

Manipulation Capabilities with Simple Hands

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

Abstract A simple hand is a robotic gripper that trades off generality in function for practicality in design and control. The long-term goal of our work is to explore that tradeoff and demonstrate broad manipulation capabilities with simple hands. This paper describes two prototype simple hands. Both hands have thin cylindrical fingers arranged symmetrically around a low friction circular palm. The fingers are compliantly coupled to a single actuator. Our experiments with both hands in a bin-picking scenario demonstrate that we can achieve robust grasp classification and in-hand localization using simple statistical techniques. We further show how the classification accuracy increases as the grasp proceeds by exploiting information obtained online. We finally evaluate the relative importance of observing the full state of the hand rather than just observing the state of the actuators.

1 Introduction

Simple hands trade off generality in function for practicality in design and control. Because simple hands have fewer actuators and sensors than complex hands, and because their control strategies are simpler than those of complex hands, simple hands inevitably are less capable than hands without such constraints. Nonetheless some manipulation capabilities remain. Our goal in this paper is to explore this tradeoff through the analysis of a simple hand [14] in a bin-picking application.

Our approach to grasping might be called, “Let the fingers fall where they may”. We close the hand, and expect the details of the grasping process to be determined

Alberto Rodriguez

Carnegie Mellon University, Pittsburgh, USA, e-mail: `albertor@cmu.edu`

Matthew T. Mason

Carnegie Mellon University, Pittsburgh, USA, e-mail: `matt.mason@ri.cmu.edu`

Siddhartha S. Srinivasa

Intel Research, Pittsburgh, USA, e-mail: `siddhartha.srinivasa@intel.com`

Report Documentation Page			Form Approved OMB No. 0704-0188		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2010		2. REPORT TYPE		3. DATES COVERED 00-00-2010 to 00-00-2010	
4. TITLE AND SUBTITLE Manipulation Capabilities with Simple Hands				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Carnegie Mellon University,Pittsburgh,PA,15213				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES 12th International Symposium on Experimental Robotics (ISER 2010), December, 2010. U.S. Government or Federal Rights License					
14. ABSTRACT A simple hand is a robotic gripper that trades off generality in function for practicality in design and control. The long-term goal of our work is to explore that tradeoff and demonstrate broad manipulation capabilities with simple hands. This paper describes two prototype simple hands. Both hands have thin cylindrical fingers arranged symmetrically around a low friction circular palm. The fingers are compliantly coupled to a single actuator. Our experiments with both hands in a binpicking scenario demonstrate that we can achieve robust grasp classification and in-hand localization using simple statistical techniques. We further show how the classification accuracy increases as the grasp proceeds by exploiting information obtained online. We finally evaluate the relative importance of observing the full state of the hand rather than just observing the state of the actuators.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 15	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

by the mechanics of the emergent interaction between hand and object. Once those details have been worked out, we use sensing and statistical techniques to determine the outcome of the grasp.

Analogously, the more common or traditional approach to grasping could be called, “Put the fingers in the right place”—use knowledge of object shape and pose, and models of the mechanics of stable grasp, to plan contact points on the object, and then drive the fingers to those contact points. Assuming accurate sensors, models and controls, all the details of the entire grasping process are determined a priori by the planner. “Put the fingers in the right place” is intensive in its dependence on accurate sensors, models, and controls. Small errors can lead to failure. The approach is mostly suited to complex hands, so that the fingers have the necessary freedoms, actuators and sensing to execute the planned motions.

On the other hand, “Let the fingers fall where they may” is well suited to simple hands, and less dependent on accurate sensing, models, or controls. Taken to the extreme, the approach would seldom work—the fingers would almost always fall someplace useless. The robot has to initiate the grasp at a promising spot, which implies at least *some* expectation of the likely evolution of the grasping process. However, that expectation can be much less detailed and error-prone than in the traditional approach. (This paper sidesteps the issue by using an application where promising spots are plentiful.)

