Autonomous Manipulation with a General-Purpose Simple Hand*

Matthew T. Mason, Alberto Rodriguez, Siddhartha S. Srinivasa and Andrés S. Vázquez

April 24, 2012

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 an approach for autonomous manipulation using a general-purpose but simple hand. We explore the approach in the context of a bin-picking task, focused on 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 describe experiments with two prototype hands to perform bin-picking of highlighter markers.

1 Introduction

Complex hands surely offer greater generality than simple hands. Yet simple hands such as the prosthetic hook have demonstrated a degree of generality yet untapped by any autonomous system, either with simple or complex hands. The goal of our research is to develop general-purpose autonomous manipulation with simple hands. Our primary motive is that simple hands are easier to study and to understand. Study of simple hands may more quickly yield insights leading to autonomous general-purpose manipulation. A secondary motive is that simple hands are often more practical. Simple hands are smaller, lighter, and less expensive than complex hands. Robots may require simple hands for the indefinite future for some applications such as micro-manipulation or minimally invasive surgery, just as humans use simple tools for many tasks.

How do we define simplicity and generality? By a *simple hand* we mean a hand with few actuators and few sensors, and with economically implemented mechanisms, so that the whole hand can be small, light, and inexpensive. Figure 1 shows two examples: a simple pickup tool,

Matthew T. Mason, Alberto Rodriguez and Siddhartha S. Srinivasa are with the Robotics Institute at Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA <albertor@cmu.edu, matt.mason@cs.cmu.edu, siddh@cs.cmu.edu>

Andrés S. Vázquez is with the Escuela Superior de Informatica at Universidad de Castilla-La Mancha, Paseo de la Universidad 4 Ciudad Real, 13071 Spain <andress.vazquez@uclm.es>

This paper is a revision of papers appearing in the proceedings of the 2009 International Symposium on Robotics Research (ISRR) and the 2010 International Symposium of Experimental Robotics (ISER).

*The International Journal of Robotics Research Vol. 31, No. 5 April 2012, pp. 688-703.

and P2, 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, and not be tuned for a specific one.

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) Bin-picking scenario and prototype gripper P2 with four fingers and angle encoders for object recognition and localization.

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 more complex and more 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 not yet achieved in autonomous robotic systems.

Autonomous manipulation with simple hands requires a new approach to grasping. Past manipulation research often adopts an approach we will call: “Put the fingers in the right place”. The robot plans a desired stable grasp configuration, including all finger contact locations, and then drives the fingers directly to those contacts. It is a direct approach to the grasping problem, but often impractical. It assumes the object shape and pose are accurately known, that a suitable grasp is accessible, and that the object does not move during the grasp.

This paper adopts a different approach: “Let the fingers fall where they may”. The robot plans a motion with the expectation that it will lead the object to a stable grasp. This expectation anticipates uncertainty in object shape and initial pose, as well as object motion during the grasp. Because there might be several possible configurations for the resting pose of the object, the hand resolves that uncertainty with sensing, during and after the grasp. Thus the approach can also be summarized as, “grab first, ask questions later”.

Machine learning is central to our approach. The grasping process is well beyond our ability to model and analyze, when the object is allowed to move and interact with fingers and with clutter, and when the object shapes and initial poses are noisy. Nonetheless, there may be enough structure to choose promising grasp motions based on experience.

Our approach was inspired by the simple pickup tool of Figure 1a. The pickup tool can grasp various shapes without careful positioning of fingers relative to the object. Our first gripper prototypes P1 and P2 in Figure 1b mimic the pickup tool but also address 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 near-blind grasping strategy of the pickup tool, augmented by simple sensors and offline training to address recognition and localization.

We test the approach on a highlighter bin-picking task: Given a bin full of randomly placed highlighters, grasp a single object and accurately estimate its pose in the hand. Bin-picking is a high clutter, high uncertainty task, with a target rich environment that simplifies the experimental testing of our approach. Bin-picking is also suited to a learning approach, since failure and iteration may be tolerated in order to produce singulated objects with great reliability. Grasp recognition, in this task, means recognizing the presence of a single marker rather than several markers or no markers. Localization means estimating the orientation of the marker flat against the palm up to the 180° near-symmetry.

Grasp recognition and localization are the product of two processes. First, the hand is designed so that irregular objects tend to fall into one of just a few discrete stable poses. Second, a modest set of sensor data is used to estimate the pose of the object, or reject the grasp. Thus a key element of the paper, addressed in Section 4, is to explore the distribution of stable poses in the object configuration space.

The main experimental results are presented in Section 6. The grasp recognition and localization systems require offline training comprising 200 random grasps, with the results presented to a machine vision system which determines ground truth. The ultimate result is that our approach can acquire singulated highlighters with error rates as low as we are able to measure, while estimating orientation with expected error of 8° or less.

The proposed approach and hand design is general-purpose in the sense that it is not specialized to any particular object shape. Our experiments use highlighter markers, but their use was not even contemplated until after the first hand design. The highlighter markers are small for the hand, hence the tendency to grasp more than one at a time. The main principle, employing stable poses isolated in the sensor space, is best suited to irregular shapes, not cylinders. Thus, prototypes P1 and P2 are general-purpose hands, although the extent of their *generality* over tasks is a harder question out of the scope of this paper.

2 Previous Work

This paper incorporates results in [46, 59]. Additional discussion of grasp characteristics and clutter is available in [47]. Related work on design of finger phalange form appears in [58].

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 [68] 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 [38] 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 (or two such hands in the case of bi-manual manipulation) rather than a toolbox of highly specialized grippers. Second, while we agree that complex hands can be used to emulate simpler effectors, we work directly with simpler effectors designed specifically to explore the issues.

2.1 Complex Hands

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

- **Anthropomorphism.** There are many reasons to emulate 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 aesthetic 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 [6, 22, 38, 54, 76].

- **Grasp taxonomies.** Rather than directly emulating the human hand, several hand designs emulate the poses taken from taxonomies of human grasp [21, 62, 51].
- **In-hand manipulation.** Controlled motion of an object using the fingers, often called “dexterous manipulation” or “internal manipulation” [48].

