Contact Sensing and Grasping Performance of Compliant Hands

Aaron M. Dollar · Leif P. Jentoft · Jason H. Gao · Robert D. Howe

Abstract Limitations in modern sensing technologies result in large errors in sensed target object geometry and location in unstructured environments. As a result, positioning a robotic end-effector includes inherent error that will often lead to unsuccessful grasps. In previous work, we demonstrated that optimized configuration, compliance, viscosity, and adaptability in the mechanical structure of a robot hand facilitates reliable grasping in unstructured environments, even with purely feedforward control of the hand. In this paper we describe the addition of a simple contact sensor to the fingerpads of the SDM Hand (Shape Deposition Manufactured Hand), which, along with a basic control algorithm, significantly expands the grasp space of the hand and reduces contact forces during the acquisition phase of the grasp. The combination of the passive mechanics of the SDM Hand along with this basic sensor suite enables positioning errors of over 5cm in any direction. In the context of mobile manipulation, the performance demonstrated here may reduce the need for much of the complex array of sensing currently utilized on mobile platforms, greatly increase reliability, and speed task execution, which can often be prohibitively slow.

Keywords Robot · Manipulation · Mobile · Tactile Sensing · Shape Deposition Manufacturing

1 Introduction

The vision of robotic assistants for domestic, health care, and workplace applications will not come to fruition without the ability to reliably grasp and manipulate typical objects in human environments. In complex unstructured settings like the home or office, object properties are frequently unknown in advance, and visual sensing is prone to error. For mobile robotics applications, the challenge of grasping objects is further complicated by imprecise knowledge of the base location, compliant ground contact, limits on applied forces, and manipulator-environment coupling during grasping.

A. M. Dollar is with the Yale School of Engineering and Applied Science, Yale University, New Haven, CT, 06510 USA (e-mail: aaron.dollar@yale.edu).

L. P. Jentoft is with the Harvard School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138 USA (e-mail: ljentoft@seas.harvard.edu).

J. H. Gao is with the Harvard School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138 USA (e-mail: jgao@fas.harvard.edu).

R. D. Howe is with the Harvard School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138 USA (e-mail: howe@seas.harvard.edu).

The traditional approach to dealing with these challenges is to implement layers of sensing and control on complex multifingered hand hardware (e.g. Jacobsen et al. 1986; Butterfass et al. 2001), with a stated goal of achieving dexterous manipulation. Twenty five years and millions of dollars invested in this approach have failed to achieve dexterity; in fact, this approach has not even produced an effective general grasping device for unstructured environments. We have taken a different approach to dealing with the uncertainties which limit grasping in complex settings. Instead of sensing, planning, and control, we focus on the mechanics of the hand itself to accomplish most of the needed "control." We have demonstrated that combining carefully selected passive joint compliance and adaptive transmissions (Dollar and Howe 2005; Dollar and Howe 2006) allows the hand to passively adapt to the object, with successful object acquisition despite large errors in object localization. These features reduce the need for complicated and expensive sensing and control and make the hand easier to operate and more reliable. This approach has been experimentally validated (Dollar and Howe 2007) and implemented on mobile robot platforms (Breazeal et al. 2008).

While compliant hands can successfully grasp objects despite object localization errors, significant forces may be applied to the object if the errors are large due to deflections of the passive joint springs (Dollar and Howe 2007). Depending on the details of the object geometry, mass distribution, and friction, these forces may displace the object before it is grasped. Minimization of these forces would enable a wider range of object to be successfully grasped. One method for limiting forces is to use tactile sensing to detect the earliest stages of contact between the fingers and object and respond accordingly. These sensors would be used in the approach phase of grasping in order to detect that the object is not located at the estimated position, and to allow the arm to then center the hand on the object. This would result in a more stable grasp, larger grasp space for the hand, and lower unbalanced contact forces.

In this paper, we investigate the use of piezoelectric polymer contact sensors embedded in each fingerpad of our compliant robot hand to enhance the grasping process in unstructured environments. A vast number of tactile sensors have been developed for robotics research, although few have been integrated with robot hands and used for control of grasping or manipulation (Tegin and Wiklander 2005;

Jacobsen et al. 1988; Lee and Nicholls 1999; Howe and Cutkosky 1992; Fearing 1987);. Here we focus on the benefits that can be obtained from the simplest type of tactile sensor, which is essentially a low cost and robust contact detector. Our hypothesis is that the combination of contact sensing and a carefully tuned compliant hand can greatly expand the range of objects that can be successfully grasped.