The authors have argued the case for simplicity in the design of robotic hands [14] and proposed a design for a simple hand aimed to reduce the set of possible grasp outcomes: thin cylindrical fingers arranged symmetrically around a low-frictional circular palm, all compliantly coupled to a single actuator. This paper compares the performance of two prototype implementations (P1 and P2) of the proposed simple hand in a bin-picking application.

The central goal of this paper concerns the last stage of the “let the fingers fall where they may” approach: determining the grasp outcome. An offline learning system runs several grasping trials, visually observing the grasp outcome and recording the corresponding hand pose. From this data it infers a map allowing it to interpret online kinesthetic data, addressing two objectives:

- *Grasp classification*: Distinguish between successful and unsuccessful grasping attempts.
- *In-hand localization*: Identify the pose of a grasped object.

The main results of the paper are the performance of prototypes P1 and P2 measured by those two objectives in a bin-picking task (Fig. 1). The hands grasp blindly inside a bin full of identical objects (whiteboard markers) with the goal of singulating an object and localizing it.

Section 3 describes our approach to the bin-picking problem and details the designs of P1 and P2. We then describe the bin-picking experimental setting in Sect. 4.

Section 4.4 addresses an interesting refinement of the approach—determining the grasp outcome before the grasping process is complete, by using the entire time series or *kinesthetic signature* of the grasping process. As the grasp proceeds and additional kinesthetic data accumulates, the confidence also increases. In some cases

it is possible to confidently predict the outcome of the grasp before the end of the grasping process.

Section 5 discusses the results obtained and lessons learned in the process of designing and experimenting with P1 and P2. We conclude in Sect. 6 with a list of ideas we want to explore in subsequent work.




Fig. 1 Bin-picking scenario

2 Related Work

Discussions of the tradeoff between generality and simplicity have been present since the first robotic hands were being developed. One early example can be found in the context of the Utah/MIT Dexterous Hand [11] where Jacobsen et al. raised the question of the relationship between cost and increased functionality by adding complexity to an end-effector. The tradeoff between generality and simplicity, however, has not become a driving factor for the design of robotic hands until recent years. The increasing interest in service and domestic robotic applications [12], in which weight, size, and cost are important factors, make of simplicity a key design goal.

After a few decades of focusing on the design of fully articulated anthropomorphic hands, there has been some recent interest towards a more minimalist approach

where fewer actuators are used to drive more degrees of freedom, and compliance takes care of shape adaptation. In exchange, part of the functionality of the hand is hard coded into its mechanical structure.

In their work on joint coupling design [8], Dollar and Howe did a comprehensive survey on underactuated hands. They classified hands based on the number of freedoms, number of actuators, coupling scheme and source of compliance. The Barrett Hand [1], an offspring of Ullrich’s seminal work on grasping with mechanical intelligence [19], is probably the most commercially successful example of an underactuated gripper. Other prolonged attempts to input simplicity into the design are the SARAH hand [13], recently turned into the commercial Adaptive Gripper [16], or the prosthetic SPRING hand [4].

The robotic hands found in the literature that are closest to our work on simple hands are: Dollar and Howe’s SDM hand [9], with four two-jointed fingers all compliantly coupled to a single actuator; Ciocarlie and Allen’s [5] two-fingered gripper with three joints per finger all compliantly coupled to a single actuator; Xu, Deyle and Kemp’s [20] end-effector designed to robustly capture a large and carefully chosen set of household objects; and Theobald et al.’s simple gripper Talon [18] with two facing sets of fingers driven by a single actuator, for grasping rocks of varying shape and size.

In our approach to grasping, we estimate the outcome of the grasp based on kinesthetic sensor data. Bicchi, Salisbury and Brock [2] explored a similar problem: assuming known finger shape and location, they estimate the contact point from a measured applied wrench, a technique known as *intrinsic contact sensing*. This contact information can be used to infer the pose of a known shape. Our work can be viewed as a generalization of intrinsic contact sensing, where we map directly from kinesthetic sensor data to object pose, bypassing the estimation of contact points. While previous work on intrinsic contact sensing has generally been model-based, our approach is based on machine learning. Assuming availability of the object for the offline learning process, this statistical data-driven approach neatly incorporates numerous sources of information that would be very challenging to capture otherwise, including the effect of the grasping motion and that of surrounding clutter.

