role to *coarse motion*. But there are many other uses of in-hand manipulation. If the hand cannot adjust the grasp, it may have to put the object down and regrasp [53, 54]. In-hand manipulation can be important in manipulating multiple objects and deformable objects, as when dealing with coins or playing cards. In-hand manipulation is a source of precision, as when aligning the ends of a pair of chopsticks by pressing them against the tabletop and allowing them to slip in the hand.

Our bin-picking application doesn’t necessarily require in-hand manipulation. And although some bin-picking systems might use in-hand manipulation, our approach doesn’t, and in-hand manipulation played no role in the design of our prototype P1.

3.4 Clutter

We define *clutter* to mean anything that might limit access to the object. Clutter is an important element of both bin-picking and object retrieval, but in the course of this work we have realized that clutter is ubiquitous and fundamental. Clutter can affect the design of the robot, the choice of grasp, the robot path, and almost every aspect of a system. Previous work has seldom addressed clutter explicitly, but there are exceptions [72, 5, 3, 53].

First consider the effect of clutter on hand design. Suppose you are designing a hand from simple geometrical elements: planes, lines, or points. If you want to capture an isolated object, stabilize it, and estimate its location, infinite planes would be ideal. A plane sweeps out lots of space, and captures and

Elements	Application	
	Grasping	Clutter
Planes	○	×
Lines	△	△
Points	×	○

Table 2: Grasping versus clutter. The better an element is for grasping, the worse for avoiding clutter. ○ = good, △ = neutral, × = bad.

stabilizes objects very economically. Four planes in a tetrahedral configuration could squeeze an object of arbitrary shape and reduce the feasible poses to a small set. This idea is not entirely impractical. The tilted tray of Grossman and Blasgen [40] used three planes, with gravity instead of a fourth plane, to provide a universal parts orienting system. You might also view a common table as a variant of the idea: a single plane moving through space with a constant acceleration of 1g to sweep up all objects placed above it, in cases where only three degrees of constraint and localization are required.

However, if you are dealing with clutter, planes are terrible. They sweep up everything and avoid nothing. For clutter, points would be ideal. These observations are captured in Table 2.

Contemplation of Table 2 leads to an impractical approach: a set of levitated independently controlled points that drift through the interstices of the clutter, approach the object, and then switch to a coordinated rigid mode to lift the object. Consider the paths that the points follow to reach the object. If the clutter is unmoving, then those paths remain clear, and we could use hyper-redundant fingers, i.e. snakes or tentacles, to reach the object. Lifting the object clear of the clutter is still an issue, but the idea does suggest a practical compromise to address the problem of grasping in clutter: use very thin fingers, approaching the object along their lengths. The idea is reflected in many common manipulation tools, such as tweezers and laparoscopic forceps. You might even view the pickup tool (Figure 1a) as a variant of the idea, a single long thin element from which several fingers deploy. These tools are all designed for high-clutter environments: the pickup tool for retrieving small parts dropped into difficult to access places, and surgical forceps for the extreme clutter encountered in surgery.

The main point of Table 2 is that clutter and grasping are in opposition. Grasping is easy, if there is no clutter. The problem arises from the need to grasp one object while avoiding others. The problem is not just to grasp, but to grasp selectively.

Almost every grasping problem involves clutter. Even for an object isolated on a table, the table is clutter. Perhaps the most clutter-free grasping problem is a door knob: an affordance specifically designed for grasping, mounted on a stalk to minimize clutter. There are only a few cases involving less clutter, such as catching butterflies or cleaning a pool, where it is practical to sweep the entire nearby volume with a net.

3.5 Object shape variation

Object shape variation includes two different ideas: Shape diversity, and shape uncertainty. Shape diversity refers to a hand’s ability to grasp a variety of shapes. Shape uncertainty refers to a hand’s ability to grasp an object when the shape is not entirely known. For example, an adjustable wrench can handle any of several size bolts, but you need pliers to handle a bolt of unknown size.

Shape diversity has been a primary concern in the research on general-purpose hands. Salisbury and

Craig [71] motivated their work by two considerations: diversity of shape, and in-hand manipulation. Numerous papers have validated their hand designs by demonstrating grasping of a variety of shapes, much like our Figure 9, and also variation in scale [75]. Talon [83] was tested on a variety of rocks, and also had a system for rating the graspability of rocks. Standardized databases of shapes [19, 74] can provide a foundation for comparing grasping systems [93, 22].

3.6 Multiple and deformable objects

Multiple and deformable objects refer to all grasping tasks that go beyond rigid objects: bulk solids, liquids, multiple rigid objects, articulated objects, and soft objects. Most work on general-purpose manipulation neglects deformable objects but there are numerous interesting exceptions. Gopalakrishnan and Goldberg [38] extend the idea of form closure to deformable objects. Aiyama et al. [1] study grasping of multiple objects. Harada and Kaneko [42] address grasping and in-hand manipulation of multiple objects. Bell and Balkcom [4] address the problem of immobilizing non-stretchable polygonal cloth shapes.

3.7 Recognition, Localization

Recognition means identifying one of some given class of parts. Localization means estimating a grasped object’s pose relative to the hand. A related problem would be estimating the shape of a grasped object. All of these problems are well established in the robotics literature, independent of our focus on grasping and manipulation. Even if we restrict ourselves to haptic sensing, there is a well established literature. Besides the previous work outlined in Section 2, Natale and Torres-Jara [63] offer data on the use of touch sensors for object recognition.