We begin this paper with a brief review of the mechanics of our hand, followed by the design and analysis of the low-threshold contact sensor. We then present an experiment that evaluates the performance of the hand under varying degrees of uncertainty in the sensed object properties. This experiment demonstrates that the use of contact events enhances grasping capabilities compared to grasp performed in a purely feedforward manner.

2 SDM Hand

Before describing the experimental work that is the focus of this paper, we provide a brief overview of the design and function of the SDM Hand (Fig. 1). An extensive description can be found in (Dollar and Howe 2007). As the name suggests, the hand was fabricated using polymer-based Shape Deposition Manufacturing (SDM) (Mertz et al. 1994; Clark et al. 2001) to provide compliance and robustness. SDM is a layered manufacturing technique with which the rigid links and compliant joints of the gripper are created simultaneously with embedded sensing and actuation components. Elastomeric flexures create compliant joints, eliminating metal bearings, and tough rigid polymers fully encase the embedded components, eliminating the need for seams and fasteners that are often the source of mechanical failure.

The preshape, stiffness, and joint coupling characteristics of the hand were determined based on the results of previously conducted optimization studies (Butterfass et al. 2001; Dollar and Howe 2005). In these simulations, the joint rest angles, joint stiffness ratio, and coupling scheme of the hand were varied and the performance analyzed to maximize the allowable uncertainty in object location and size as well as minimize contact forces.

2.1 Finger design

The concave side of each finger link contains a soft fingerpad to maximize friction and contact area, thereby increasing grasp stability. Links are connected via elastomer joint flexures, designed to be compliant in the plane of finger motion and stiffer out of plane. Due to the molding process used to create them, the SDM fingers, with embedded sensors and actuation components, are a single lightweight part (39 grams each), with no fasteners or adhesives.

The polyurethane used for the joints of the fingers demonstrates significant viscoelastic behavior, providing both compliance and passive damping to the hand. The damping in the joints is necessary to reduce joint oscillations and permit the use of low joint stiffness. The joints can also undergo large deflections while remaining completely functional. The advantages of this property are clear when considering the damage that can result due to large contact forces that can occur with unplanned contact during use of traditional stiff robotic hands.

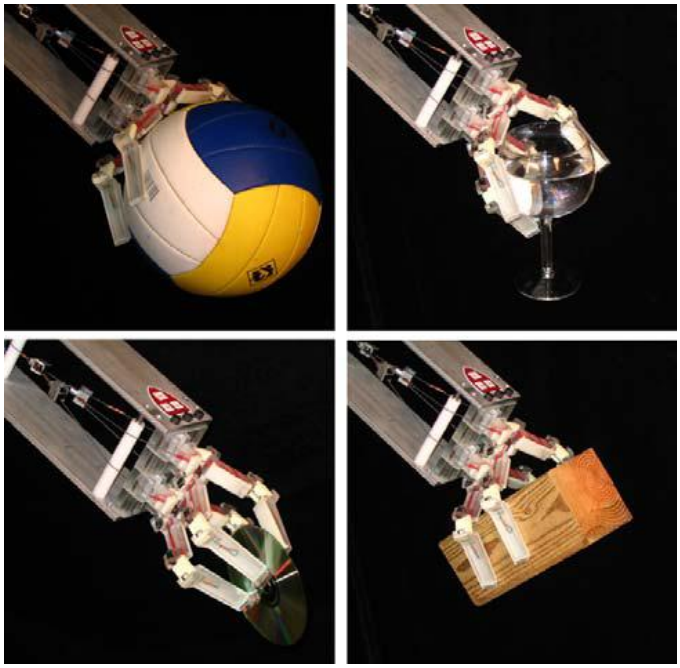

Fig. 1 SDM Hand grasping a variety of objects

2.2 Actuation

For actuation, each finger has a pre-stretched, nylon-coated stainless steel cable anchored into the distal link, and running through low-friction tubing to the base. The transmission of the hand is arranged such that the compliance in the fingers is in parallel with the actuator. Before the hand is actuated, the tendon cable remains slack and the finger is in its most compliant state. This method permits the use of actuators that are not backdrivable and prevents the inertial load of the actuator from increasing the passive stiffness. After actuation, the stiff tendon takes much of the compliance out of the fingers, resulting in a stiffer grasp with greater stability. This arrangement of the compliance in parallel with the actuation is a key factor in the effective performance of the hand.