In-hand sensor information has previously been shown to improve grasping performance of simple hands [10]. However, some degree of manipulability is always lost when opting for a simple rather than a complex hand. In this paper we show how with enough sensor information, very simple hands are still capable of accomplishing complex tasks.

3 Simple Hand

The main objective of this work is to explore manipulation capabilities with simple hands. In particular, we have chosen a bin-picking task, where the goal is to singulate an object from a bin and to localize its pose in the hand. In this section we describe the approach used to address the bin-picking problem and the design of P1 and P2.

3.1 Approach

The following are the key elements of our approach to the bin picking problem:

- Simple hand designed with a low-friction palm and fingers to reduce the number of stable poses of the object in the hand.
- Grasp first and ask questions later. We close the hand until some stall torque is exceeded and then analyze the outcome of the grasp. (Sect. 4.4 addresses an early termination refinement.)
- Offline learning of the map from kinesthetic sensor data to stable grasp poses with a data-driven approach.
- Repeat strategy until the learning algorithm detects the successful grasp of a single object in a predictable pose.

Central to our approach is the notion of grasp stability. The gripper design needs to take into account the fact that, after the grasping process, we have to answer questions regarding the outcome of the grasp. By reducing the number of stable poses, the hand design simplifies the mapping from hand poses to possible grasp outcomes, and facilitates learning.

To explore the mapping from hand poses to object poses, we analyze the stable grasp of a simple object, a sphere. We model the interaction of the hand with the object as N linear springs in parallel, all connected to the actuator, as in Fig. 2. After driving the actuator of the hand to a stall torque and given the geometry of the hand and object, we can identify the statically stable position of the fingers and the compression of each spring.




Fig. 2 Compliance model: parallel compliance scheme that models the interaction of the hand with the object. In our simulations we normalize the constant of the finger springs to $k_f = 1 \text{ lb.rad}^{-1}$ and the motor is driven to a stall torque of $\tau = 10 \text{ lb.in.}$

The compliance model yields the total potential energy as the sum of the potential energies stored in the N springs. In the presence of any dissipative force, stable poses occur at local minima of the potential energy. Figure 3 shows plots of the potential energy of grasps of the simple hand, both with 3 and 4 fingers, grasping a sphere translating in their palm.

The plots present a unique stable grasp of the sphere, i.e. a unique local minimum of the potential energy function, both for the three-fingered and four-fingered cases. The plots reveal a few interesting points. First, the potential wells in the immediate vicinity of the equilibrium are comparable. Adding fingers is not sufficient to