3.8 Placing.

By “placing” we refer broadly to a variety of downstream requirements: dropping onto a conveyor, placing into a dishwasher, throwing, assembly, handing off to another (human or robotic) hand, and so forth. This paper does not address placing, but in a real bin-picking application, or in a domestic object retrieval application, placing is important.

3.9 Other issues and requirements

There are many different applications of grasping, with wildly varying requirements. Some applications have product inspection processes integrated with a hand, such as checking electrical continuity of an automotive part. Others have unusual force or compliance requirements, such as grasping an object for a machining operation.

We have also neglected numerous other practical requirements: specific requirements for speed, payload, weight, cost, robustness, suitability for special environments such as space or clean rooms.

4 Approach: Let the fingers fall where they may

This section outlines our approach, illustrated by a classic robotic manipulation problem: picking a single part from a bin full of randomly posed parts. Our approach is inspired by the pickup tool shown in Figure 1a, which is very effective at capturing one or several parts from a bin, even when operated

blindly. Rather than attempting to choose a part, estimate its pose, and plan for a stable grasp, we propose to execute a blind grasp, let the gripper and object(s) settle into a stable configuration, and only then address the problem of determining whether a single object was captured and estimating the object pose. Hollerbach [45] labeled the idea using very similar terms—“in which the details work themselves out as the hand and object interact rather than being planned in advance”—as *grasp strategy planning*, in contrast with *model-based planning* which computes exact grasp points based on prior knowledge of the object and environment.

The main problem addressed by this paper is to determine singletude and object pose within the grasp. We propose to use knowledge of stable grasp configurations to simplify both problems. The knowledge of those stable configurations is gained through a set of offline experiments that provide enough data to model the map between hand pose and object pose. This leads to a novel design criterion: to minimize the number of stable grasp poses, which ultimately has implications on both gripper design and gripper control.

Our initial gripper concept departs from the pickup tool that inspired it. For industrial bin picking the pickup tool has some deficiencies, including a tendency to capture several parts at a time, in unpredictable poses. And while the fingers of spring steel, with the bent tips, and the motion emerging along their lengths are all very effective and interesting, we want to begin with a design that is easier to analyze and simulate, and which supports estimation of pose. Our initial concept design has three rigid cylindrical fingers attached by revolute joints with encoders to a disk-shaped frictionless palm (Figure 1b). As with the pickup tool, all three fingers are compliantly coupled to a single actuator.

The initial gripper design is not a general-purpose gripper. It is a research prototype, designed to facilitate exploration of the “let the fingers fall where they may” approach, and to deal effectively with some other issues, such as penetrating a heap of parts in a bin. The design focuses on that approach, even to the detriment of generality. Future efforts may address the tradeoffs and compromises required to incorporate this approach into a true general-purpose gripper or hand.

The key elements of the approach are:

- Low-friction palm and fingers so that for irregular objects there are only a few stable grasp configurations;
- Blind or nearly blind grasping motions;
- Recognition of the presence of a single object, and pose estimation, from joint encoder values;
- Offline training either in simulation or in the real world;
- Iteration of a reach, grasp, withdraw, and classify strategy, terminating when a successful grasp of a single object in a recognized pose is achieved.

One important detail is the design of the compliant coupling of finger joints to actuator. It may seem that blind grasping of unknown shapes requires soft springs and large finger deflections. However, large finger deflections can yield unstable grasps. See Figure 2 for an example.

Rather than relying on large finger spring deflections to accommodate unknown shapes, we drive the actuator until some stall torque is exceeded. In our numerical studies of potential energy fields in the next section, we model this process by a series-parallel compliance. Three stiff finger springs in parallel are connected in series to a soft motor spring, as shown in Figure 3.

The remainder of the paper reports our numerical, analytical, and experimental results, aimed at evaluating the efficacy of the approach.

Figure 2: This simple example illustrates stability problems with large finger spring deflections. If the finger rest positions are at the disk center, the elastic potential energy field is a saddle. Adding two more fingers at top and bottom yields a constant potential field and a globally metastable disk pose. (Units are dimensionless here and throughout the paper, normalized so that the palm radius is 1 and the motor spring stiffness is 0.01.)

Figure 3: The diagram illustrates the series-parallel compliance scheme used to model the behavior of the grasp when we drive the motor until some stall torque is exceeded. The spring constant of the motor, k_m , and the spring constants of the fingers, k_f , are such that $k_f \gg k_m$.

5 Stable pose field

Unlike perhaps all previous work, we seek to *minimize* the number of stable grasp poses, in order to maximize the information implied by the existence of a stable grasp. Ideally there would be exactly one non-trivial local minimum in the elastic potential energy field. (By “non-trivial” we mean to exclude all out-of-reach poses with no finger contact.) With our prototype gripper P1 this ideal cannot be attained, due to the symmetric arrangement of the fingers. Whether there is any hand design, and any object, which attains the ideal of a single minimum is unknown. The related problem of an object with a single stable pose on a horizontal plane in a gravitational field has only recently been solved [31].

For the case of P1, the ideal of a single local minimum cannot be attained. But the question remains, how close can we get? In this section we example the stable pose fields of several different objects of varying scale.

Our model of P1 is a planar frictionless palm, and three frictionless line segment fingers, with the compliant coupling illustrated in Figure 3.