A single actuator drives the four fingers (eight joints) of the hand. This not only makes the gripper simpler and lighter, but it also allows the gripper to be self-adapting to the target object. Fig. 2 details the actuation scheme, by which motion of the distal links can continue after contact on the coupled proximal links occurs, allowing the finger to passively adapt to the object shape. Additionally, the pulley design in this scheme allows the remaining fingers to continue to enclose the object after the other fingers have been immobilized by contact, ensuring that an equal amount of tension is exerted on each tendon cable, regardless of finger position or contact state.

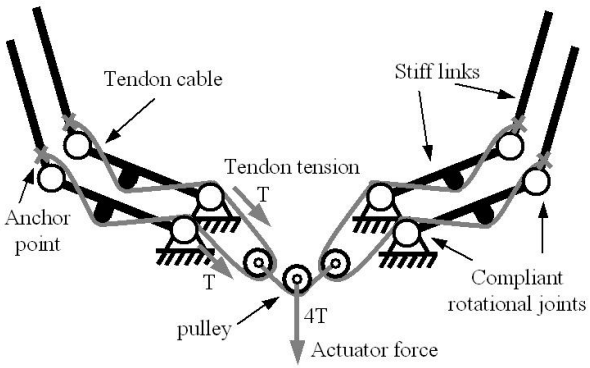

Fig. 2 Actuation schematic of the hand

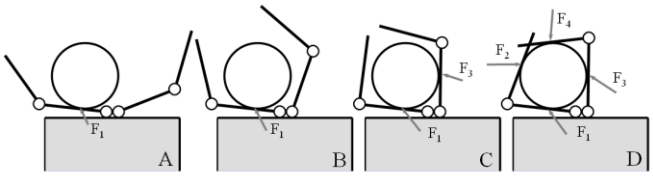

Fig. 3 Example grasp scenario

Note that the tendon cable is fixed only to the outer link of each finger, and freely moves over all other finger components without exerting torque or enforcing direct motion. This actuation scheme is similar to that used in (Hirose 1978).

Fig. 3 details an example grasp to demonstrate the adaptability of the transmission design. The grasper is unactuated until contact with a target object is sensed and a successful grasp is predicted based on any available sensory information. This initial contact may produce a small contact force (Fig. 3A). When the gripper is actuated, forces are exerted at the initial contact point while the second finger is brought into contact (Fig. 3B). Finger motion continues until the distal links on both fingers contact the object. Finally, the forces at the distal links increase as the grip on the object is secured (Fig. 3D). This process is completed in a purely feed-forward manner, with the actuator simply powered at a constant torque. A video detailing some of the performance characteristics of the hand can be seen at http://www.eng.yale.edu/adollar/SDM_Hand.avi.

3 Piezofilm Contact Sensor

One of the most important parameters to detect in manipulation or legged locomotion is the transition from noncontact to contact at the end effector. This event signals a change in the mechanical state of the robot-environment system and typically triggers a change in controller behavior. A wide variety of sensor can be used to accomplish contact detection. We selected a piezoelectric polymer film element (model DT1-028K/L, MSI sensors, Hampton, VA, USA, terminated with a 10MΩ load resistor) because of its high sensitivity, low cost, and excellent durability (Howe and Cutkosky 1993). These sensors are molded into the compliant