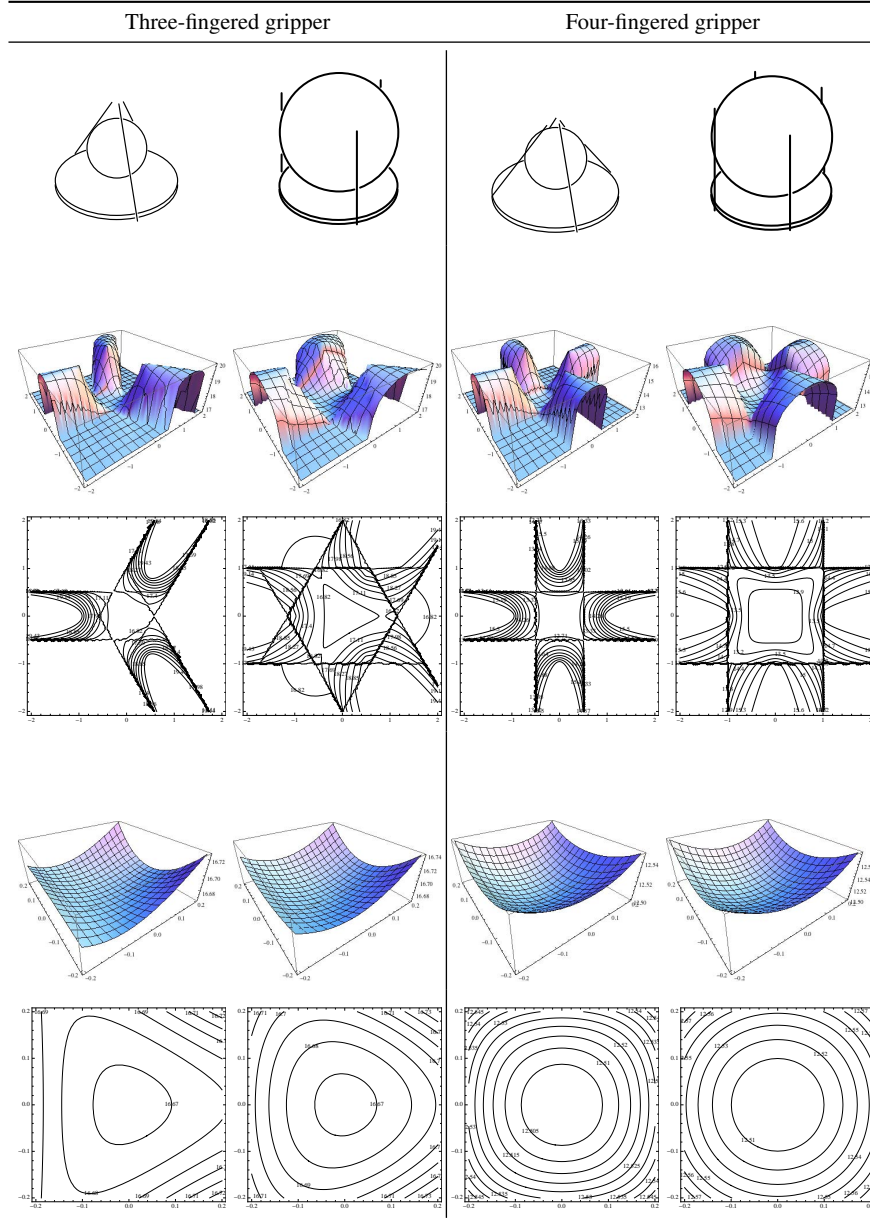


Fig. 3 Potential field surface and contour plots of the grasp of a sphere with **left)** three fingered version of the simple gripper and **right)** four fingered version of the simple gripper. The radius of the palm measures 1 inch while the radius of the small and large spheres measure 0.5 inches and 1 inch respectively. The spring constants are normalized to $k_f = 1 \text{ lb.rad}^{-1}$ and the hands are driven to a stall torque of $\tau = 10 \text{ lb.in}$. The lower half of the figure is the zoomed version of the upper half.

steepen the potential well and thereby increase stiffness of the grasp. However, the global structure is altered, yielding a larger basin of attraction which is less easily escaped.

Nonetheless, the primary design goal is to support grasp recognition and pose estimation. With some object shapes, the addition of a fourth finger should “sharpen” the bottom of the potential well, adding precision to both grasp recognition and pose estimation. This is not the case with a sphere as the plot well illustrates.

3.2 Prototypes P1 and P2

Prototype P1, Fig. 4, has three fingers. The actuation is transmitted from the motor to the fingers through a series of gears. Torsional springs coupling the fingers with the gear assembly introduce compliance which allows for moderate conformability of the hand.

Prototype P2, also in Fig. 4, 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, at the same time, equalize the transmission ratio from the vertical motion of the leadscrew to the rotational motion of the finger. One of the links in each finger linkage is elastic (black link in the close up of the transmission mechanism of P2 in Fig. 4) and provides moderate conformability to the hand.