To calculate potential energy for a given shape at a given pose, the first problem is to determine which fingers are in contact, which can be formulated as a linear complementarity problem [67]. The naive algorithm is exponential in the number of contacts, but 2^3 is only 8, so we embrace naiveté. We examine every possible combination to find the one that yields positive finger deflections for the fingers in contact, and no penetration for the fingers not in contact. Given the corresponding finger and motor

spring deflections, potential energy is calculated by summing the elastic potential energy across all four springs.

Figure 4 shows the potential energy field for a sphere of varying scale, projected onto the $x - y$ plane. Symmetry dictates that the potential will be constant with rotation of the sphere, and stability with z translation is straightforward if the gravity vector is vertical (up or down) and the weight is small. It follows that the stable grasp determines translation but not rotation. The figure also shows that there are limits in scale. The small ball is stabilized only by assuming extremely small finger deflections, with correspondingly high stiffnesses.

Figures 5 and 6 show the potential field for a grasped cylinder. We examine only the planar configuration space for the cylinder, parameterized by (x, y, α) . (For the idealized model of P1, the only out-of-plane stable pose would be standing on end.) In Figure 5, the planar configuration space is projected onto various two-dimensional subspaces, so that the potential field can be displayed. Figure 6 shows some level sets of potential energy, plotted against the full three-dimensional planar configuration space. Due to the translational symmetry of the cylinder, local minima are not isolated points in the configuration space, but line segments in the $x-y$ plane, with constant α . The potential surfaces in Figures 5a and 5b are ruled surfaces, except when a finger contacts the end of the truncated cylinder. Each subfigure of Figure 5 includes a local minimum of the potential field, corresponding to one of the two most stable grasps of a cylinder: the cylinder aligned with the direction of one of the fingers or the cylinder perpendicular to one of the fingers. These are, in fact, the two most frequent poses observed in the experiments described in the next section.

The last example shape is a polyhedron (Figure 7). We examine a planar configuration space obtained by assuming one large face is in contact with the palm. The figure shows level surfaces of the potential field. As expected with an irregular object, the stable poses are isolated points (Figure 7).

It isn't known whether an irregular polyhedron exists that would not produce isolated stable poses. One illustrative example is an interesting singular shape developed by Farahat and Trinkle [34], which could exhibit planar motions while maintaining all three contacts without varying the finger angles, but Farahat and Trinkle's example does not correspond to a stable pose of P1.

6 Experimental prototype

This section describes the design and construction of P1, a gripper prototype (Figure 8) based on the proposed simple hand concept. The main purpose of the gripper is to assess the ideas presented in Sections 3, 4 and 5. The three main guidelines followed in the design are:

- A circular low friction flat palm. Avoiding friction is meant to produce fewer stable grasp poses and wider capture regions which ultimately facilitates grasp recognition and localization. The palm is covered with a thin sheet of Teflon.
- Three thin cylindrical fingers arranged symmetrically around the palm. Three is the minimum number of required frictionless contacts to produce a stable grasp in the plane for several objects.
- All three fingers are compliantly coupled to a single actuator. Between the actuator and each finger there is a torsion spring that gives the desired compliance to the gripper.

While we are interested in the behavior of the gripper across different scales, we chose the dimensions of P1 so that we could build it with off-the-shelf components. The palm was laser cut measuring 4 inches in diameter. Fingers are made out of 3/16 inch stainless steel rods measuring 2.5

Figure 4: Potential fields for small offsets. The lower half of the figure is zoomed. The spring deflections are $\pi/10$ for the large and medium spheres, $\pi/10,000$ for the small sphere.

Figure 5: 2D slices of the potential energy of the grasp of a cylinder. (a) X-Y Slice: Cylinder is parallel to one of the fingers. (b) X-Y slice: Cylinder is perpendicular to one of the fingers. (c) X- α slice: Y is constant and equal to zero and the cylinder moves in X and rotates from 0 to π .

Figure 6: Three isocontours of the potential energy of the grasp of a cylinder close to the minimum potential.

Figure 7: Potential function for the simple gripper grasping a prism formed by extruding a triangle with side lengths of 3, 4, and 5. Height of the prism is one. The plot (b) shows three closed contours enclosing isolated potential wells corresponding to three stable grasps obtained with the prism flat against the palm. There will be at least three more isolated stable grasps when the triangle is flipped over, and there may be other stable grasps we have not identified. The plot (c) is a closer view of one potential well, with some additional contours plotted.

Figure 8: Two views of P1, our implementation of the simple hand concept.

inches long. The gripper is actuated by a DC motor that, through a series of gears, transmits the power to the three fingers. The actuator is controlled open loop and driven to stall whenever we want to grasp. While all of our bin-picking experiments are with whiteboard markers, Figure 9 suggests that P1 will work with a variety of objects.

Experiments reveal some discrepancies between P1 and the idealized model described in Section 4. In order to avoid stall configurations of the hand with triple fingertip contact we introduced an offset of a few degrees between the resting positions of the three fingers. The knuckle or hub of the fingers also proved to be an interference during the grasping process.

P1 is a minimal implementation of the proposed simple gripper and only the first of a series of prototypes to come. Still, it has been useful in two ways. First, it has helped us to realize the importance of the mechanical design as part of the search for simplicity, in particular that we should also address complexity of fabrication. At the same time, P1 has allowed us to verify or refine some of the ideas that arise from the theoretical study of stability in Section 5.

Figure 9: P1 providing stable grasps for different objects.