Two of these three elements were already considered in 1962 with the Belgrade hand [68], which appears to have been cast from a human hand, and which emulated six of the seven grasp poses in Schlesinger’s taxonomy [62].

Okada’s work [53] 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 approach to generality: *in-hand manipulation*. Salisbury [60] 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. There are other approaches to in-hand manipulation, involving fewer actuators [9, 7], but Salisbury’s approach requires 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) [51]. In-hand manipulation is usually performed with precision grasps. Salisbury’s design [60] was optimized in part for in-hand manipulation, using a specific fingertip grasp of a one inch diameter sphere. Both the DLR-Hand II [14] and the UPenn Hand [69] 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.

2.2 Simple Hands

Thus far we have looked primarily at complex hands, but simple hands also have a long history. 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 [2].

The Handey system also used a simple parallel-jaw gripper exhibiting impressive generality [43, 44]. It grasped a range of plane-faced polyhedra.

Our gripper concept is similar to Hanafusa and Asada’s [33], 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.

Theobald et al. [67] 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.

Simple hands are widely used in industrial automation. Hands 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 [26, 50].

Our work is related to other industrial automation problems, such as the design of work-holding fixtures [10, 11, 73, 74, 65] and parts orienting [32, 40, 30, 16, 27], 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 [45, 30]. Broad arguments in favor of simplicity in industrial automation algorithms and hardware are advanced under the name “RISC Robotics” [15, 16].

Some recent work inspired by service applications has addressed generality with simple designs, by directly testing and optimizing systems to a variety of environments, rather than referencing human hands or grasp taxonomies. Xu, Deyle and Kemp [75] 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 [17]. Similarly, Ciocarlie and Allen’s work [18] is tuned to obtain a collective optimum over a test suite of 75 grasps applied to 15 objects. Saxena, Driemeyer and Ng [61] 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 [24] adopt a single generic shape (a disk) and explore variations in pose.

Underactuation is an interesting way to address generality and simplicity. One may achieve some of the advantages of having several degrees of freedom, while still retaining a small number of motors. Our prototypes P1 and P2 use underactuation to drive three and four fingers, respectively, through a single motor. Hirose and Umetani [34, 35] designed a soft gripper, controlling as many as 20 freedoms with just two motors.

Dollar et al. [24] demonstrates a simple planar two-fingered gripper, with two compliantly-coupled joints per finger, and explores 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 [25].

Brown et al. [12] takes underactuation to an extreme. An elastic bag containing a granular material such as coffee grounds can be switched from a deformable state to a rigid state by applying vacuum—a virtually infinite number of degrees of freedom actuated by a single motor. The device can be used as a gripper which closely conforms to a broad range of shapes and grasps them securely. It is an interesting case for any discussion of generality and simplicity.

2.3 Intermediate Complexity Hands

One interesting entry in the discussion of simple and complex hands is the work of Ulrich and colleagues [72, 70, 69, 71] leading to the UPenn hand and ultimately to the Barrett Hand [4]. 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

	Stability	Capture	In-hand manipulation	Object shape variation	Multiple and deformable objects	Recognition, localization	Placing	Clutter
Bin picking	✓	✓		✓		✓	✓	✓
Pickup Tool	✓	✓		✓	✓			
P1	✓	✓		✓		✓		✓
P2	✓↑↑	✓↑↑		✓		✓		✓↓↓

Table 1: Dimensions of general-purpose grasping to characterize manipulation tasks and systems. A check broadly indicates either a task requirement or a hand capability. ↑↑ indicates an improvement of P2 with respect to P1, and ↓↓ otherwise.

stereotypical grasp poses, favoring enveloping grasps over fingertip grasps.

The UPenn and Barrett designs also use underactuation. 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.

Laliberté et al. [41] also develop designs based on underactuated fingers, with two or three freedoms actuated by a single motor, culminating in a 10 degree-of-freedom hand with just two motors, and the commercial Adaptive Gripper [57].

2.4 Dimensions of General-Purpose Grasping

In [47] the authors propose a list of eight characteristics of general-purpose grasping to be used to characterize either the requirements of an application or the capabilities of a hand: stability, capture, in-hand manipulation, object shape variation, multiple/deformable objects, recognition/localization, placing and clutter. In Table 1 we make use of that set of general-purpose *dimensions* to compare the requirements of bin-picking and the designs of the “pickup tool”, and the prototype grippers described in this paper P1 and P2.

In particular, *clutter* is a key characteristic of the bin-picking task, and *recognition and localization* are key capabilities of the “let the fingers fall where they may” approach. This section reviews previous work in clutter, recognition, and localization.

Previous work has seldom addressed clutter explicitly, but there are exceptions. Freddy II [2] 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. Handey [43] 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. Berenson and Srinivasa [5] developed an algorithm for planning stable grasps in cluttered environments, and Dogar and Srinivasa [23] use pushing to avoid clutter.

Many others have explored haptic object recognition and localization. Lederman and

Klatzky [42] 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, sometimes including contact normal, sampled from the object surface [29, 31, 1]. 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 [7] 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 [63] 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 [39] fused contact data with system dynamics to estimate object pose and motion, and Moll and Erdmann [49] also fused haptic data with system dynamics and controls to estimate both the shape and the pose of a body in three dimensions. Natale and Torres-Jara [52] show an example of the use of touch sensors for object recognition.

3 Approach: Let the Fingers Fall Where they May

This section outlines our approach to grasping, illustrated by a classic robotic manipulation problem: picking a single part from a bin full of randomly posed parts. The key elements of the approach are:

- **Simple control strategy:** In the traditional approach to grasping, robotic hands try to “put the fingers in the right place”, which means driving the fingers directly to the ultimate desired stable grasp pose. This assumes that the object shape and pose are accurately known, that the grasp pose is accessible, and that the object will not move during the grasp. Instead, we “let the fingers fall where they may”, which means using a grasp motion chosen so that the gripper and object settle into a stable configuration. The object shape and pose need not be precisely known, and the approach accommodates motion of the object, even in clutter. Hollerbach [36] describes the idea thus: “in which the details work themselves out as the hand and object interact rather than being planned in advance” and calls the approach *grasp strategy planning*, in contrast with *model-based planning* which computes exact grasp points based on prior knowledge of the object and environment.
- **Simple mechanism:** The simple control strategy encourages a simple hand design.