In [14] 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, clutter, object shape variation, multiple/deformable objects, recognition/localization and placing. We make use here of that set of general-purpose *dimensions* to compare the designs of P1 and P2.

Table 1 characterizes the bin-picking task as well as P1 and P2 in terms of those characteristics. Due 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 Sect. 4 that for whiteboard 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 for our approach. We need to know the state of the hand for mapping it to known stable poses of the object. For that, P2 is equipped with absolute encoders on each finger and in the actuator. The combination of finger encoders and motor encoder, gives us an estimate of the compression of the compliant source for each finger. In Sect. 4.2 we evaluate the improvement in grasp classification that full observability of the hand pose yields with respect to just the state of the actuator. Fig. 5 shows the grasp signature of a typical grasp motion and an estimate of finger deviation from resting position due to hand-object interaction.




Fig. 4 Prototypes P1 and P2. **Top** 3D model and close up of the transmission mechanism. **Mid** Side view. **Bottom** Front view.

Table 1 Dimensions of general-purpose grasping. A checkmark indicates either a task requirement or a hand capability. $\uparrow\uparrow$ indicates an improvement of P2 with respect to P1, and $\downarrow\downarrow$ otherwise.

	Stability	Capture	In-hand manipulation	Clutter	Object shape variation	Multiple and deformable objects	Recognition, localization	Placing
Bin picking	✓	✓		✓			✓	✓
P1	✓	✓		✓			✓	
P2	✓ $\uparrow\uparrow$	✓ $\uparrow\uparrow$		✓ $\downarrow\downarrow$	✓		✓	

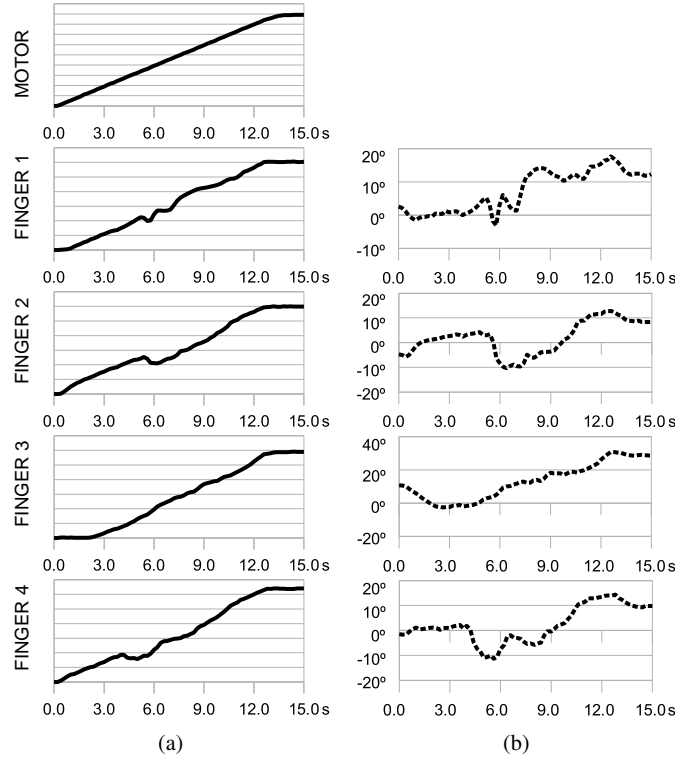


Fig. 5 Grasp signature of a representative singulation attempt. In the absence of any external disturbance, motor and finger encoders should be proportional. Deviations from that proportionality are correlated with external forces applied to the fingers. **a)** P2’s motor and finger encoder signals during a complete grasp motion. Y units are encoder “ticks”, normalized for visualization purposes. **b)** Estimate of finger deviation from the resting position due to hand-object interaction.

4 Experiments

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. 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.

The experimentation is divided in two parts: First an offline learning process creates a data-driven model of the relationship between signature and outcome of the grasp. Second, once the model is estimated, the robot grasps blindly inside the bin 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.