7 Experimental results

In this section we test the performance of P1 in a bin-picking scenario (Figure 1b). We have selected bin picking because of the many challenges it presents, and because of its similarity to object retrieval for a domestic service robot. A hand that addresses bin picking, for a variety of different object shapes, might be a good candidate for retrieving an object from a cluttered kitchen drawer.

Bin picking is a challenging grasping application, combining high clutter with high pose uncertainty. It has been the focus of numerous research efforts for several decades, yet successful applications are rare [46, 82, 10]. Ironically, while we selected bin picking in part because of the high uncertainty, many researchers deprecate factory automation for the lack of uncertainty.

In our experiments the goal is to blindly grasp an object from a bin full of identical objects, and determine the outcome from sensor readings. We evaluate the performance of P1 in two subtasks:

- *Grasp Classification*: distinguish grasps that hold a single object from grasps that do not.
- *In-hand Localization*: if a single object is grasped, estimate its pose.

We attached P1 to a 6-axis industrial manipulator, and programmed it to move the gripper in and out of the bin full of markers. We hand labeled successful and unsuccessful grasps, providing ground

truth of both the number of markers and the pose of a single marker. We collected the data from 200 repetitions of the experiment. Figure 10 shows a small collection of representative grasp types.

Next we show the results of the experiments both in grasp recognition and localization.

7.1 Grasp Recognition: Object singulation

In this experiment we train a classifier to recognize the number of markers grasped, and evaluate the classifier’s accuracy.

We will call a grasp *successful* if it acquires a single marker and *failed* otherwise. Our primary concern is not the frequency of successful grasps, but the ability of the system to differentiate them. Table 3 shows the distribution of the number of markers grasped during the 200 trials. We hand-designed the path that the gripper follows to get in and out of the bin. To deal with clutter, while approaching the bin, the gripper oscillates its orientation along the vertical axis with decreasing amplitude. During departure the gripper vibrates to reduce the effect of friction and settle the object closer to the potential energy minimum.

0 markers	1 marker	2 markers	3 markers
57	83	43	17
(28.5%)	(41.5%)	(21.5%)	(8.5%)
	✓		

Table 3: Results of singulation experiments. Distribution of the number of markers grasped.

In order to model the map between sensor inputs and successful/failed grasps, we trained a Support Vector Machine with a Gaussian kernel [25, 15]. The classifier constructs a decision boundary in the sensor space—the three finger encoder values in the case of P1—by minimizing the number of misclassifications and maximizing the margin between the correctly classified examples and the separation boundary. Figure 11 shows the separation boundary found by the classifier.

To evaluate the performance of the system we randomly divide the dataset into two subsets of equal size, training and testing. We train the classifier with the first set and evaluate its accuracy by predicting success/failure on the grasps in the second set. The hyperparameters C and γ are tuned with leave-one-out cross validation on the training set. After several iterations of the randomly dividing the data, training the classifier and estimating its accuracy, the average accuracy obtained was 92.9%.

For the singulation system to be useful in a real bin-picking application it should be optimized to maximize *precision*—the ratio of true positives to those classified as positive—even to the detriment of *recall*—the ratio of true positives to the total number of positive elements in the dataset. Figure 12 shows the relationship between precision and recall, obtained by varying the relative weights of positive and negative examples when training the SVM. By choosing a working point with high enough precision, the system can potentially be turned into a rather slow, but accurate singulation system by rejecting all but the most confident outcomes.

7.2 In-hand Localization: Pose estimation

In the second part of the experiment we estimate marker pose. We limit the analysis to those grasps that have correctly isolated a marker, and suppose that it lies flat on the palm of the gripper.

The objective is to estimate the orientation of the marker, up to the 180 degree symmetry of the marker. We divide the set of successful grasps into training and testing sets and use Locally Weighted

Figure 10: Representative grasp types. The first row shows one-marker grasps. Within the first row, the second case has two point contacts and one line contact, while the first and third cases have three point contacts. The distinction between the first and the third cases is whether the fingers close where the cap meets the body of the marker. First and second grasp examples in the first row correspond to the two stable poses revealed by the stable pose field analysis in Figure 5 (a) and (b). The second and third rows show two- and three-marker grasps.

Figure 11: Perspective view and 2D projection, showing the decision boundary found by the Support Vector Machine in the space of finger encoders in degrees. Dots inside the closed region are grasps classified as successful.

Regression [23] with a Gaussian kernel. That is, the orientation of the marker is estimated as a weighted average of the orientation of the markers in the training set, where the weights vary exponentially with the distance between training and testing grasps in the input sensor space. Grasps that are closer in sensor space will have more weight when estimating the orientation.

Using leave-one-out cross validation, the average error is 13.0° . Figure 13 shows both the linear histogram of the error and its distribution in a polar chart.

As in the previous section, we can be more conservative and allow the system to reject outcomes if its confidence is low. We monitor the distance of any given testing grasp to the closest neighbors in the sensor space. Whenever that distance is too big, it means there is insufficient information to infer marker orientation confidently. By setting a threshold on that distance we can effectively change the tradeoff between average error and recall as shown in Figure 14.

By using an appropriate working point in the plot of Figure 14 we can lower the average error at the cost of having to discard not recognizable grasps. Figure 15 illustrates the result where we have selected a confidence level yielding an expected error of 8° . The left half of the figure shows the singulated grasps, and the right half shows the same grasps after a rotation to make the marker horizontal. Deviations from the horizontal correspond to errors in the regression.

