Generality and Simple Hands

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

Abstract While complex hands seem to offer generality, simple hands are more practical for most robotic and telerobotic manipulation tasks, and will remain so for the foreseeable future. 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 and object localization. The central idea is to use knowledge of stable grasp poses as a cue for object localization. This leads to some novel design criteria, such as a desire to have only a few stable grasp poses. We explore some of the design implications for a bin-picking task, and then examine some experimental results to see how this approach might be applied in an assistive object retrieval task.

1 Introduction

This is the first paper from a project that aims to develop robot grippers that are simple, yet also general and practical. By "simple", we mean hands with a few actuators, a few simple sensors, and without complicated mechanisms, so that the whole hand would be small, light, and inexpensive.

Can a hand be both simple and general? Perhaps generality requires complexity. Some arguments in favor of complexity are:

Matthew T. Mason
Carnegie Mellon University, Pittsburgh, PA, e-mail: `matt.mason@ri.cmu.edu`

Siddhartha S. Srinivasa
Intel Research, Pittsburgh, PA, e-mail: `siddhartha.srinivasa@intel.com`

Andres Vazquez
Universidad de Castilla-La Mancha, Spain, e-mail: `andress.vazquez@uclm.es`

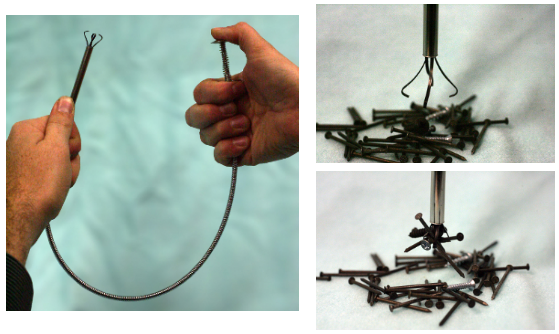

Fig. 1 The common "pickup tool" is very simple, but also very effective in achieving stable grasps over a broad class of shapes. Four fingers of spring steel are compliantly driven by a single actuator.

- Grippers for manufacturing automation are often simple and 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: humans using prosthetic hooks, teleoperated systems with simple pincer grippers, or the simple pickup tool shown in Fig. 1. We conclude that while there is a tradeoff between simplicity and generality, the details of that tradeoff are important and poorly understood. Simple grippers offer a level of generality that is yet untapped in autonomous robotic systems.

To explore the tradeoff between simplicity and generality, we list capabilities required of a general purpose gripper (Section 1.2), and discuss clutter (Section 1.3). However, it simplifies the discussion if we begin with a specific example, the Pickup Tool of Fig. 1, and a simple hand concept inspired by the pickup tool (Section 1.1). The rest of the paper describes simulation and analysis of the simple hand grasping a variety of shapes, and an experimental study motivated by the simple hand concept.

1.1 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 Fig. 1, 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 calculate 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 estimating the object pose (as well as whether a single object was captured).

The main problem addressed by this paper is how to determine whether a single object was captured, and how to estimate the object pose. We propose to use a table of the stable poses with corresponding finger encoder values, produced offline either by experiment or in simulation.

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 design has rigid cylindrical fingers attached to a disk-shaped palm by revolute joints with encoders (Fig. 2). As with the pickup tool, all three fingers are compliantly coupled to a single actuator.

For the bin-picking task the goal is to retrieve a single familiar object from a bin of objects, and to accurately estimate its pose in the hand. For the sphere shown in the figure, estimation of position would suffice. For a non-spherical object we would require orientation as well as position. 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;
- A table of stable grasp configurations and corresponding encoder values to determine whether a single object is present, and to estimate its pose;
- Offline training of pose and singulation classifiers, either in simulation or in the real world;
- Iteration of a reach, grasp, withdraw, and classify strategy until a successful grasp of a single object in a recognized pose is achieved.

Figures 2 and 3 show a dynamic simulation of the concept, applied to a bin picking problem. The system performed blind grasps, identified successful single object grasps, and estimated pose modulo up to symmetries. We also tried some variations: shaking the bin, and sweeping the bin with the fingers, preparatory to the grasping operation. Both behaviors improved performance, and the two behaviors together brought the success rate close to 50%. More advanced versions of this scenario would include the use of visual guidance, of strategically planned pregrasp poses, and a variety of motion and perceptual strategies that could be employed.

The central idea is to use knowledge of stable poses to determine whether a single part has been captured, and to localize that part in the grasp. The general idea is well known in parts orienting research (see [15, 11] for two examples) and has even been used in the context of simple hands ([18, 10]). Yet we believe the idea can be taken further, perhaps even leading to practical designs working across a broad range of manipulation tasks.