The hand does not need several degrees of freedom per finger as in the traditional approach. We adopt a gripper concept inspired by the pickup tool in Figure 1a, which is very effective at capturing parts from a bin, even when operated blindly. Gripper prototypes P1 and P2 in Section 5 have low friction palm and fingers so that, for irregular objects, there are only a few stable grasp configurations. When a single object is captured, we expect the fingers to drive the object to one of those stable configurations.

- **Statistical model of grasp outcome:** We learn a data-driven model of the relationship between kinesthetic sensor feedback and grasp outcome. We introduce the concept of *grasp signature* to refer to the time history of the entire grasp process as perceived by the hand’s own sensors. The proposed gripper design simplifies the prediction of grasp outcome from grasp signature: By reducing the number of stable poses, in-hand object localization requires minimal sensing.
- **Iteration:** To address the stochastic nature of our approach the robot iteratively grasps and classifies, terminating when a single object is captured in a recognized pose.

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 from kinesthetic sensors to object pose. This leads to a novel design criterion—to minimize the number of stable grasp poses—which ultimately has implications both for gripper design and gripper control.

Our initial gripper concept departs from the pickup tool that inspired it. The basic concept is a planar disk-shaped frictionless palm, with rigid cylindrical fingers evenly spaced around the palm (Figure 1b). The fingers are attached by revolute joints with encoders. The fingers are compliantly coupled to a single motor. The design is generic and simple, manifestly not guided by the geometry of any particular object to be grasped. It is also easy to simulate and analyze.

The compliant actuation scheme for our prototype grippers is illustrated in Figure 2. The fingers are coupled through linear springs to a motor which is driven to a predetermined stall torque τ_m . Variation in object shape is accommodated mostly by motor travel, rather than by finger joint compliance. Softer finger springs would be an alternative way of accommodating varying sizes, and would have the additional advantage of being more sensitive to contact forces. Unfortunately, excessively floppy fingers sometimes yield no stable grasps at all.

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

4 Distribution of Stable Poses

This section focuses on the set of feasible stable configurations of hand and object, and their distribution in the object configuration space. If there are only a few different stable poses, well separated in the configuration space, then it is likely that very little sensor data is required to localize the object.

One consequence of this approach is that by *minimizing* the number of stable grasp poses, we maximize the information implied by the existence of a stable grasp. The ideal case would

Figure 2: The diagram illustrates the parallel compliance-actuation scheme used to model hand/object interaction. Units are dimensionless throughout the analysis in Section 4 so that the palm radius is 1 and the constant of finger springs is $k_f = 1$. When closing the hand, the motor is driven to a stall torque τ_m .

be a single stable pose with a large capture region. Because of the symmetric design of our grippers, that ideal will never be attained. Symmetric objects also depart from the ideal, giving rise to continua of quasi-stable poses, and corresponding pose ambiguities. Friction and sensor noise also add to the difficulties. In practice, even with asymmetric objects, equivalent grasps of an object will be observed as a cluster rather than a point in the sensor space, and precision is compromised.

We model the handling forces by a potential field, following the example of [33]. Assuming some dissipation of energy, stable poses correspond to minima in the potential field. We can also get some idea of localization precision by examining the shape of the potential field. V-shaped wells, or deep narrow wells, are less susceptible to noise than broad shallow U-shaped wells.

Our model of the hand is depicted in Figure 2. Fingers are modeled as lines, and finger-finger interactions are not modeled. For n fingers there are n springs and the motor, giving $n + 1$ sources of potential energy which account for the total energy of the system U . We assume the motor is a constant torque source τ_m , corresponding to a potential $U_m = \tau_m \cdot \theta_m$, where θ_m is motor position. The finger spring rest positions are also given by θ_m , so each finger potential is given by $U_i = k_f(\theta_i - \theta_m)^2/2$, where k_f is the finger stiffness and θ_i is the finger angle. The total potential energy is:

$$U = U_m + \sum_{i=1}^n U_i = \tau_m \cdot \theta_m + \sum_{i=1}^n \frac{1}{2} k_f (\theta_i - \theta_m)^2 \quad (1)$$

By examining the distribution of local minima in that potential field, and the shape of the potential wells, we seek some insights into hand design, finger stiffness, and choice of stall torque.

4.1 Examples of Potential Fields

This section shows the potential energy and corresponding stable poses for three objects: a sphere, a cylinder and a polyhedron; both with the three and four fingered prototype grippers. To calculate the potential energy for a given object geometry and pose, the first step is to determine the motor angle θ_m and finger angles $\theta_1 \dots \theta_n$, yielding a linear complementarity problem [55]. Appendix A describes the solution in detail.

Figure 3: Potential field of a sphere grasped by a (a) three-fingered and (b) four-fingered versions of the proposed simple hand. The plots illustrate the variation of the potential energy of the grasp with translation of the sphere in the $x - y$ plane. The radius of the sphere is 0.5 while the radius of the palm is 1. The hand is driven to a stall torque of $\tau_m = 0.1$. The contour plots illustrate that $(0, 0)$ is an isolated stable pose for both grippers.

Figure 3 shows the potential field of a sphere of radius equal to half the radius of the palm, projected onto the $x - y$ plane. Because of symmetry, x and y are the only coordinates of the pose of the sphere meaningful to the grasping process, and the only ones that can be derived from the knowledge of a stable pose.

The plots in Figure 3 present a unique stable grasp of the sphere both for the three-fingered and four-fingered cases. With some object shapes, the addition of the fourth finger should “sharpen” the bottom of the potential well, increasing the stiffness of the grasp and adding precision to both grasp recognition and pose estimation. As the plot illustrates, this is not the case with the sphere. However, the global structure of the potential field is altered when adding the fourth finger, yielding a somewhat larger basin of attraction.