8 Conclusion

Generality will always be out of reach. Even the most general known hand—the human hand—has to be augmented with many different tools to perform the tasks of interest to humans. Current

Figure 12: Precision/Recall plot of the singulation detection in the bin-picking application.

Figure 13: Linear and polar histogram of the errors in degrees obtained during the regression of the orientation of the marker.

autonomous robotic systems are at the opposite extreme, unable to exhibit even modestly general behavior, regardless of the potential generality of the hardware.

We have learned many lessons from our analysis, simulation, and experiments. There is a disparity between the performance of the pickup tool, which inspired our design, and the results we obtained with our design. It became evident from the first five minutes with our prototype that friction would be a problem, exacerbated by indentation of plastic or wood objects by our hard thin steel fingers. Further, the use of three fingers on smooth or convex polyhedral objects implies higher-order contact.

Figure 14: Tradeoff between the average error in estimating the orientation of the marker and the recall of a system.

Nonetheless the experimental results are encouraging. We were surprised by the rate at which the gripper acquired, recognized, and localized single objects.

This paper focused on a “grasp first, ask questions later” approach. Blind grasping, expectation based on training, and haptic sensing are combined to address a realistic manipulation problem with high clutter and high uncertainty. The approach leads in interesting and novel directions, such as using a slippery hand to minimize the number of stable poses. While slipperiness may not really be a good idea for a general-purpose hand, the insight is valid: sometimes stability and capture are *not* the highest priority. Broad swaths of stable poses in the configuration space can work against recognition and localization.

This suggests an important next step, which is to address another task domain. While we feel we have learned something about generality by addressing bin picking, it has led us in directions that might prove limiting in other task domains. One example is the use of frictionless hands, and our neglect of power grasps. While our experiments showed that capture and stability are adequate for the bin picking task, in other task domains that may not be so. As a first attempt to address other domains, we have begun to explore the utility of our approach with a Barrett Hand in a domestic object retrieval task.

There are several other directions for future work. Regarding the simple gripper design, the most obvious is suggested by using Table 1 to compare P1’s capabilities with the bin-picking requirements. P1 lacks a checkmark where bin-picking has one: placing. P1 needs a placing capability. Obviously it can just drop the part, but if the singulation and localization are to serve a purpose, it needs to place the part more carefully. Other directions for future work include morphology, including palm and finger form, and number of fingers.

Finally, a more objective and refined understanding of the dimensions of generality would help to make tradeoffs between generality and simplicity when designing a robotic hand. The field would benefit from a consensus on measures of generality to enable comparison between specific hand designs or task domains.

(a) Random subset of successful grasps

(b) Grasps after orientation correction

Figure 15: Orientation correction: Images of the grasps have been rotated based on their estimated orientation in order to orient the marker horizontally.

Acknowledgements

We wish to thank Mike Erdmann, Garth Zeglin, Amir Degani, Koen Ruskamp, Tomas Lozano-Perez, Tom Fuhlbrgge, David Bourne, and Mark Cutkosky.

References

- [1] Y. Aiyama, M. Minami, and T. Arai. Manipulation of multiple objects by two manipulators. In *Robotics and Automation, 1998. Proceedings. 1998 IEEE International Conference on*, volume 4, pages 2904–2909 vol.4, 16-20 1998.
- [2] P. Allen and K. Roberts. Haptic object recognition using a multi-fingered dextrous hand. In *1989 IEEE Int. Conf. on Robotics and Automation, 1989. Proceedings.*, pages 342–347, 1989.
- [3] A. P. Ambler, H. G. Barrow, C. M. Brown, R. M. Burstall, and R. J. Popplestone. A versatile computer-controlled assembly system. In *Third Int. Joint Conf. on Artificial Intelligence*, pages 298–307, 1973.
- [4] M. Bell and D. Balkcom. Grasping Non-stretchable Cloth Polygons. *The International Journal of Robotics Research*, 2009.
- [5] D. Berenson and S. Srinivasa. Grasp synthesis in cluttered environments for dexterous hands. In *IEEE-RAS Int. Conf. on Humanoid Robots*, December 2008.
- [6] A. Bicchi. On the closure properties of robotic grasping. *The International Journal of Robotics Research*, 14(4):319, 1995.
- [7] A. Bicchi. Hands for dexterous manipulation and robust grasping: A difficult road toward simplicity. *IEEE Tran. on Robotics and Automation*, 16(6):652–662, 2000.
- [8] A. Bicchi and V. Kumar. Robotic grasping and contact: A review. In *IEEE International Conference on Robotics and Automation*, 2000.
- [9] A. Bicchi, J. Salisbury, and D. Brock. Contact sensing from force measurements. *The Int. Journal of Robotics Research*, 12(3):249, 1993.
- [10] K. Boehnke. Object localization in range data for robotic bin picking. In *IEEE International Conference on Automation Science and Engineering, 2007. CASE 2007*, pages 572–577, 2007.
- [11] D. L. Brock. Enhancing the dexterity of a robot hand using controlled slip. In *IEEE Int. Conf. on Robotics and Automation*, pages 249–251, 1988.
- [12] R. Brost and A. Christiansen. Probabilistic analysis of manipulation tasks: A conceptual framework. *The International Journal of Robotics Research*, 15(1):1, 1996.
- [13] R. Brost and K. Goldberg. A complete algorithm for designing planar fixtures using modular components. *IEEE Transactions on Robotics and Automation*, 12(1):31–46, 1996.
- [14] R. C. Brost and R. R. Peters. Automatic Design of 3-D Fixtures and Assembly Pallets. *The International Journal of Robotics Research*, 17(12):1243–1281, 1998.
- [15] C. J. Burges. A Tutorial on Support Vector Machines for Pattern Recognition. *Data Mining and Knowledge Discovery*, 2:121–167, 1998.
- [16] J. Butterfass, M. Grebenstein, H. Liu, and G. Hirzinger. DLR-Hand II: Next generation of a dextrous robot hand. *IEEE Int. Conf. on Robotics and Automation*, 1:109–114, 2001.