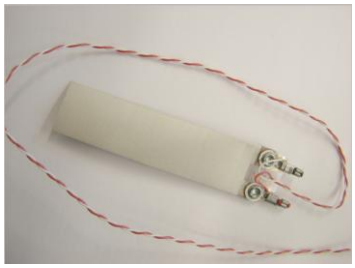

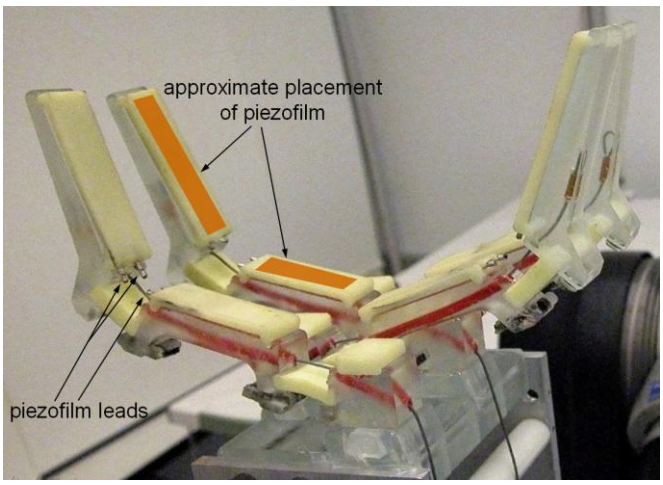

Fig. 4 Piezofilm element (top) and approximate placement within the fingerpads of the SDM Hand (embedded approximately 3mm below the surface)

fingerpads of the SDM Hand (Fig. 4). These sensors generate an electrical charge in proportion to the applied strain, have excellent frequency response and high sensitivity, but have no static response. By embedding the flexible sensor just under the contact surface, it senses the transient when the fingerpad is deformed on initial contact as well as when contact is removed. The sensor responds to all strain changes on the piezofilm element – forces applied normal to the fingerpad surfaces lead to bending strain and sheer forces lead to axial strain.

To determine the sensitivity and resolution of the sensor to contact transitions, a series of small loads placed on the fingerpad were quickly removed from the sensor with a fall time of under 10 ms. The loads were applied via a small spherical indenter (0.64 cm diameter). As shown in Fig. 5, these stimuli produced approximately 1.38 volts per Newton. The data points are the average of five trials (standard deviation shown as error bars). The RMS sensor noise was 21mV, or approximately 0.015N.

Fig. 6 shows a series of sensor responses to a typical grasping operation performed with the SDM Hand attached to a manipulator arm. The top plot shows three distinct contact events in which a fingerpad contacts an object during object acquisition. These events show an initial negative response at contact with a positive peak generated when the contact is removed. The height and sharpness of the peaks are dependent on how quickly the contact force is applied.

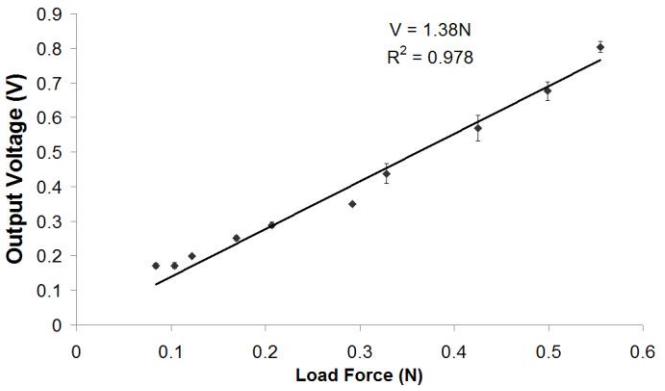

Fig. 5 Piezofilm sensor output vs. contact force with linear fit

In addition to the noncontact-contact transition, the sensor responds to changes in load on the finger surface during grasping and manipulation. The middle plot of Fig. 6 shows the sensor output as the fingers of the hand are closing around the object to secure the grasp, with the base of the hand remaining stationary. The signal has smaller amplitude due to the slower speed at which the fingers close. The oscillations seen in this signal are a result of vibrations induced as the remaining fingers contact and apply force to the target object.

The bottom plot in Fig. 6 shows the sensor response as the manipulator arm moves the object while grasped by the SDM Hand. The first transient shows the sensor response as the object is lifted off the table surface, where the changing load forces cause stress changes within the contact sensor. The portions of the signal marked "Motion up" and "Motion down" denote when the manipulator is moving the SDM Hand vertically up in the air and back down again, where small vibrations due to controller action are apparent. The final transient occurs when the object comes back into contact with the table.

The results of these tests with the embedded piezofilm contact sensor show that the sensor can rapidly respond to low force contact transients. This allows a manipulator to react quickly to minimize contact forces with the object or environment, yet still operate at a reasonable speed. Similar sensors have been developed for contact and transient detection, as well as perception of small shapes and incipient slips (Howe and Cutkosky 1996). Integration with the SDM fabrication process allows optimization of the overall finger mechanics and sensor response.