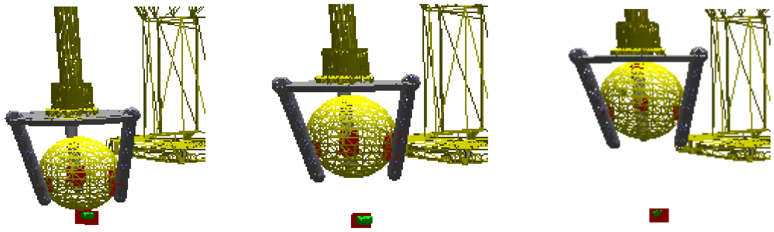

(a) (b) (c)

Fig. 2 Grasping of a Sphere

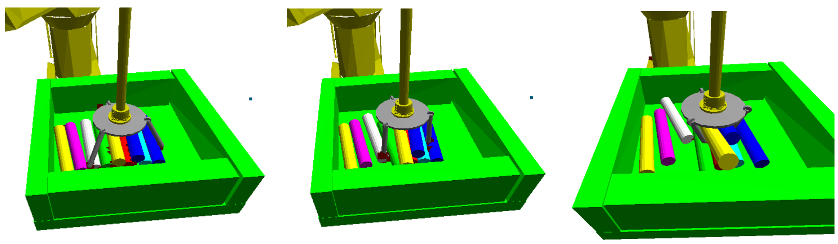

(a) (b) (c)

Fig. 3 A blind grasp capturing two cylinders. The failure is detected and the operation will be repeated. The cylinders were heaped up by a blind sweeping operation prior to the grasp.

For the approach to work in its simplest form, the stable poses must map to isolated points in the space of encoder values. Hence the primary focus of the paper is the set of stable poses. We consider how hand design parameters affect the stable pose space, and we also examine the stable pose space for a typical assistive home robot in an object retrieval task, in Sections 2 and 3. First, we return to our discussion of generality in hand function.

1.2 Dimensions of Grasping

The bin-picking application described above is the framework that motivates several problems, some of which are addressed in the rest of the paper. In this section we place those problems in a broader context. While a precise definition of generality in hand function is elusive, we can at least identify the dimensions. What are the requirements of a grasp? What are the problems that a grasp might solve?

Capture and *Stability.* The main theme of grasping research seems to be stability, but capture seems equally important.
In-hand manipulation and *Grasp adjustment.* Controlled motion of the object in the grasp. Typical prior work has employed fingertip grasps with several actuators

to achieve controllability, but there are other strategies including, for example, tapping two chopsticks against a tabletop to align the ends.

Clutter and *Singulation.* Clutter refers to everything that might limit access to the object. See Section 1.3 for further discussion. Singulation means extracting one part at a time.

Shape diversity and *Shape uncertainty.* Shape diversity: Does the grasp work over a broad range of shapes? Shape uncertainty: does the grasp work when the shape isn't entirely known? There is a difference. A net can capture an object without even knowing what the shape is. A modular fixturing kit can be configured to grasp a wide range of shapes, but not without knowledge of the shape.

Perceptual issues: *Localization, Object recognition, Shape estimation.* Localization means estimation of object pose in the hand. Object recognition assumes there is some class of parts, perhaps even finite, which includes the present object.

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.

Other task specific 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. Some applications may involve complex processes involving multiple objects, such as dealing cards.

We can use the above list to characterize different tasks. The bin picking application, by its nature, poses challenges in clutter and singulation. Our approach to it, employing a simple hand iteratively employing a blind strategy, entails additional challenges of capture and localization. Since we propose to use a single hand design over a range of parts, we introduce the shape diversity issue.

In contrast, assistive robotic object retrieval, the application addressed in Section 3, poses challenges in capture, clutter, and shape diversity. Typical households could not afford to have one robot to retrieve spoons, another to retrieve forks, and so on.

In this paper we focus on capture, clutter, stability, and pose estimation, but ultimately a general-purpose gripper must address the entire list of requirements. Then there are many additional pragmatic issues: cost, weight, ruggedness, and ease of programming. Those are perhaps the main motivation for the use of simple grippers.

1.3 Grasping versus clutter

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 almost ubiquitous. Previous work has seldom addressed clutter explicitly, but clutter often affects the design of the robot, the choice of grasp, the robot path, in fact just about every aspect of a system.