Figure 4 shows the potential field of a cylinder with both the three- and four-fingered hands. Cylinder location is represented by (r, α) , where r is the minimum distance from the cylinder axis to the palm center and α is the axis angle relative to the x axis. Translation of the cylinder along its length is unobservable, and therefore not meaningful for this analysis. As a consequence, each local minimum corresponds to a continuum of quasi-stable poses in the $x - y$ plane that allow translation along the cylinder’s axis. Figure 4a shows six stable poses for the case of the three-fingered gripper and Figure 4b shows four for the case of the four-fingered one, corresponding to the different interlocked configurations of fingers and cylinder. These are, in fact, the most frequent poses observed in the experiments described in Section 6.

Figure 4: Potential field of a cylinder grasped by a (a) three-fingered and (b) four-fingered version of the proposed simple hand. The plots illustrate the variation of the potential energy of the grasp with displacements of the cylinder in the $r - \alpha$ space, where r and α are the radial coordinates of the axis of the cylinder from the center of the palm. The hand is driven to a stall torque of $\tau_m = 0.1$. The contour plots yield six and four stable poses for the three- and four-fingered simple hands respectively.

Figure 5: Potential field of a scaled 3-4-5 polyhedron grasped by a (a) three-fingered and (b) four-fingered version of the proposed simple hand. The plots illustrate the variation of the potential energy of the grasp with displacements of the polyhedron in the $x - y$ plane, while holding orientation constant. The hand is driven to a stall torque of $\tau_m = 0.2$.

Figure 5 shows the potential field of a scaled 3-4-5 polyhedron, with both the three- and four-fingered hands. Assuming that a triangular face lies flat on the palm, the set of possible displacements is three-dimensional, and cannot be reduced to a two-dimensional plot as we did for the sphere and the cylinder. Figure 5 shows a slice of the potential energy where the orientation of the polyhedron is held constant. For that specific orientation, the hand yields a stable pose.

As expected with an irregular object, the stable poses of the polyhedron are isolated points. 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 [28], which would 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 the proposed simple hand.

4.2 Stall Torque and Stability

The proposed strategy to grasp an object consists of closing the hand by driving the motor to stall and letting hand and object settle into a stable configuration. In this section we show that stability of that process depends on motor stall torque.

While it might seem intuitive that the stronger the grasp, the more stable it is, the reality is not so simple. In the compliant actuation scheme in Figure 2, high motor torque implies

Figure 6: Potential fields for a sphere held by P1. Stall torque increases by a factor of 10 from 0.01 at the top to 10 at the bottom. The grasps are stable for lower values of motor torque, and unstable for higher values.

high spring compression which might yield unstable grasps, as shown by Baker, Fortune and Grosse [3] in the context of the Hanafusa and Asada hand [33].

To illustrate this effect we show in Figure 6 the potential field of a three-fingered hand grasping a sphere for four increasing values of the stall torque. Only the first two cases yield stable grasps.

5 Experimental Prototypes

This section describes the design and construction of prototypes P1 and P2 (Figure 7 and Figure 8). The main purpose is to explore the ideas in Section 3, in particular the “let the fingers fall where they may” approach, and at the same time deal with the particular constraints of the bin picking task such as singulation of objects and heap penetration. The three main guidelines followed in the design are:

Figure 7: Side and frontal view, and transmission mechanism of gripper prototype P1.

- A circular flat low friction palm. Avoiding friction is meant to produce fewer stable grasp poses and yield wider capture regions. For both prototypes, the palm is covered with a thin sheet of Teflon.
- Thin cylindrical fingers arranged symmetrically around the palm.
- All fingers are compliantly coupled to a single actuator.

While we are interested in the behavior of the grippers across different scales, we chose the dimensions so that we could build them mostly with off-the-shelf components. The highlighter markers were selected because they are inexpensive, readily available in large numbers, and about the right size. In fact they are a bit small, but that suits the experiment since it enables the hand to grasp several at a time.

The palm was laser cut measuring 2 inches in radius. Fingers are made out of 3/16 inch stainless steel rods measuring 2.5 inches long. Our desire to minimize friction met with limited success. Bench tests yielded a coefficient of friction of 0.13 (friction angle 7°) between palm and marker, and a coefficient of friction of 0.34 (friction angle 19°) between finger and marker.

Prototype P1 (Figure 7) has three fingers. The gripper is actuated by a DC motor that transmits the power to the fingers through a geartrain. The actuator is controlled open loop and driven to stall. Torsional springs coupling the fingers with the gear assembly introduce compliance which allows for moderate conformability of the hand. While all of our bin-picking experiments are with highlighter markers, our analysis in Section 4 with different object shapes suggests that the gripper will work with a variety of objects.

Prototype P2 (Figure 8) has four fingers. The actuation is transmitted through a leadscrew connecting the motor to an individual linkage for each finger. The linkage has been optimized to maximize the stroke of the fingers and to yield a nearly uniform transmission ratio from leadscrew to finger rotation. One link in each finger linkage is elastic and provides compliance to the gripper.

Figure 8: Side and frontal view, and transmission mechanism of gripper prototype P2.

As explained in Section 4.1, owing to the fourth finger, P2 has a theoretical advantage both in its capture region and grasp stability, as measured by the basin of attraction. However, we can also expect it to perform worse in the presence of clutter. And while P2 might have improved recognition and localization for some objects, we shall see in Section 6 that for highlighter markers the performance is worse, which we attribute to the “self-clutter” effect: fingers interfere more often with each other when the hand has four fingers than when it has three.

Sensing the state of the hand is key if we want to “grasp first and ask questions later”, so P2 has absolute encoders on each finger and the actuator. Figure 9 shows an example of finger signatures for successful and failed grasp attempts.

P1 and P2 are minimal implementations of the proposed simple gripper and only the first of a series of prototypes to come. Still, they have been useful in two ways. First, they have 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, both grippers have allowed us to verify or refine some of the ideas that arise from the theoretical study of stability in Section 4.

6 Experimental Results

In this section we describe the implementation and results obtained in our approach to the bin-picking problem. Bin-picking is characterized by high clutter and high pose uncertainty, making it a challenging task for the conventional model-driven “put the fingers in the right place” approach. It has been the focus of numerous research efforts for several decades, yet successful applications are rare [37, 66, 8]. As we shall see the “let the fingers fall where they may” approach handles high clutter and pose uncertainty, and also benefits from the target rich environment inherent to bin-picking.