The reading from each sensor was converted to a signal/noise value and thresholded to yield a binary contact value for use by the positioning algorithm used in the following grasping study. The baseline noise value was calculated by averaging the absolute value of the sensor reading with a first-order IIR low-pass filter with a cutoff frequency of 0.1 Hz. The sensor readings during experiment were filtered to reduce noise with another first-order IIR lowpass filter (cutoff frequency 500 Hz) and then divided by the baseline noise reading to generate a signal/noise value appropriate for thresholding.

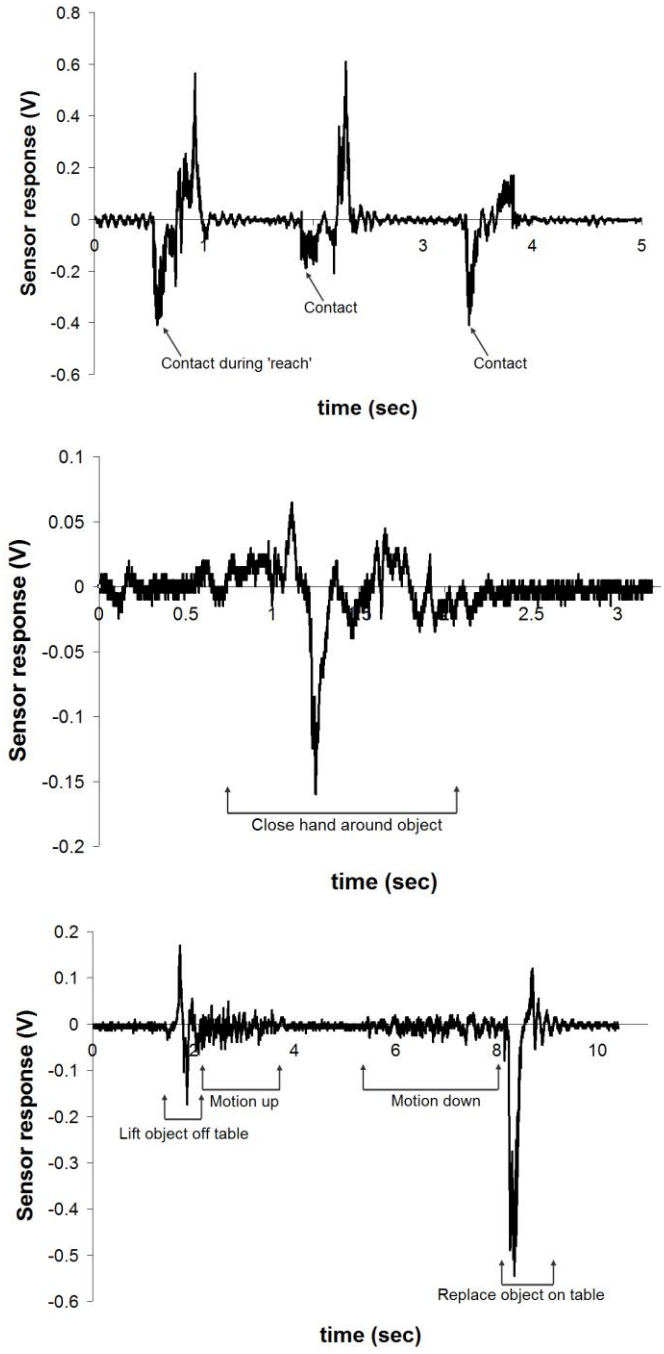

Fig. 6 Piezofilm contact sensor output for various phases of the grasping process: initial contacts during reach (top), increasing grasp force during object acquisition (middle), and internal forces during object lift and manipulation (bottom)

4 Experimental Setup

Using feedback from the contact sensors, we created an algorithm that uses contact with the target object to re-center the hand with respect to the target object given some initial positioning error. To evaluate its effectiveness in unstructured environments, we measured the ability of the algorithm to generate a successful grasp when a target object's actual