First consider the effect of clutter on hand design. Suppose you are designing a hand from simple geometrical elements: points, lines, or planes. 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 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 [11] 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 the table below. Higher dimensional elements are better for grasping; lower dimensional elements are better for avoiding clutter.

	grasping	clutter
points	bad	good
lines	okay	okay
planes	good	bad

This table suggests 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 laporoscopic forceps. You might even view the pickup tool (Fig. 1) 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 forceps for the extreme clutter encountered in surgery.

Clutter and grasping are in opposition. Grasping is easy, if there is no clutter. Clutter is easily avoided, if no grasp is required. 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. The only cases involving less clutter are catching butterflies or cleaning a pool, where it is practical to sweep the entire nearby volume with a net.

1.4 Previous work

There is a long history of research in robotic hands, and related issues in perception, planning, and control. See [13] for an early but still highly relevant overview, and [3] for a more recent overview. Generality and simple hands are particularly relevant issues for domestic and assistive robotics. See [7, 24, 19] for some examples.

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, grasping a variety of shapes [1]. Freddy II is also one of the rare robotics research projects to address grasping in clutter. The Handey system also used a simple parallel-jaw gripper exhibiting impressive generality [17].

Our gripper concept is similar to Hanafusa and Asada's [12], who analyzed the stability of planar grasps using three compliant fingers. Our concept is further informed by the analysis of Baker, Fortune and Grosse [2] who developed some interesting observations on grasp stability by compliant fingers. Since that time, most previous work on stable grasping has assumed a rigid grasp, or has assumed compliance primarily at the finger tips [20, 16], or has addressed the choice of control strategies for complex grippers.

Our approach contrasts with that often referred to as "dexterous" hands: the use of several fingers, with several actuators per finger, with resulting freedom in the placement of fingers, programmable compliance, and fully controllable in-hand manipulation [23, 14, 4]. This contrast is also reflected in our choice of hard frictionless contacts, despite several advantages of soft high-friction fingertips. As noted in [6] and elsewhere, soft high-friction contact reduces surface forces, increases stability, and increases the effective coupling between finger and object. However, our approach does not necessarily benefit from increased stability and coupling between finger and object.

Our approach has its roots in the parts orienting research community [11, 15, 10, 5, 8], where the ideas of using knowledge of stable poses and iterative randomized blind manipulation strategies are well known. As far as we know the present paper is the first to apply these ideas to problems including capture and singulation in cluttered environments.

One central question is the existence and characterization of stable poses without friction, and without force closure, sometimes referred to as higher-order form closure, or second-order stability. See [9] for examples involving planar polygons, and [21, 22] for closely related research.

2 The set of stable poses

This section addresses the set of stable poses of an object grasped by the simple gripper described earlier. The issue is not just whether a certain pose is stable, but whether the subspace of stable poses has a structure that supports pose estimation. The ideal would be a single stable pose with a large basin of attraction, requiring

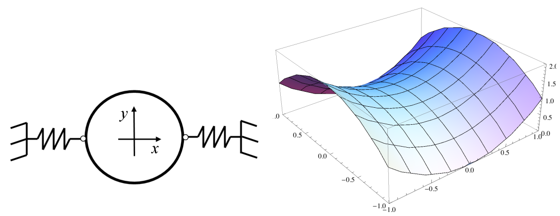

Fig. 4 Potential function for a Hanafusa and Asada-style hand with two fingers grasping a disk. With offsets that would bring the fingers to rest at the center, the potential function is a saddle. Adding two more fingers yields a constant potential function.

zero sensory information to determine the pose. The worst case is symmetric objects where the stable pose space includes submanifolds of poses, all of them mapping to the same finger encoder values. Friction also poses a challenge, yielding connected regions of stable poses, rather than isolated points, but without the stationary sensor mapping.

Nonetheless, existence of stable grasps is an issue, and has some implications for the design of our hand. For a blind grasp to accomodate a range of shapes, fingers require a large range of motion. We consider two alternatives. Option one – *large offsets* – obtains the entire range of motion from individual finger springs. Bounds on contact force require that spring be soft, and that the spring offsets be sufficient to move the finger well past the minimum desired motion. Option two – *small offsets* – would use stiff finger springs with very small deflections. The large of range of motion could be obtained by driving the actuator until a threshold force is exceeded, or by a soft spring in series with the actuator. The following sections address both options for a disk in two dimensions, and for a sphere and a polyhedron in three dimensions.