- [17] J. Canny and K. Goldberg. A risc approach to robotics. *Robotics Automation Magazine, IEEE*, 1(1):26–28, mar 1994.
- [18] J. Canny and K. Goldberg. “RISC” industrial robotics: Recent results and open problems. In *IEEE Int. Conf. on Robotics and Automation*, pages 1951–1958, 1994.
- [19] Y. Choi, T. Deyle, T. Chen, J. Glass, and C. Kemp. A list of household objects for robotic retrieval prioritized by people with als. In *IEEE 11th International Conference on Rehabilitation Robotics*, pages 510–517, 2009.
- [20] Y. S. Choi, T. Chen, A. Jain, C. Anderson, J. Glass, and C. Kemp. Hand it over or set it down: A user study of object delivery with an assistive mobile manipulator. In *Robot and Human Interactive Communication, 2009. RO-MAN 2009. The 18th IEEE International Symposium on*, pages 736–743, September 2009.
- [21] M. Ciocarlie and P. Allen. Data-driven Optimization for Underactuated Robotic Hands. In *IEEE Int. Conf. on Robotics and Automation*, pages 1292–1299, May 3–8 2010.
- [22] M. T. Ciocarlie and P. K. Allen. Hand Posture Subspaces for Dexterous Robotic Grasping. *The International Journal of Robotics Research*, 28(7):851–867, 2009.
- [23] W. S. Cleveland and S. J. Devlin. Locally Weighted Regression: An Approach to Regression Analysis by Local Fitting. *Journal of the American Statistical Association*, 83:596 – 610, 1988.
- [24] R. Cole and C. Yap. Shape from probing. *Journal of Algorithms*, 8(1):19–38, 1987.
- [25] C. Cortes and V. Vapnik. Support-Vector Networks. In *Machine Learning*, pages 273–297, 1995.
- [26] M. Cutkosky. On grasp choice, grasp models, and the design of hands for manufacturing tasks. *IEEE Trans. on Robotics and Automation*, 5(3):269–279, 1989.
- [27] A. De Santis. *Modelling and Control for Human–Robot Interaction*. PhD thesis, Universita degli Studi di Napoli Federico II, November 2007.
- [28] M. Dogar and S. Srinivasa. Push-grasping with dexterous hands: Mechanics and a method. In *Proceedings of 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2010)*, October 2010.
- [29] A. Dollar and R. Howe. Towards grasping in unstructured environments: Grasper compliance and configuration optimization. *RSJ Advanced Robotics*, 19(5):523–543, 2005.
- [30] A. Dollar and R. Howe. Simple, robust autonomous grasping in unstructured environments. In *Robotics and Automation, 2007 IEEE International Conference on*, pages 4693–4700, 10-14 2007.
- [31] G. Domokos and P. Várkonyi. Geometry and self-righting of turtles. *Proceedings of the Royal Society B: Biological Sciences*, 275(1630):11, 2008.
- [32] J. Engelberger, D. Lock, and K. Willis. *Robotics in practice: management and applications of industrial robots*. Amacom, 1980.

- [33] M. Erdmann and M. Mason. An exploration of sensorless manipulation. *IEEE Journal of Robotics and Automation*, 4(4):369–379, 1988.
- [34] A. Farahat, P. Stiller, and J. Trinkle. On the geometry of contact formation cells for systems of polygons. *IEEE Tran. on Robotics and Automation*, 11(4):522–536, 1995.
- [35] R. Fearing. Simplified grasping and manipulation with dextrous robot hands. *IEEE Journal of Robotics and Automation*, 2(4):188–195, 1986.
- [36] P. Gaston and T. Lozano-Perez. Tactile recognition and localization using object models: The case of polyhedra on a plane. *IEEE Trans. Pattern Anal. Mach. Intellig.*, 6(3):257–265, 1984.
- [37] K. Goldberg. Orienting polygonal parts without sensors. *Algorithmica*, 10(2):201–225, 1993.
- [38] K. G. Gopalakrishnan and K. Goldberg. D-space and Deform Closure Grasps of Deformable Parts. *The International Journal of Robotics Research*, 24(11):899–910, 2005.
- [39] W. Grimson and T. Lozano-Pérez. Model-based recognition and localization from sparse range or tactile data. *The Int. Journal of Robotics Research*, 3(3):3–35, 1984.
- [40] D. Grossman and M. Blasgen. Orienting mechanical parts by computer-controlled manipulator. *IEEE Tran. on Systems, Man, and Cybernetics*, 5(5):561–565, 1975.
- [41] H. Hanafusa and H. Asada. Stable prehension by a robot hand with elastic fingers. In *Proc. Seventh Int. Symp. Industrial Robots*, pages 361–368, 1977.
- [42] K. Harada and M. Kaneko. Enveloping grasp for multiple objects. In *1998 IEEE International Conference on Robotics and Automation, 1998. Proceedings*, pages 2409–2415, 1998.
- [43] S. Hirose and Y. Umetani. The development of soft gripper for the versatile robot hand. *Mechanism and Machine Theory*, 13:351–359, 1978.
- [44] S. Hirose and Y. Umetani. The kinematics and control of a soft gripper for the handling of living or fragile objects. In *Proc. of the Fifth World Congress on the Theory of Machines and Mechanisms*, pages 1175–1178, 1979.
- [45] J. Hollerbach. Grasping in human and robot hands. *Vision and Action: The Control of Grasping*, pages 243–274, 1990.
- [46] K. Ikeuchi. Generating an interpretation tree from a CAD model for 3D-object recognition in bin-picking tasks. *International Journal of Computer Vision*, 1(2):145–165, 1987.
- [47] S. Jacobsen, J. Wood, D. Knutti, and K. Biggers. The Utah/MIT dextrous hand: Work in progress. *The Int. Journal of Robotics Research*, 3(4):21–50, 1984.
- [48] Y. Jia and M. Erdmann. Geometric sensing of known planar shapes. *The International Journal of Robotics Research*, 15(4):365, 1996.
- [49] Y. Jia and M. Erdmann. Pose and motion from contact. *The Int. Journal of Robotics Research*, 18(5):466, 1999.