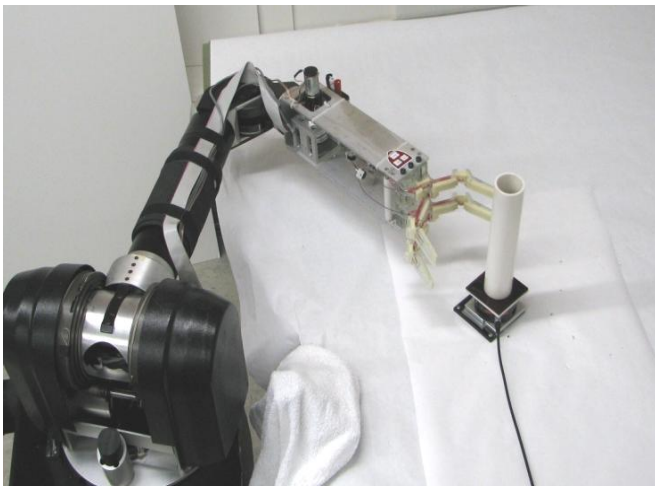

Fig. 7 Overhead view of WAM robot configuration. The table on the right half of the image is level with the ground and the arm operates in the plane until lifting the object

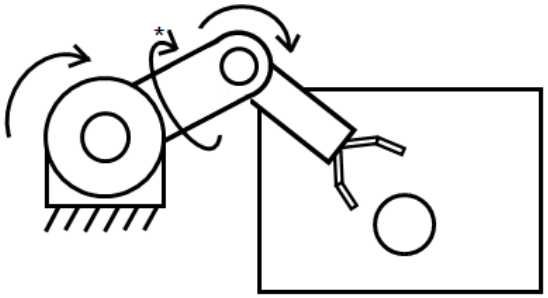

Fig. 8 Diagram of the robot manipulator arm with degrees of freedom. Arm operates in the plane of the work table until object is lifted, utilizing the shoulder roll joint, indicated by *

position is offset from its expected location. The results of the reactive algorithm are compared to those of a basic algorithm that merely grasps at the expected position of the object. Both algorithms are evaluated in terms of the grasp success and the magnitude of the planar force exerted on the object during the grasp.

4.1 Robot manipulator

The SDM Hand was mounted on a low-impedance robotic arm for positioning (Fig. 7) (Whole-Arm Manipulator (WAM), Barrett Technology, Cambridge, MA, USA). The robot was configured to operate in a planar configuration during the approach phase of the grasp, with the shoulder roll used to lift target objects after grasp (Fig. 8). Positioning commands were given in Cartesian coordinates and converted to trajectories in joint space, with a PID loop control running at 1000 Hz on a coprocessor (DS1103 PPC, dSpace Inc., Novi, MI). To increase performance and allow for the use of lower gains, the robot controller uses a feedforward model of the forces on the arm (before contact with the object),

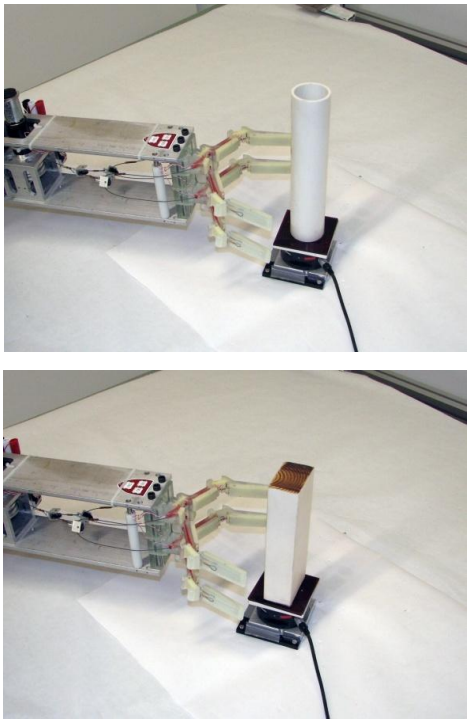

Fig. 9 Images of the two test objects and their orientations

including compensation for torque ripple, gravity, and friction. The arrival of the end-effector at a commanded position was defined as being within 1mm of the desired position according to the forward kinematics based on the joint angle readings.

Since there is no wrist, orientation of the hand was not controlled and was determined based on the kinematics of the manipulator at the target position.