2.1 Unstable grasp of a disk in 2D

While it may seem obvious that four fingers spaced equally about a disk would give a stable grasp, the reality is not so simple. Large offsets may yield an unstable grasp [2]. First consider a two-finger grasp (Fig. 4). While the finger compliances stabilize motions aligned with the fingers, the curvature of the disk yields a negative stiffness for transverse motions. The potential function is a saddle and the grasp is unstable. If we add two more fingers to obtain a four-finger grasp, the negative stiffness for one finger pair can nullify the positive stiffness of the other pair, yielding a metastable grasp.

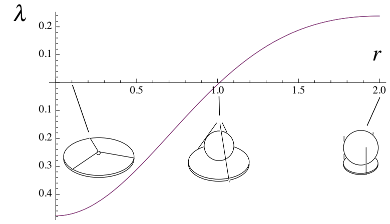

Fig. 5 Local stability of sphere grasped with large finger spring offsets. The stiffness matrix eigenvalues are negative for small spheres and positive for large spheres.

2.2 Stable grasp of a sphere

This section explores the stability of a sphere being grasped by a palm and three line fingers. The main focus is to examine variations of scale. We also reexamine the effect of large offsets versus small.

The hand is the simple gripper of Fig. 2, except the fingers are lines rather than cylinders. To find the total potential energy we start with the potential of a single finger. For a sphere of radius r located at (x,y,z), and a finger nominally aligned with the x-axis, the finger angle is:

$$\theta = \pi/2 - \tan^{-1}(z/x) - \tan^{-1}(\sqrt{r^2-y^2}/\sqrt{x^2+y^2+z^2-r^2})$$

The total potential is $U = 1/2k\Sigma(\theta_i - \theta_{i0})^2$, assuming stiffnesses k and offsets θ_{i0}. We assume the sphere is in palmar contact, i.e. $z = r$, and write the potential as a function of x and y. The gradient of this potential is the force in the x-y plane, and the Hessian of the potential is the stiffness matrix. With the sphere placed at the origin the gradient is zero as expected. We can determine stability by examining the eigenvalues of the Hessian. Fig. 5 shows the eigenvalues when the springs have large offsets—large enough to close the fingers all the way to the palm. The two eigenvalues are equal, and are negative for small spheres, and positive for large spheres. The grasp is unstable for small spheres, and stable for large. Note the contrast with the planar analysis of the previous section. For the largest sphere considered, the fingers are exactly vertical, mimicking the planar three finger grasp. However, the planar grasp is unstable, where the three-dimensional grasp is stable. Evidently when the sphere moves toward the gap between two fingers, the third finger's force inclines downward, so that the planar component drops off rapidly enough to eliminate the instability.

The eigenvalues of the Hessian give us the local stability, but tell us nothing about the basin of attraction. In Figs. 6 and 7 we plot the potential surface for three different sphere sizes and both large and small spring offsets. These plots tell a more complicated story. The global structure more closely resembles a three-lobed saddle, sometimes called a monkey saddle. Because a monkey saddle is fundamentally a third-order surface, the second-order analysis of the Hessian determines local stability. But the size of the basin of attraction, and the robustness of the stability, are determined by the size of the sphere and the spring offsets. As expected small offsets improve the basin of attraction for the large sphere, and can stabilize the grasp even for the small sphere. The medium size sphere corresponds to the crossover point of Fig. 5, and is now shown to be unstable, although it is stabilized with small offsets.