4.1 Experimental Setting

In the experimental setup, the gripper is attached to a 6 DOF industrial manipulator. A preprogrammed plan moves the gripper in and out of the bin repetitively while the gripper opens and closes. At each iteration we record the state of the gripper over the entire grasp motion, and also note the outcome of the grasp—the number of markers grasped and their pose within the gripper.

The system architecture is built within the framework Robot Operating System (ROS) [15]. The system runs a sequential state machine that commands four sub-systems interfaced as ROS nodes:


- *Robot controller*: Provides an interface for absolute positioning of the robotic arm holding the gripper.
- *Grasp controller*: Interfaces the motor controller that drives the gripper. It also logs the signature of the grasp by capturing the state of the motor and finger encoders along the entire grasp motion.
- *Vision system*: Provides ground truth for the learning system both on the number of markers grasped and their position within the hand.
- *Learning system*: After offline training, the learning system classifies grasps as singulated or not singulated as well as gives an estimation of the orientation of the marker within the hand 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 as a strategy for dealing with clutter. During departure, the gripper vibrates to reduce the effect of friction and to help the object settle in a more stable position.


For each of the prototypes, we run 200 repetitions of the experiment. The grasp signature and outcome of those experiments make up the dataset used to evaluate the system in terms of singulation detection in Sect. 4.2 and pose estimation in Sect. 4.3. Table 2 shows the distribution of the number of markers grasped both with P1 and P2 and Fig. 6 shows the most representative types of singulated grasps obtained.

Table 2 Distribution of the number of markers grasped in the 200 runs of the bin-picking experiment.

	0 markers	1 marker	2 markers	3 markers	4 markers
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 %)



(a) P1



(b) P2

Fig. 6 Representative types of singulated grasps for a) P1 and b) P2.

4.2 Experimental Results: Grasp Classification

In this section we detail the analysis on the classification between successful and failed grasps. We use a supervised learning approach to learn the distinction based on the signature of the grasp. The signature of P1 is the final pose of the three fingers. P2 has a much more complete signature—the value of the four fingers and motor encoders during the entire grasp motion.

After labeling each run of the experiment as success or failure we train a Support Vector Machine (SVM) with a Gaussian kernel [7, 3] to correctly predict singulation. In the case of P2 the dimension of the signature is too large for the amount of training data captured and we use Principal Component Analysis (PCA) [17] to compress the signature, reduce its dimensionality and speed up the learning process.

The performance of the system is evaluated using leave-one-out cross-validation. The hyperparameters C and γ are tuned using 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 (92.9%) and P2 (90.5%).

To evaluate the relative importance of observing the full state of the hand (motor + finger encoders) with respect to observing only the state of the actuators (motor encoder), we train a new SVM for P2 where the feature vector contains only the signature of the motor and no information about the position of the fingers. The accuracy detecting singulation decreases in this case from 90.5% to 82%.

4.3 Experimental Results: In-hand Localization

In this section we regress the orientation of the grasped marker with the signature of the grasp. We focus only on those grasps that have correctly isolated a marker, i.e. the second column in Table 2, and assume that the marker lies flat on the palm of the gripper. Judging by the outcomes of the singulated grasps, the assumption holds well for P1 and is violated occasionally for P2, where the marker is sometimes caught on top of a finger or on one of the “knuckles” at the finger base.

Due to the almost cylindrical shape of the marker, we only attempt to estimate its orientation up to the 180 degree symmetry. We use Locally Weighted Regression [6] to regress the orientation. The orientation of the marker is estimated as a weighted average of the closest examples in the training set, where the weights depend on the distance between signatures.

The leave-one-out cross-validation error 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 it performs so much worse is unexpected. Section 5 discusses some possible reasons.