Figure 9: Side by side comparison of the grasp signature (only 4 finger encoders) of representative (a) successful and (b) failed grasps with P2. The fingers begin the grasp perpendicular to the palm (0°) and reach the final position shown in the figures.

The experimentation is divided in two parts: First, an offline learning process creates a data-driven model of the mapping from grasp signature to grasp outcome. Second, the robot attempts grasps until it detects a singulated object in a recognizable pose. Grasp classification and in-hand localization capabilities are key to the success of our approach. In the next sections we evaluate and compare the performance of P1 and P2 in both capabilities.

6.1 Experimental Setting

We test our prototype grippers with a 6 DOF articulated industrial manipulator. A preprogrammed plan moves the gripper in and out of the bin iteratively while the gripper opens and closes. For each iteration we record the final state of the hand for P1, and the entire grasp signature for P2. We also record the grasp outcome—the number of markers grasped and their pose in the gripper.

The system architecture is built using Robot Operating System (ROS) [56]. The system runs a sequential state machine that commands four subsystems interfaced as ROS nodes:

- **Robot Controller:** Interface developed for absolute positioning of an ABB robotic arm.
- **Grasp Controller:** Interfaces the motor controller that drives the gripper. It also logs the grasp signature by capturing the state of the motor and finger encoders during the entire grasp motion.
- **Vision Interface:** Provides ground truth for the learning system, including the number of markers grasped and their position within the hand.
- **Learning Interface:** After offline training, the learning system classifies grasps as singulated or not singulated, and estimates marker orientation for singulated grasps.

The robot follows a preprogrammed path to get in and out of the bin. While approaching the bin, the gripper slowly oscillates its orientation along the vertical axis with decreasing amplitude. The oscillation allows the fingers to penetrate the bin contents without jamming. The penetration depth was hand-tuned to improve the chances of obtaining a single marker. During departure, the gripper vibrates to reduce the effect of remaining friction and help the object settle in a stable configuration. Contact forces are not easily determined but in bench tests we observed forces ranging from three to seven newtons. During the experiments, we occasionally shook the bin to randomize and improve statistical independence of successive trials.

For each prototype we run 200 repetitions of the experiment. The grasp signatures and outcomes make up the dataset used to evaluate the system in terms of singulation detection in Section 6.2 and pose estimation in Section 6.4. Table 2 shows the distribution of the number of markers grasped both with P1 and P2 and Figure 10 shows some representative singulated grasps.

6.2 Experimental Results: Grasp Classification

This section analyzes experimental performance of grasp classification. We use a supervised learning approach to classify grasps as either successful (singulated) or failed, based on grasp signature.

Number of markers	0	1	2	3	4
P1	57 28.5%	83 41.5%	43 21.5%	17 8.5%	0 0.0%
P2	37 18.5%	84 42.0%	49 24.5%	27 13.5%	3 1.5%

Table 2: Distribution of the number of markers grasped. The dataset captured to evaluate the system comprises 200 grasp attempts for each prototype gripper.

Figure 10: Representative singulated grasps for (a) P1 and (b) P2.

After labeling each run of the experiment as success or failure we train a Support Vector Machine (SVM) with a Gaussian kernel [20, 13] to correctly predict singulation. In the case of P1, the classifier constructs a decision boundary in the sensor space—the three finger encoder values of the final gripper pose—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.

For P2 the grasp signature is of much higher dimension. We use Principal Component Analysis (PCA) [64] to project the grasp signature onto a smaller set of linearly uncorrelated features, reducing its dimension and enabling learning with a relatively small set of training examples.

The performance of the system is evaluated using leave-one-out cross-validation. The hyperparameters C and γ are tuned using 10-fold cross-validation on the training set in each training round. The parameter C controls the misclassification cost while γ controls the bandwidth of the similarity metric between grasp signatures. Both parameters effectively trade off fitting accuracy in the training set vs. generalizability. The analysis yields similar accuracies for P1 and P2: 92.9% and 90.5% respectively.

To compare observing the full state of the hand (motor and finger encoders) with observing only the motor encoder, we train a new SVM for P2 where the feature vector contains only the motor signature. The accuracy detecting singulation decreases in this case from 90.5% to 82%.

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 for P1. The SVM optimized for accuracy achieves a recall of 0.89, but a precision of only 0.875—one out of eight positives is a false positive. By choosing an appropriate working point on the precision-recall curve, we can increase precision, reduce false positives, and obtain a slow but accurate singulation system.

6.3 Experimental Results: Early Failure Detection

While the grasp signature of P1 contains only the final sensor values, the grasp signature of P2 contains the entire time series. This gives us the possibility of *early failure detection*. Sometimes it becomes clear long before the end of the grasp, that the grasp is doomed. If the robot can detect failure early, it can also abort the grasp early and retry.

To test early failure detection, we trained a classifier to predict success or failure at several times during the grasp motion. At each instant we train the classifier using only information available prior to that instant. Fig. 13 shows classifier accuracy as it evolves during the grasp, from random at the beginning, to the already mentioned 90.5% at the end.

6.4 Experimental Results: In-hand Localization

In this experiment we estimate the marker orientation based on the grasp signature. We limit the analysis to those grasps that have singulated a marker, and assume that it lies flat on the palm of the gripper. This assumption holds well for P1 but is violated occasionally for P2,

Figure 11: Perspective view and 2D projection (fingers 2 and 3) of the decision boundary found by the Support Vector Machine in the P1 finger encoder space. Dark dots are successful grasps and clear dots are failed grasps. The interior of the bounded region is classified as success. The finger symmetry of the proposed design in Figure 2 is not reflected in the plots. We introduced an offset between the finger rest positions to avoid the situation where all fingers make simultaneous contact and block the grasp.

Figure 12: Precision/Recall plot for singulation detection using P1.

Figure 13: Early failure detection: evolution of the classifier accuracy for P2.