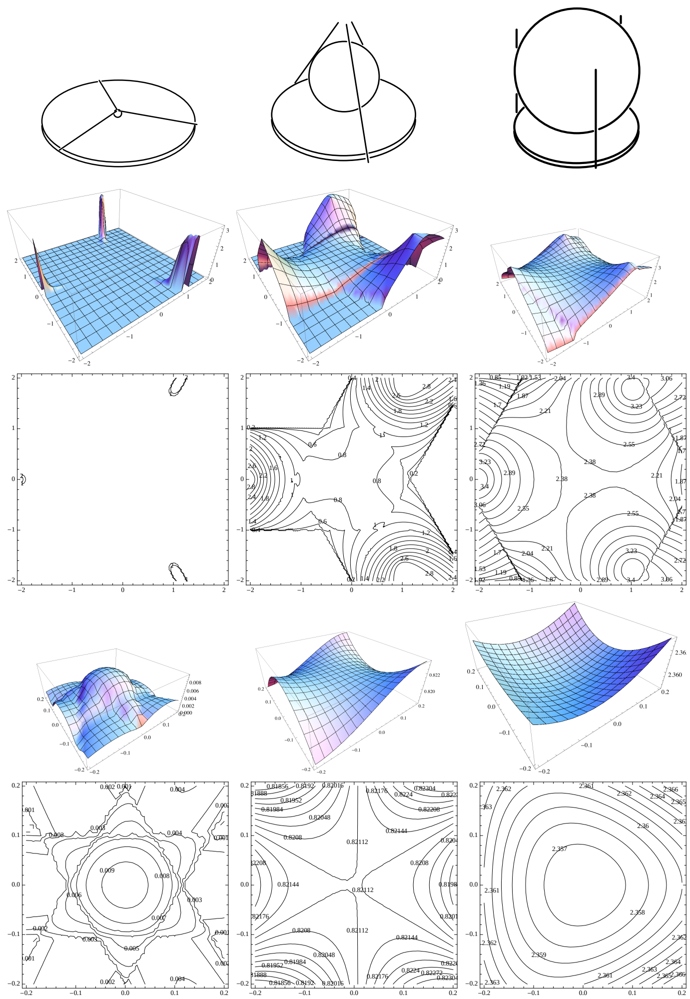

Fig. 6 Potential fields for large offsets. The offsets would close the fingers to the palm. The lower half of the figure is zoomed in to show detail not revealed in the upper half.

The obvious and unsurprising conclusion is that large offsets tend to yield unstable grasps. This is consistent with the general trend of grasp stability analysis over the years, which has focused on stiff grasps and small offsets. However, our goals are different. We do not aim to stabilize as many poses as possible. We would prefer a few stable grasps with large basins of attraction.

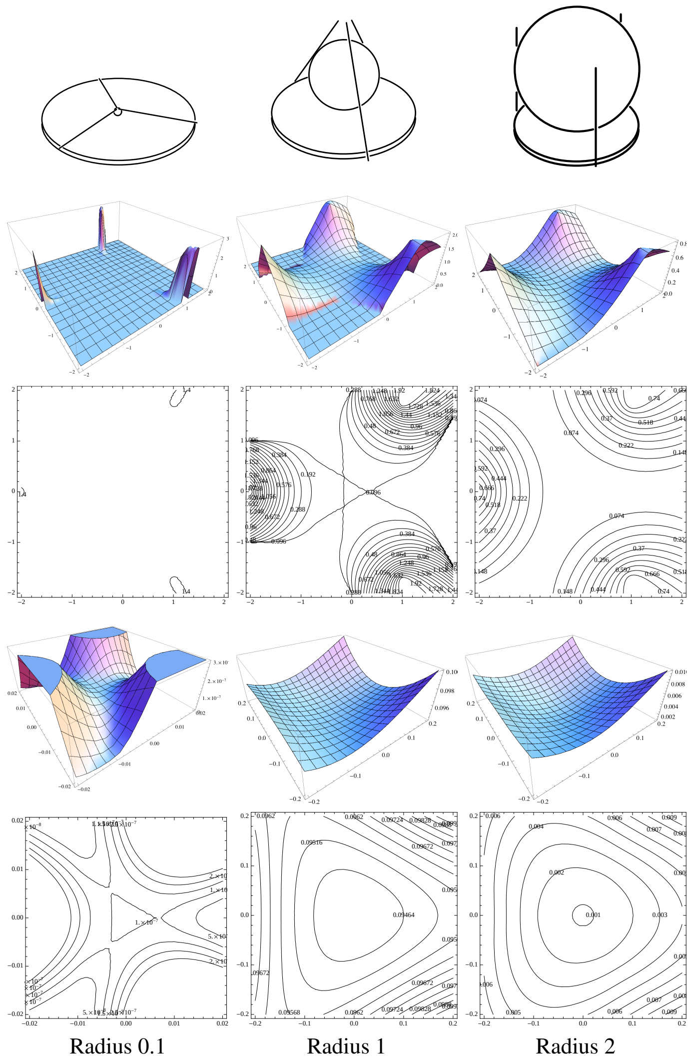

Fig. 7 Potential fields for small offsets. The lower half of the figure is zoomed. The spring deflections are $\pi/10$ for the large and medium spheres, $\pi/10,000$ for the small sphere.

One way to address this is to have very stiff finger springs, and to place a force threshold on the actuator, so that the actuator stalls with small offsets. We will explore a similar idea, which is to have four springs: three stiff finger springs and one soft actuator spring. A very soft actuator spring with a very large offset provides a reasonable approximation of a stalling actuator.