4.2 Workspace

Two objects were tested with both the feed-forward and reactive sensor control algorithm: a 48mm diameter cylindrical PVC tube and a wood block with a cross-section of 38mm x 89mm, oriented with the wider face in the plane of the palm of the hand (Fig. 9). These objects were mounted on a 6-axis force/torque sensor (Gamma model, ATI Industrial Automation, Inc, Apex, NC, USA, 0.1 N resolution). This sensor is used to measure the contact forces on the objects during the grasping task. Planar forces were sampled at 1KHz; forces outside the plane of the workspace and torques were ignored, and a 20-sample (0.02s) median filter was applied to reduce noise.

Objects were mounted to the force sensor mount via a square peg, such that position and orientation in the plane were fixed, yet the object could be lifted up out of the mount after grasping. In actual unstructured grasping tasks, even small forces can dislodge some objects, particularly if they are lightweight or top-heavy. Predicting whether the object will move requires specification of detailed parameters such as mass distribution, three dimensional geometry, and frictional properties at the contact with the environment and with the

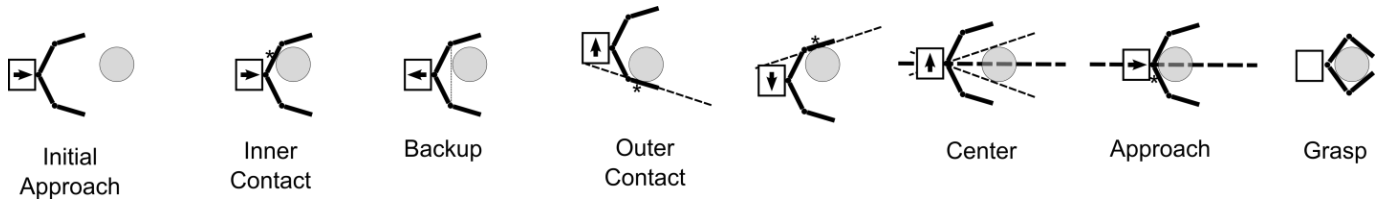

Fig. 10 Reactive control algorithm. If the initial contact is on the outer link, the process begins at the fourth cell, "outer contact". Robot motion is indicated by the arrow at the base of the hand. The comparable process for the Feed-forward algorithm is shown in Fig. 3.

fingers. This results in a large parameter space, and testing controller performance across this range is impractical.

Fortunately, it is not necessary to directly test the entire parameter space. By measuring the force applied by the hand to a fixed object, a prediction can be made as to whether an unfixed object might move for a given condition. The lower the applied force, the larger the range of objects that will not be moved, making applied force a good metric for grasping performance. For any given object, these experimental results can be used to predict if the object would have moved in a specific condition by comparing the force required to overcome friction and displace it with the experimental force on the "fixed" object.

Maximum force applied to the "fixed" object is then a conservative indicator of controller quality, since some objects might be successfully grasped even if a high enough force is applied to cause motion (e.g. if the object simply slides towards the other finger). Combining the maximum net force measure with the assumption that the object does not move reduces the parameter space to a tractable size but preserves the key result.

5 Experimental Procedure

The experiment begins by finding the 'zero position' for the particular object and location. This position was taken as the point at which the hand contacts the object without any deflection, centered on the object, representing the ideal positioning of the hand under perfect visual sensing (hand is centered on the object) and perfect contact sensing with zero manipulator inertia (allowing the manipulator to stop at the instant of initial contact) as in (Dollar and Howe 2007).

The *y* direction was taken as the normal vector to the palm of the hand at the zero configuration, with *x* being taken in the plane of the hand, parallel to the ground. To simulate errors in object location estimates that would occur in unstructured environments, the robot was positioned at 10mm increments from the zero position in the positive *x* (symmetry in the positive and negative *x* direction was assumed) and positive and negative *y* directions (grasping behavior is not symmetric in *y*). Forces on the object and whether the grasp was successful were recorded for each of these positions. In doing so, we evaluate the range of positions offset from the target object for which a successful grasp can be achieved, representing the allowable positioning error for the grasper and