Figure 14: Error distribution in the regression of the orientation of the marker for singulated grasps with P1.

where the marker is sometimes caught on top of a finger or on top of a “knuckle” at the finger base.

We use Locally Weighted Regression [19] to estimate the orientation as a weighted average of the closest examples in the training set, where the weights depend exponentially on the distance between signatures. Because of the cylindrical shape of the marker, we only attempt to estimate its orientation up to the 180 degree symmetry.

Figure 14 shows a polar chart of the error distribution for P1. The leave-one-out cross-validation errors obtained for P1 and P2 are 13.0 degrees and 24.1 degrees respectively. While no improvement of P2 over P1 can be expected for cylindrical shapes, the fact that it performs so much worse is unexpected. The most likely explanation relates to the self-clutter effect described earlier. Our numerical analysis models the fingers as lines, and neglects interactions among the fingers. In reality the fingers are thin cylinders, and they tend to stack up in unpredictable order. Marker interactions with the knuckles may contribute to the problem. All of these considerations introduce noise which appears to be worse for the four-fingered design of P2.

As in Section 6.2, we can be more cautious and allow the system to reject grasps when 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 15. By using an appropriate working point in Figure 15 we can lower the average error at the cost of expected number of retries to get a recognizable grasp. Figure 16 illustrates the resulting predictions of the system for a working point 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 using estimated orientation to rotate the image to horizontal. Deviations from the horizontal correspond to errors in the regression.

7 Discussion

This paper focused on a “let the fingers fall where they may”—“grasp first, ask questions later” approach to grasping. Near-blind grasping, expectation based on offline training, and haptic sensing are combined to address a 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. Whether a slippery hand is a good idea or not,

Figure 15: Tradeoff between the average error in estimating the orientation of the marker and the recall of the bin-picking system for P1.

Figure 16: Orientation correction: (a) Random subset of successful grasps (b) Images of the grasps have been rotated based on their estimated orientation in order to orient the marker horizontally.

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. Thus the paper can be viewed as an exploration of design for perception.

Perhaps more significantly, the paper can be viewed as an exploration of design for learning. The grasping process is generally very complicated. The map from initial conditions to grasp outcome can be extraordinarily complex, but for the simple hands studied in this paper, the map is so simple that its main features can be learned in just 200 trials.

It might be possible to analyze the mechanics of our prototypes and develop a direct algorithm for interpreting the sensor data. Indeed, similar algorithms have been developed for pushing, squeezing, and grasping planar parts with parallel jaw grippers [45, 30]. However, those algorithms depend on numerous simplifying assumptions, including that the part is already singulated. Generalizing those algorithms to parts that are not singulated, and modeling the dynamic interactions with surrounding clutter, is well beyond the state of the art. The learning approach can handle even the extreme clutter challenge presented by bin-picking. Likewise the learning approach can deal with three dimensions, arbitrary shapes, unknown or uncertain shapes, and many other variations not addressed by research on direct algorithms.

There are many interesting areas for future research:

- **Placing.** Table 1 shows a checkmark in the “placing” column for bin-picking, but no checkmark for our prototype hands. Early experiments show that our prototypes can place highlighters, but much remains to be done.
- **Morphology.** Our prototype hands have generic designs for palm and fingers. It is obvious that the stable pose distribution could be improved by applying some well-known design features, such as cutting v-shaped grooves across the palm.
- **Generalizations.** Generality is one of the main goals of our work, so we plan to explore bin-picking with other shapes, and also to explore entirely different task domains, such as object retrieval by a domestic service robot.
- **Dimensions of grasping.** Finally, we would like to improve on Table 1. 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 Grasp Potential Energy: A Linear Complementarity Problem

Here we solve the linear complementarity problem that arises from the compliant coupling scheme in Figure 2. We assume a known object shape and pose, and calculate the potential energy of the grasp. Recall that θ_m is the motor position as well as the resting position of each finger, and $\theta_1 \dots \theta_n$ are the finger positions. Also let the finger limit angle θ_i^l be the angle at which finger i would make first contact either with the object or the palm when closing.

We adopt the convention that θ_m and θ_i increase in the closing direction. If finger i is not in contact, then $\theta_i = \theta_m$ and that finger’s torque is zero. If it is in contact, then $\theta_i = \theta_i^l \leq \theta_m$.

In other words,

$$\theta_i = \min(\theta_m, \theta_i^l) \quad (2)$$

The motor torque τ_m balances the sum of the finger torques:

$$\tau_m = - \sum_{i=1}^n \tau_i = - \sum_{i=1}^n k_f \cdot (\theta_i - \theta_m) \quad (3)$$

Imposing the motor torque to equal the desired stall torque τ_s , constrained by Equation 2 yields a linear complementarity problem. To simplify the solution we renumber the fingers so that the finger limit angles θ_i^l are ordered from smallest to largest. Now suppose we increase motor angle θ_m until the motor reaches τ_s . Before the first contact, the motor torque is 0. After the first contact and before the second, the torque increases linearly from 0 to $k_f \cdot (\theta_2^l - \theta_1^l)$. By repeating the same process we find a series of torque limits, the highest torque attained before the next finger contact:

$$T_0 = 0, \quad T_i = \sum_{j=1}^i k_f \cdot (\theta_{i+1}^l - \theta_j^l) \quad (4)$$

Let i_m be the largest i such that T_i is smaller or equal to the desired stall torque τ_s . At the limit angle θ_{i_m} only fingers 1 to i_m contact the object, each finger i providing torque $k_f (\theta_{i_m}^l - \theta_i^l)$ for a total of T_{i_m} . In the final grasp configuration, the remaining torque until motor stall $\tau_s - T_{i_m}$ is split evenly between those fingers. The final resting pose of the motor is then:

$$\theta_m = \theta_{i_m}^l + \frac{(\tau_s - T_{i_m})}{k_f} \quad (5)$$

We then substitute the value of the final motor angle in Equation 2 to get the finger angles $\theta_1 \dots \theta_n$ and finally in Equation 1 to obtain the grasp potential energy.