Introducing a fourth spring couples the fingers, and complicates the derivation of a potential function. In short, we now have a linear complementarity problem. The

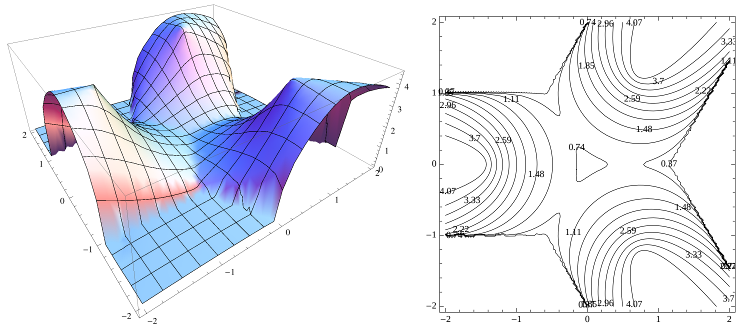

Fig. 8 Potential surface and contours for sphere of radius one grasped by gripper with four springs: each finger springs with stiffness 10 and a motor spring of stiffness 1.

naive algorithm is exponential in the number of contacts, but 2^3 is only 8, so we enthusiastically adopt the naive approach: For every subset of the three fingers, we assume that subset is in contact, and the rest follows easily. Fig. 8 plots the potential surface and several contours, for the medium sphere, where the finger springs are ten times stiffer than the motor spring.

2.3 Stable grasp of a 3-4-5 prism

This section explores the stability of a polyhedron grasped by the three finger hand, with three springs, with large offsets. We exhibit the level surfaces of the potential field in a three-dimensional slice of the six-dimensional configuration space. As expected with an irregular object, the stable poses are isolated points (Fig. 9).

Not all objects will exhibit isolated equilibria. That is obviously the case with symmetries such as the spheres examined in the previous section. Even for polyhedra, it is possible to define a shape and a positive-length path for that shape, while maintaining contact with a palm and three line fingers. The shape is a variation on the example of [9], a planar polygon with three concurrent edges, extruded to a polyhedral prism.

3 Simulation and experiments with the Barrett Hand

This section explores the application of principles detailed in the previous sections to the commercially-available Barrett hand. In contrast to the compliant 3DOF simple hand, the Barrett hand (Fig. 10) is stiff, it is not frictionless and it has 7DOFs, 4 of which are active. We are not using a system carefully engineered to match our assumptions. Rather, we are asking whether our approach is applicable when the assumptions are violated. In particular, we explore the construction of grasp tables, and we examine a part of the space of stable poses. Our results are twofold: (1) departures from our assumptions do present challenges; and (2) even so, the structure of the set of stable poses will support partial estimation of pose.

3.1 Grasp tables

A grasp table is a set of samples of the mapping from grasp inputs to grasp outcomes. Depending on the problem, the inputs may include the relative pose of the object and the hand, the shape of the hand, the material properties of the surfaces, and the expected clutter. The output may include the pose of the object and the hand at the end of the grasping operation, a score of the goodness of the grasp, the location of contact points on the grasp, readings of sensors instrumented on the hand, the object or the environment, or even just a Boolean value denoting grasp success or failure.

Constructing a realistic grasp table depends crucially on the engine that takes the input, simulates contact physics, and produces the output. In this paper, we use real world experiments to produce a small grasp table (Section 3.2) and then address construction of grasp tables using kinematic simulations (Section 3.3).

In the past, grasp tables have served the purpose of precomputed and cached templates or behaviors that the robot searches through online to find a good match for a given scenario. The tables have usually been relatively small, with $20-50$ entries. However, localization introduces an additional mapping: from the space of outputs to the sensor space. Successful localization requires the inversion of the above mapping, *i.e.*, to estimate the relative pose of the object in the hand at the end of the grasping operation, from just the senor readings. Section 3.3 describes the complexities introduced by this additional mapping, and our experimental results.

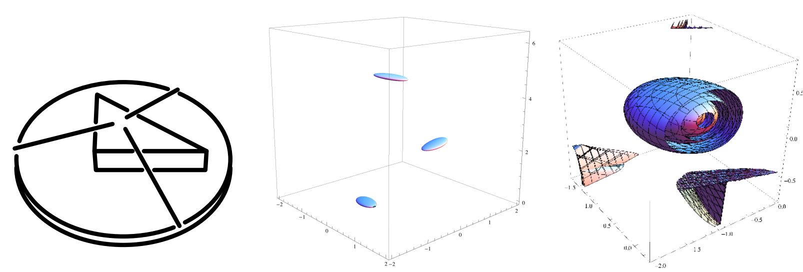

(a) (b) (c)

Fig. 9 Potential function for the simple three finger hand grasping a prism formed by extruding a triangle with side lengths of 3, 4, and 5. Height of the prism is one. The plot 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 first plot shows all three potential wells represented by a single contour. The second plot is a closer view of one potential well, with some additional contours plotted.