- [50] C. Kemp, A. Edsinger, and E. Torres-Jara. Challenges for robot manipulation in human environments [Grand Challenges of Robotics]. *IEEE Robotics & Automation Magazine*, 14(1):20–29, 2007.
- [51] D. Kriegman. Let them fall where they may: Capture regions of curved objects and polyhedra. *The Int. Journal of Robotics Research*, 16(4):448–472, 1997.
- [52] S. Lederman and R. Klatzky. Haptic perception: A tutorial. *Attention, Perception & Psychophysics*, 71(7):1439, 2009.
- [53] T. Lozano-Perez, J. Jones, E. Mazer, P. O’Donnell, W. Grimson, P. Tournassoud, and A. Lanusse. Handey: A robot system that recognizes, plans, and manipulates. In *Robotics and Automation. Proceedings. 1987 IEEE International Conference on*, volume 4, pages 843 – 849, mar 1987.
- [54] T. Lozano-Pérez, J. Jones, P. O’Donnell, and E. Mazer. *Handey: A robot task planner*. MIT Press Cambridge, MA, USA, 1992.
- [55] T. Lozano-Perez, M. T. Mason, and R. H. Taylor. Automatic synthesis of fine-motion strategies for robots. *International Journal of Robotics Research*, 3(1), 1984.
- [56] M. Mason and J. Salisbury Jr. *Robot Hands and the Mechanics of Manipulation*. The MIT Press, Cambridge, 1985.
- [57] M. T. Mason. *Mechanics of Robotic Manipulation*. MIT Press, Cambridge MA, 2001.
- [58] C. Melchiorri and M. Kaneko. Robot hands. In *Springer handbook of robotics*, pages 345–260. Springer-Verlag New York Inc, 2008.
- [59] B. Mishra, J. T. Schwartz, and M. Sharir. On the existence and synthesis of multifinger positive grips. *Algorithmica*, 2(4):541–558, 1987.
- [60] M. Moll and M. Erdmann. Reconstructing the shape and motion of unknown objects with active tactile sensors. *Algorithmic Foundations of Robotics V, Springer-Verlag*, pages 293–309, 2004.
- [61] G. Monkman, S. Hesse, and R. Steinmann. *Robot grippers*. John Wiley and Sons, 2007.
- [62] J. Napier, J. Napier, and R. Tuttle. *Hands*. Princeton University Press, 1993.
- [63] L. Natale and E. Torres-Jara. A sensitive approach to grasping. In *Proceedings of the sixth international workshop on epigenetic robotics*, pages 87–94. Citeseer, 2006.
- [64] V. Nguyen. Constructing stable grasps. *The Int. Journal of Robotics Research*, 8(1):26–37, 1989.
- [65] T. Okada. Object-handling system for manual industry. *Systems, Man and Cybernetics, IEEE Transactions on*, 9(2):79–89, February 1979.
- [66] D. Osswald, J. Martin, C. Burghart, R. Mikut, H. Wörn, and G. Bretthauer. Integrating a flexible anthropomorphic, robot hand into the control, system of a humanoid robot. *Robotics and Autonomous Systems*, 48(4):213–221, 2004.
- [67] J. Pang and J. Trinkle. Complementarity formulations and existence of solutions of dynamic multi-rigid-body contact problems with coulomb friction. *Mathematical Programming*, 73(2):199–226, 1996.