Acknowledgements

We wish to thank Mike Erdmann, Garth Zeglin, Amir Degani, Robbie Paolini, Matt Bernstein, Alex Zirbel, Billy Okal, Koen Ruskamp, Tomas Lozano-Perez, Tom Fuhlbrigge, David Bourne, and Mark Cutkosky for their precious advise and help.

This work was supported by the National Science Foundation [NSF-IIS-0916557], the Defense Advanced Research Projects Agency [DARPA-BAA-10-28] and Army Research Laboratory [W911NF-10-2-0016], Intel, Junta de Comunidades de Castilla-La Mancha and Caja Madrid. This work does not necessarily reflect the position or the policy of the U.S. Government, DARPA or ARL. No official endorsement should be inferred.

References

- [1] P. Allen and K. Roberts. Haptic object recognition using a multi-fingered dextrous hand. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 342–347, 1989.
- [2] A. Ambler, H. Barrow, C. Brown, R. Burstall, and R. Popplestone. A versatile computer-controlled assembly system. In *International Joint Conferences on Artificial Intelligence (IJCAI)*, pages 298–307, 1973.
- [3] B. Baker, S. Fortune, and E. Grosse. Stable prehension with a multi-fingered hand. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 570–575, 1985.
- [4] Barrett Technologies. The Barrett hand. <http://www.barrett.com/robot/products-hand.htm>.
- [5] D. Berenson and S. Srinivasa. Grasp synthesis in cluttered environments for dexterous hands. In *IEEE/RAS International Conference on Humanoid Robots (Humanoids)*, 2008.
- [6] A. Bicchi. Hands for dexterous manipulation and robust grasping: A difficult road toward simplicity. *IEEE Transactions on Robotics and Automation*, 16(6):652–662, 2000.
- [7] A. Bicchi, K. Salisbury, and D. Brock. Contact sensing from force measurements. *International Journal of Robotics Research*, 12(3):249, 1993.
- [8] K. Boehnke. Object localization in range data for robotic bin picking. In *IEEE International Conference on Automation Science and Engineering (CASE)*, pages 572–577, 2007.
- [9] D. L. Brock. Enhancing the dexterity of a robot hand using controlled slip. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 249–251, 1988.
- [10] 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.
- [11] R. Brost and R. Peters. Automatic design of 3-D fixtures and assembly pallets. *International Journal of Robotics Research*, 17(12):1243–1281, 1998.
- [12] E. Brown, N. Rodenberg, J. Amend, A. Mozeika, E. Steltz, M. R. Zakin, H. Lipson, and H. M. Jaeger. Universal robotic gripper based on the jamming of granular material. *Proceedings of the National Academy of Sciences*, 107(44):18809–18814, 2010.
- [13] C. Burges. A tutorial on support vector machines for pattern recognition. *Data Mining and Knowledge Discovery*, 2:121–167, 1998.
- [14] J. Butterfass, M. Grebenstein, H. Liu, and G. Hirzinger. DLR-Hand II: Next generation of a dextrous robot hand. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 109–114, 2001.
- [15] J. Canny and K. Goldberg. A RISC approach to robotics. *IEEE Robotics Automation Magazine*, 1(1):26–28, 1994.

- [16] J. Canny and K. Goldberg. “RISC” industrial robotics: Recent results and open problems. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 1951–1958, 1994.
- [17] 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 International Conference on Rehabilitation Robotics (ICORR)*, pages 510–517, 2009.
- [18] M. Ciocarlie and P. Allen. Data-driven optimization for underactuated robotic hands. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 1292–1299, 2010.
- [19] W. Cleveland and S. Devlin. Locally weighted regression: An approach to regression analysis by local fitting. *Journal of the American Statistical Association*, 83:596 – 610, 1988.
- [20] C. Cortes and V. Vapnik. Support-vector networks. In *Machine Learning*, volume 30, pages 273–297, 1995.
- [21] M. Cutkosky. On grasp choice, grasp models, and the design of hands for manufacturing tasks. *IEEE Transactions on Robotics and Automation*, 5(3):269–279, 1989.
- [22] A. De Santis. *Modelling and Control for Human–Robot Interaction*. PhD thesis, Università degli Studi di Napoli Federico II, 2007.
- [23] M. Dogar and S. Srinivasa. Push-grasping with dexterous hands: Mechanics and a method. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2123–2130, 2010.
- [24] A. Dollar and R. Howe. Towards grasping in unstructured environments: Grasper compliance and configuration optimization. *RSJ Advanced Robotics*, 19(5):523–543, 2005.
- [25] A. Dollar and R. Howe. Simple, robust autonomous grasping in unstructured environments. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 4693–4700, 2007.
- [26] J. Engelberger, D. Lock, and K. Willis. *Robotics in practice: Management and applications of industrial robots*. Amacom, 1980.
- [27] M. Erdmann and M. Mason. An exploration of sensorless manipulation. *IEEE Journal of Robotics and Automation*, 4(4):369–379, 1988.
- [28] A. Farahat, P. Stiller, and J. Trinkle. On the geometry of contact formation cells for systems of polygons. *IEEE Transactions on Robotics and Automation*, 11(4):522–536, 1995.
- [29] P. Gaston and T. Lozano-Pérez. Tactile recognition and localization using object models: The case of polyhedra on a plane. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6(3):257–265, 1984.