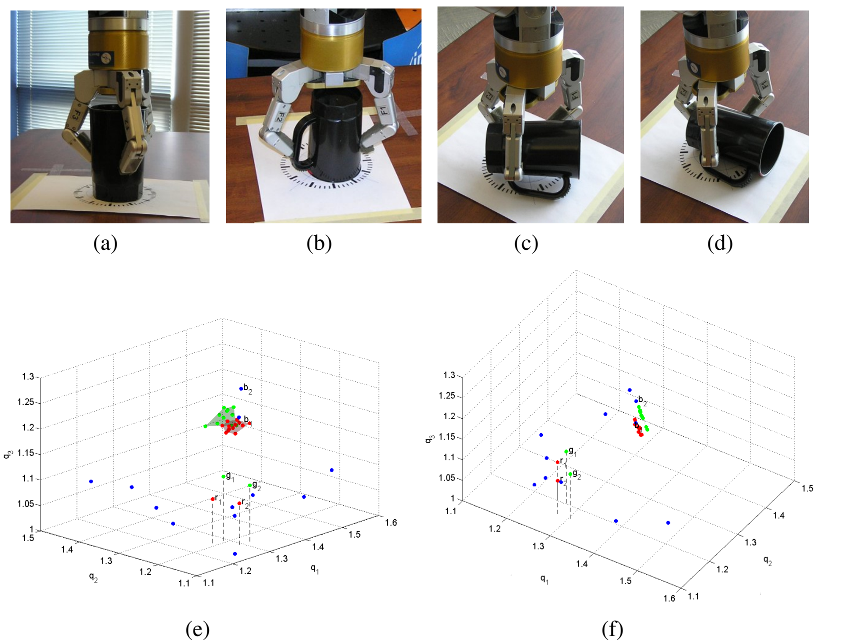

Fig. 10 An experimentally derived grasp table

3.2 Experimentally developed grasp table

Fig. 10 shows an experimentally produced grasp table. The initial mug position was fixed, while several stable orientations were sampled. As the hand closed, the object was constrained only by the table, simulating the effect of the palm. Three clusters appear in the sensor space (Fig. 10(e), Fig. 10(f)): red points for the upright mug, green points for the inverted mug, and blue points for the mug on its side. The upright and inverted poses are separate because the mug has a slight taper. Besides the clusters, we observe some outliers, for example point G1 where one finger landed on the handle instead of the body. (Fig. 10(b)), and point B1 (Fig. 10(d)) where capture failed.

3.3 Grasp tables via kinematic simulation

To further our goal of producing a dense grasp table comprising of many thousands of samples, we augment real-world samples with those derived from simulation.

In our experiments, we observed significant movement of the mug before it settled into a grasp. Simulating such contact dynamics is computationally expensive and inaccurate. Instead, we chose a two-step approach of building the grasp table kinematically (Fig. 11(a)) followed by a refinement step motivated by Section 2

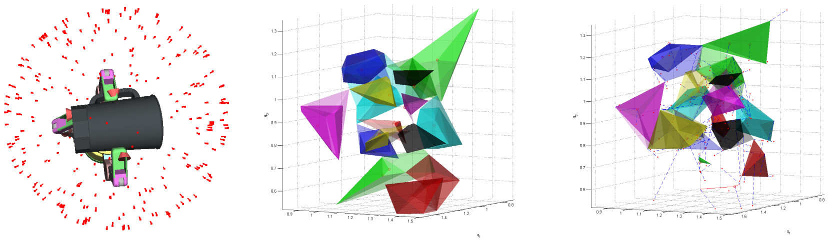

(a) (b) (c)

Fig. 11 Grasp table refinement: (a) Input samples, (b) kinematic grasp table clustered in sensor space, (c) clusters shrink and collapse after minimum energy refinement

Fig. 12 Right: Pose ambiguity in the kinematic grasp table, Left: Predicted pose from grasp refinement and real-world experiment