4.4 Experimental Results: Early Failure Detection

P2 captures the state of the hand during the entire grasp motion. This gives us the possibility of detecting *early failure*. There are situations where it becomes clear long before the end of the grasp that the grasp is not proceeding as it should to correctly singulate an object. If we can detect that, it can potentially be exploited for early abort and retrial by confidently discarding unpromising grasps at different instants in the grasp process.

We put in practice the early failure detection idea by training a classifier to predict success or failure at several points during the grasp motion. At each instant we train the classifier using only information available prior to that instant. Fig. 7 shows the

accuracy of the singulation prediction as it evolves during the grasp, from random at the beginning, to the already mentioned 90.5% at the end.




Fig. 7 Early failure detection: evolution of the accuracy of prediction of the outcome of the grasp.

5 Discussion

The main objective of this paper is to show manipulation capabilities with simple hands. We have chosen a very particular scenario: bin-picking of whiteboard markers. The approach used is to blind-grasp inside the bin until recognized successful singulation of a marker in a recognizable pose. Success relies on two enabling capabilities: detection of object singulation and regression of the pose of the object within the hand.

We have performed experiments with two prototype simple hands P1 and P2. Both are based on the same concept, thin cylindrical fingers symmetrically arranged around a circular frictionless palm. The main difference is that P1 has three fingers while P2 has four. Experimental results show similar accuracies for both prototypes in singulation detection, both greater than 90%. On the other hand, pose regression gives an estimation error of 13.0 degrees for P1 and 24.1 degrees for P2. The big difference between the performances of P1 and P2 comes as a surprise to us, although might reflect the fact that sometimes simpler is better.

The idealized model used in the analysis of grasp stability has two simplifications: that the fingers are infinitesimally thin and that they do not interfere with each other. After careful examination of grasp outcomes, we observed that fingers interfere with each other much more often for P2 than for P1. Figure 6 shows examples of how, even for the most common grasps of P2, fingers are resting on top of other fingers, instead of on top of the object or the palm, as our idealized model assumes. The different possible intertwined configurations for the finger contacts introduces noise into the learning process. Another source of noise that seems to have a greater

effect on P2 than P1 is the marker caught on top of a finger or “knuckle” rather than lying flat on the palm of the gripper.

We have seen that the observation of the full pose of an underactuated hand improves the results in grasp classification. Adding the finger encoder signals when training the classifier increases the experimental accuracy for P2 from 82% to 90.5%.

Finally we have also measured the system accuracy for detecting *early failure*, i.e. situations where it becomes clear before the end of the grasp that the gripper is not going to correctly singulate an object. Early failure detection enables early abort and retrieval, reducing the time to a successful grasp.

6 Future Work

P1 and P2 are prototypes, from which we hope to learn how to build a better P3. While our evaluation is ongoing, we have already learned some valuable lessons.

We have yet to observe any improvement from adding a fourth finger. Still, analysis predicts that four fingers should give an improvement with respect to three in extracting information from kinesthetic sensor data, for at least some shapes. We would like to refine the design of P2 for that improvement not to be masked by the *self-clutter* effect between the fingers.

We finish with a list of design issues that we might address with future prototypes:

1. Non-interfering fingers. Whether we have three or four fingers, it seems clear that the approach would benefit from fingers that do not interfere with each other. The most straightforward way of doing it is by shortening their length, with the consequent shrinking of the capture region. Other options include fingers that retract or bend while they close.
2. Explore palm and finger form design to get more pronounced V-shaped potential fields and increase grasp stability.
3. Placing. Bin picking is not complete without a placing strategy. The designs of P1 and P2 do not address it.
4. Variable stiffness. Stiff fingers yield great stability while soft fingers can be used as sensors. Variable compliance with stiffening springs would have both benefits.

Acknowledgements We would like to thank Garth Zeglin, Amir Degani, Matt Bernstein, Alex Zirbel and Larry Hayhurst. This work was sponsored in part by the Army Research Laboratory under Cooperative Agreement Number W911NF-10-2-0016, DARPA, and NSF. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory, DARPA, or the U.S. Government. Alberto Rodriguez is the recipient of a graduate fellowship sponsored by Caja Madrid.