- [30] K. Goldberg. Orienting polygonal parts without sensors. *Algorithmica*, 10(2):201–225, 1993.
- [31] W. Grimson and T. Lozano-Pérez. Model-based recognition and localization from sparse range or tactile data. *International Journal of Robotics Research*, 3(3):3–35, 1984.
- [32] D. Grossman and M. Blasgen. Orienting mechanical parts by computer-controlled manipulator. *IEEE Transactions on Systems, Man, and Cybernetics*, 5(5):561–565, 1975.
- [33] H. Hanafusa and H. Asada. Stable prehension by a robot hand with elastic fingers. In *International Symposium of Industrial Robots (ISIR)*, pages 361–368, 1977.
- [34] S. Hirose and Y. Umetani. The development of soft gripper for the versatile robot hand. *Mechanism and Machine Theory*, 13:351–359, 1978.
- [35] S. Hirose and Y. Umetani. The kinematics and control of a soft gripper for the handling of living or fragile objects. In *World Congress on the Theory of Machines and Mechanisms*, pages 1175–1178, 1979.
- [36] J. Hollerbach. Grasping in human and robot hands. *Vision and action: The control of grasping*, pages 243–274, 1990.
- [37] 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.
- [38] S. Jacobsen, J. Wood, D. Knutti, and K. Biggers. The Utah/MIT dextrous hand: Work in progress. *International Journal of Robotics Research*, 3(4):21–50, 1984.
- [39] Y. Jia and M. Erdmann. Pose and motion from contact. *International Journal of Robotics Research*, 18(5):466, 1999.
- [40] D. Kriegman. Let them fall where they may: Capture regions of curved objects and polyhedra. *International Journal of Robotics Research*, 16(4):448–472, 1997.
- [41] T. Laliberte and C. Gosselin. Simulation and design of underactuated mechanical hands. *Mechanism and machine theory*, 33(1-2):39–57, 1998.
- [42] S. Lederman and R. Klatzky. Haptic perception: A tutorial. *Attention, perception & psychophysics*, 71(7):1439–1459, 2009.
- [43] T. Lozano-Pérez, J. Jones, E. Mazer, P. O’Donnell, W. Grimson, P. Tournassoud, and A. Lanusse. Handey: A robot system that recognizes, plans, and manipulates. In *IEEE International Conference on Robotics and Automation (ICRA)*, volume 4, pages 843–849, 1987.
- [44] T. Lozano-Pérez, J. Jones, P. O’Donnell, and E. Mazer. *Handey: A robot task planner*. MIT Press Cambridge, MA, USA, 1992.
- [45] M. Mason and K. Salisbury. *Robot Hands and the Mechanics of Manipulation*. The MIT Press, Cambridge, 1985.

- [46] M. Mason, S. Srinivasa, and A. Vazquez. Generality and simple hands. In *International Symposium of Robotics Research (ISRR)*, 2009.
- [47] M. Mason, S. Srinivasa, A. Vazquez, and A. Rodriguez. Generality and simple hands. Technical Report CMU-RI-TR-10-40, Robotics Institute, November 2010.
- [48] C. Melchiorri and M. Kaneko. Robot hands. In *Springer Handbook of Robotics*, pages 345–360. Springer-Verlag, 2008.
- [49] M. Moll and M. Erdmann. Reconstructing the shape and motion of unknown objects with active tactile sensors. In *Workshop on the Algorithmic Foundations of Robotics (WAFR)*, pages 293–309, 2004.
- [50] G. Monkman, S. Hesse, and R. Steinmann. *Robot grippers*. John Wiley and Sons, 2007.
- [51] J. Napier and R. Tuttle. *Hands*. Princeton University Press, 1993.
- [52] L. Natale and E. Torres-Jara. A sensitive approach to grasping. In *International Workshop on Epigenetic Robotics*, pages 87–94, 2006.
- [53] T. Okada. Object-handling system for manual industry. *IEEE Transactions on Systems, Man and Cybernetics*, 9(2):79–89, 1979.
- [54] 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.
- [55] 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.
- [56] M. Quigley, B. Gerkey, K. Conley, J. Faust, T. Foote, J. Leibs, E. Berger, R. Wheeler, and A. Ng. ROS: an open-source robot operating system. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2009.
- [57] Robotiq. The Adaptive Gripper. <http://robotiq.com/robot-hand>, 2010.
- [58] A. Rodriguez and M. Mason. Grasp invariance. In *Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 2010.
- [59] A. Rodriguez, M. Mason, and S. Srinivasa. Manipulation capabilities with simple hands. In *International Symposium of Experimental Robotics (ISER)*, 2010.
- [60] K. Salisbury and J. Craig. Articulated hands: Force control and kinematic issues. *International Journal of Robotics Research*, 1(1):4–17, 1982.
- [61] A. Saxena, J. Driemeyer, and A. Ng. Robotic grasping of novel objects using vision. *International Journal of Robotics Research*, 27(2):157–173, 2008.
- [62] G. Schlesinger. Der Mechanische Aufbau der Kunstlichen Glieder. Ersatzglieder und Arbeitshilfen, part II, 1919.

- [63] D. Siegel. Finding the pose of an object in a hand. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 406–411, 1991.
- [64] L. Smith. A tutorial on principal components analysis. Cornell Univeristy, 2002.
- [65] A. Sudsang, N. Srinivasa, and J. Ponce. On planning immobilizing grasps for a reconfigurable gripper. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 106–113, 1997.
- [66] R. Tella, J. Birk, and R. Kelley. General purpose hands for bin-picking robots. *IEEE Transactions on Systems, Man and Cybernetics*, 12(6):828–837, 1982.
- [67] D. Theobald, W. Hong, A. Madhani, B. Hoffman, G. Niemeyer, L. Cadapan, J. Slotine, and K. Salisbury. Autonomous rock acquisition. In *AIAA Forum on Advanced Developments in Space Robotics*, 1996.
- [68] R. Tomovic and G. Boni. An adaptive artificial hand. *IRE Transactions on Automatic Control*, 7(3):3–10, 1962.
- [69] N. Ulrich. Grasping with mechanical intelligence. MS-CIS-89-51, University of Pennsylvania, 1989.
- [70] N. Ulrich and V. Kumar. Grasping using fingers with coupled joints. In *ASME Mechanisms Conference*, 1988.
- [71] N. Ulrich, V. Kumar, R. Paul, and R. Bajcsy. Grasping with mechanical intelligence. In *Symposium on Robots, Manipulators and Systems*, 1990.
- [72] N. Ulrich, R. Paul, and R. Bajcsy. A medium-complexity compliant end effector. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 434–436, 1988.
- [73] A. Wallack and J. Canny. Planning for modular and hybrid fixtures. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 520–527, 1994.
- [74] A. Wallack and J. Canny. Modular fixture design for generalized polyhedra. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 830–837, 1996.
- [75] Z. Xu, T. Deyle, and C. Kemp. 1000 Trials: An empirically validated end effector that robustly grasps objects from the floor. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 2160–2167, 2009.
- [76] 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.