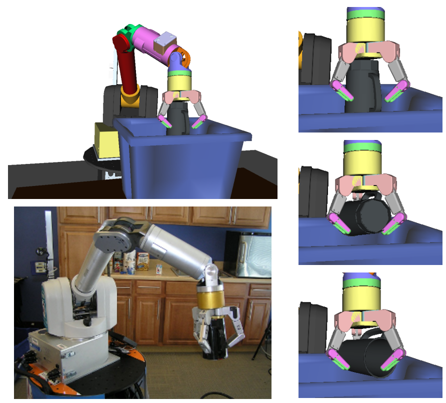

in which the kinematic samples (Fig. 11(b)) are perturbed into the nearest statically stable minimum energy configuration using simulated annealing(Fig. 11(c)). Our preliminary results are encouraging: refinement eliminates unlikely grasps like Fig. 12(right-mid) and Fig. 12(right-bottom) while preserving likely grasps like Fig. 12(left-top) which are experimentally verifiable (Fig. 12(left-bottom)).

Our ultimate goal is to infer object pose from sensor readings. For the Barrett hand with moderate friction, we observed significant pose ambiguity with the kinematic grasp table (Fig. 12(right-all)). With refinement, some of the unstable grasps were filtered out, resulting in fairly accurate pose estimation (Fig. 12(left)). However, the ambiguity could not be eliminated in many other configurations. This suggests two potential strategies: (1) to avoid ambiguous configurations altogether, and (2) to disambiguate by further sensing or sensorless action. Our future work will explore both strategies.

Acknowledgements This work was supported by NSF, DARPA, Junta de Comunidades de Castilla-La Mancha, and Intel. We wish to thank Tom Fuhlbrigge, David Bourne, Mike Erdmann, Alberto Rodriguez, Garth Zeglin, Amir Degani, Koen Ruskamp, Tomas Lozano-Perez. This work is sponsored in part by the Defense Advanced Research Projects Agency. This work does not necessarily reflect the position or the policy of the Government. No official endorsement should be inferred.

References

1. 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.
2. B. Baker, S. Fortune, and E. Grosse. Stable prehension with a multi-fingered hand. In *IEEE Int. Conf. on Robotics and Automation.*, volume 2, pages 570–575, 1985.
3. 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.
4. 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.
5. 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.
6. M. Cutkosky and P. Wright. Friction, stability and the design of robotic fingers. *The Int. Journal of Robotics Research*, 5(4):20–37, 1986.
7. A. Dollar and R. Howe. Towards grasping in unstructured environments: Grasper compliance and configuration optimization. *RSJ Advanced Robotics*, 19(5):523–543, 2005.
8. M. Erdmann and M. Mason. An exploration of sensorless manipulation. *IEEE Journal of Robotics and Automation*, 4(4):369–379, 1988.
9. 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.
10. K. Goldberg. Orienting polygonal parts without sensors. *Algorithmica*, 10(2):201–225, 1993.
11. D. Grossman and M. Blasgen. Orienting mechanical parts by computer-controlled manipulator. *IEEE Tran. on Systems, Man, and Cybernetics*, 5(5):561–565, 1975.
12. H. Hanafusa, H. Asada, et al. Stable prehension by a robot hand with elastic fingers. In *Proc. Seventh Int. Symp. Industrial Robots*, pages 384–389, 1977.
13. J. Hollerbach. Grasping in human and robot hands. *Vision and Action: The Control of Grasping*, pages 243–274, 1990.
14. 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.
15. 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.
16. Q. Lin, J. Burdick, and E. Rimon. Computation and analysis of natural compliance in fixturing and grasping arrangements. *IEEE Tran. on Robotics*, 20(4):651–667, 2004.
17. T. Lozano-Pérez, J. Jones, P. O'Donnell, and E. Mazer. *Handey: A robot task planner*. MIT Press Cambridge, MA, USA, 1992.
18. M. Mason and J. Salisbury Jr. *Robot hands and the mechanics of manipulation*. The MIT Press, Cambridge, 1985.
19. H. Nguyen and C. C. Kemp. Bio-inspired assistive robotics: Service dogs as a model for human-robot interaction and mobile manipulation. In *The Second IEEE/RAS-EMBS Int. Conf. on Biomedical Robotics and Biomechatronics*, pages 542–549, 2008.
20. V. Nguyen. Constructing stable grasps. *The Int. Journal of Robotics Research*, 8(1):26–37, 1989.
21. E. Rimon and J. Burdick. Mobility of bodies in contact. I. A 2nd-order mobility index formultiple-finger grasps. *IEEE Tran. on Robotics and Automation*, 14(5):696–708, 1998.
22. 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.
23. J. Salisbury and J. Craig. Articulated hands: Force control and kinematic issues. *The Int. Journal of Robotics Research*, 1(1):4–17, 1982.
24. 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.