- [68] F. Reuleaux. *The Kinematics of Machinery*. MacMillan, 1876. Reprinted by Dover, 1963.
- [69] E. Rimon and J. Burdick. Mobility of bodies in contact. I. A 2nd-order mobility index for multiple-finger grasps. *IEEE Tran. on Robotics and Automation*, 14(5):696–708, 1998.
- [70] E. Rimon, R. Mason, J. Burdick, and Y. Or. A general stance stability test based on stratified morse theory with application to quasi-static locomotion planning. *IEEE Tran. on Robotics*, 24(3):626–641, 2008.
- [71] J. Salisbury and J. Craig. Articulated hands: Force control and kinematic issues. *The Int. Journal of Robotics Research*, 1(1):4–17, 1982.
- [72] A. Saxena, J. Driemeyer, and A. Ng. Robotic grasping of novel objects using vision. *The Int. Journal of Robotics Research*, 27(2):157–173, 2008.
- [73] G. Schlesinger. Der Mechanische Aufbau der Kunstlichen Glieder. Ersatzglieder und Arbeitshilfen, part II, 1919.
- [74] P. Shilane, P. Min, M. Kazhdan, and T. Funkhouser. The princeton shape benchmark. In *Shape Modeling International*, pages 167–178, 2004.
- [75] T. Shirai, M. Kaneko, K. Harada, and T. Tsuji. Scale-dependent grasps. In *Proc. of the Int. Conf. on Advanced Mechatronics*, pages 197–202. Citeseer, 1998.
- [76] D. M. Siegel. Finding the pose of an object in a hand. In *1991 IEEE Int. Conf. on Robotics and Automation*, pages 406–411, 1991.
- [77] S. Srinivasa, D. Ferguson, C. Helfrich, D. Berenson, A. Collet, R. Diankov, G. Gallagher, G. Hollinger, J. Kuffner, and M. Weghe. HERB: a home exploring robotic butler. *Autonomous Robots*, 28(1):5–20, 2010.
- [78] C. Stanger, C. Anglin, W. Harwin, and D. Romilly. Devices for assisting manipulation: a summary of user task priorities. *IEEE Transactions on Rehabilitation Engineering*, 2(4):256–265, 1994.
- [79] A. Sudsang, N. Srinivasa, and J. Ponce. On planning immobilizing grasps for a reconfigurable gripper. In *Intelligent Robots and Systems, 1997. IROS '97., Proceedings of the 1997 IEEE/RSJ International Conference on*, volume 1, pages 106–113 vol.1, 7-11 1997.
- [80] B. Technologies. The Barrett hand. <http://www.barrett.com/robot/products-hand.htm>.
- [81] M. Teichmann and B. Mishra. Reactive Robotics I: Reactive Grasping with a Modified Gripper and Multifingered Hands. *The International Journal of Robotics Research*, 19(7):697–708, 2000.
- [82] R. Tella, J. Birk, and R. Kelley. General purpose hands for bin-picking robots. *Systems, Man and Cybernetics, IEEE Transactions on*, 12(6):828–837, 1982.
- [83] D. Theobald, W. Hong, A. Madhani, B. Hoffman, G. Niemeyer, L. Cadapan, J. Slotine, and J. Salisbury. Autonomous rock acquisition. In *Proceedings of the AIAA Forum on Advanced Developments in Space Robotics, Madison, Wisconsin*. Citeseer, 1996.

- [84] R. Tomovic and G. Boni. An adaptive artificial hand. *Automatic Control, IRE Transactions on*, 7(3):3–10, 1962.
- [85] J. P. Trevelyan, S.-C. Kang, and W. R. Hamel. Robotics in hazardous applications. In B. Siciliano and O. Khatib, editors, *Handbook of Robotics*, chapter 48. Springer, 2008.
- [86] N. Ulrich and V. Kumar. Grasping Using fingers with Coupled joints. In *Proceedings of the 20th Biennial ASME Mechanisms Conference*, Orlando, Florida, September 1988.
- [87] N. Ulrich, V. Kumar, R. Paul, and R. Bajcsy. Grasping with mechanical intelligence. In *Proc. CISM-IFTOMM VIIIth Symposium on Robots, Manipulators, and Systems. Cracow, Poland*, 1990.
- [88] N. Ulrich, R. Paul, and R. Bajcsy. A medium-complexity compliant end effector. In *1988 IEEE Int. Conf. on Robotics and Automation, 1988. Proceedings.*, pages 434–436, 1988.
- [89] N. T. Ulrich. Grasping with mechanical intelligence. MS-CIS-89-51, University of Pennsylvania, 1989.
- [90] A. Wallack and J. Canny. Planning for modular and hybrid fixtures. In *Robotics and Automation, 1994. Proceedings., 1994 IEEE International Conference on*, pages 520–527 vol.1, 8-13 1994.
- [91] A. Wallack and J. Canny. Modular fixture design for generalized polyhedra. In *Robotics and Automation, 1996. Proceedings., 1996 IEEE International Conference on*, volume 1, pages 830–837 vol.1, 22-28 1996.
- [92] A. S. Wallack and J. F. Canny. Generalized Polyhedral Object Recognition and Localization Using Crossbeam Sensing. *The International Journal of Robotics Research*, 16(4):473–496, 1997.
- [93] Z. Xu, T. Deyle, and C. Kemp. 1000 trials: An empirically validated end effector that robustly grasps objects from the floor. In *Robotics and Automation, 2009. ICRA '09. IEEE International Conference on*, pages 2160–2167, May 2009.
- [94] B. Yamauchi. PackBot: A versatile platform for military robotics. In *Proceedings of SPIE*, volume 5422, pages 228–237. Citeseer, 2004.
- [95] L. Zollo, S. Roccella, E. Guglielmelli, M. Carrozza, and P. Dario. Biomechatronic design and control of an anthropomorphic artificial hand for prosthetic and robotic applications. *IEEE/ASME Transactions on Mechatronics*, 12(4):418–429, 2007.