References

1. Barret Technologies: The Barret hand (2001). URL <http://www.barrett.com/robot/products-hand.htm>
2. Bicchi, A., Salisbury, J.K., Brock, D.L.: Contact Sensing from Force Measurements. *The International Journal of Robotics Research* **12**(3), 249 (1993)
3. Burges, C.J.: A Tutorial on Support Vector Machines for Pattern Recognition. *Data Mining and Knowledge Discovery* **2**, 121–167 (1998)
4. Carrozza, M.C., Suppo, C., Sebastiani, F., Massa, B., Vecchi, F., Lazzarini, R., Cutkosky, M.R., Dario, P.: The SPRING Hand: Development of a self-Adaptive Prosthesis for Restoring Natural Grasping. *Autonomous Robots* **16**(2), 125–141 (2004)
5. Ciocarlie, M., Allen, P.: Data-driven optimization for underactuated robotic hands. In: *International Conference on Robotics and Automation*, pp. 1292–1299. IEEE (2010)
6. Cleveland, W.S., Devlin, S.J.: Locally Weighted Regression: An Approach to Regression Analysis by Local Fitting. *Journal of the American Statistical Association* **83**(403), 596–610 (1988)
7. Cortes, C., Vapnik, V.: Support-vector networks. *Machine Learning* **20**(3), 273–297 (1995)
8. Dollar, A.M., Howe, R.D.: Joint coupling design of underactuated grippers. In: *30th Annual Mechanisms and Robotics Conference*, pp. 903–911. Harvard University (2006)
9. Dollar, A.M., Howe, R.D.: The Highly Adaptive SDM Hand: Design and Performance Evaluation. *The International Journal of Robotics Research* **29**(5), 585–597 (2010)
10. Dollar, A.M., Jentoft, L.P., Gao, J.H., Howe, R.D.: Contact sensing and grasping performance of compliant hands. *Autonomous Robots* **28**(1), 65–75 (2010)
11. Jacobsen, S., Wood, J., Knutti, D., Biggers, K.: The Utah/M.I.T. Dextrous Hand: Work in Progress. *The International Journal of Robotics Research* **3**(4), 21–50 (1984)
12. Kemp, C., Edsinger, A., Torres-Jara, E.: Challenges for Robot Manipulation in Human Environments [Grand Challenges of Robotics]. *IEEE Robotics & Automation Magazine* **14**(1), 20–29 (2007)
13. Laliberte, T., Birglen, L., Gosselin, C.: Underactuation in robotic grasping hands. *Machine Intelligence & Robotic Control* **4**(3), 1–11 (2002)
14. Mason, M.T., Srinivasa, S.S., Vazquez, A.S., Rodriguez, A.: Generality and Simple Hands. Technical Report CMU-RI-TR-10-40, Robotics Institute, Carnegie Mellon University (2010)
15. Quigley, M., Gerkey, B., Conley, K., Faust, J., Foote, T., Leibs, J., Berger, E., Wheeler, R., Ng, A.: ROS: an Open-Source Robot Operating System. In: *International Conference on Robotics and Automation*. IEEE (2009)
16. Robotiq: The Adaptive Gripper (2010). URL <http://robotiq.com/robot-hand/>
17. Smith, L.: A tutorial on principal components analysis (2002)
18. Theobald, D.A., Hong, W.J., Madhani, A., Hoffman, B., Niemeyer, G., Cadapan, L., Slotine, J.J., Salisbury, K.: Autonomous Rock Acquisition. In: *AIAA Forum on Advanced Developments in Space Robotics* (1996)
19. Ulrich, N.T.: Grasping with mechanical intelligence (1989)
20. Xu, Z., Deyle, T., Kemp, C.C.: 1000 Trials: An Empirically Validated End Effector that Robustly Grasps Objects from the Floor. In: *International Conference on Robotics and Automation*, pp. 2160–2167. IEEE (2009)