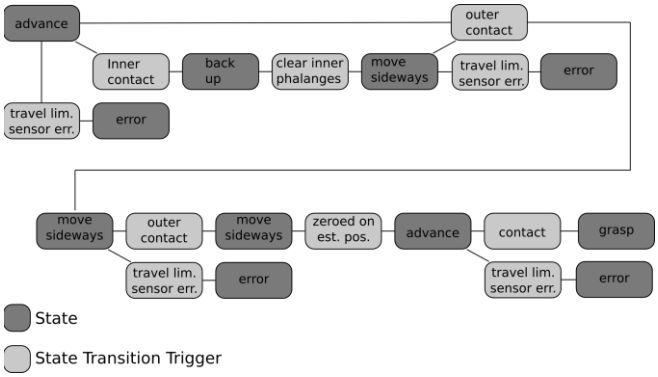

Fig. 11 State diagram showing the complete decision process for the reactive control algorithm used in this study

control algorithm. A successful grasp was defined as one where the object was able to be successfully lifted out of the force sensor mount without slipping out of the hand.

For each object, a fixed "start" position for the hand was calculated, offset from the object's zero position by 100mm in the *y* direction. This is the hand position from which the manipulator begins during each grasp trial, and from which it moves to each target location on the 10mm grid as described above.

Two grasp algorithms were tested. In the "feed-forward" algorithm, the hand moves to this target position and immediately closes the fingers, attempting to grasp the object and lift it out of the socket. This is the method utilized in (Dollar and Howe 2007).

The second algorithm, "reactive control", utilizes sensed contact with the target object to reposition the hand such that the object is centered in the grasp to increase stability of the grasp and balance contact forces. This algorithm is a straightforward implementation of more generalized frameworks for sensor-based control of robot hands (e.g. Tomovic et al. 1987; Brock 1993; Tremblay et al. 1995; Natale and Torres-Jara, 2006).

Fig. 10 describes our basic "reactive control" algorithm in which the hand is moved towards the target position until contact is registered on one of the fingerpad contact sensors. A more complete description can be seen in Fig. 11. The location of this contact is used to determine a line in the plane of the workspace that represents a bound on one edge of the object (Fig. 10). The hand is then moved in the *x* direction until contact is made on the opposing side of the hand, with the resulting contact location used to determine a second bounding edge of the object. The manipulator then centers the hand on

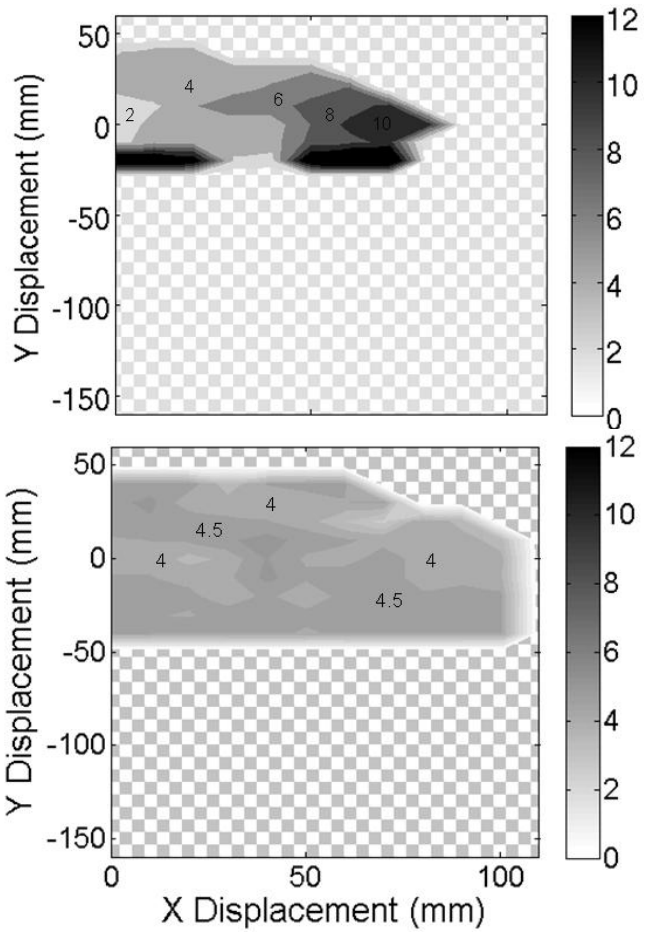

Fig. 12 Maximum force results for the cylindrical object. Top: "feed forward" algorithm, bottom: "reactive control." Contours are in Newtons, with various magnitudes called out for easy interpretation

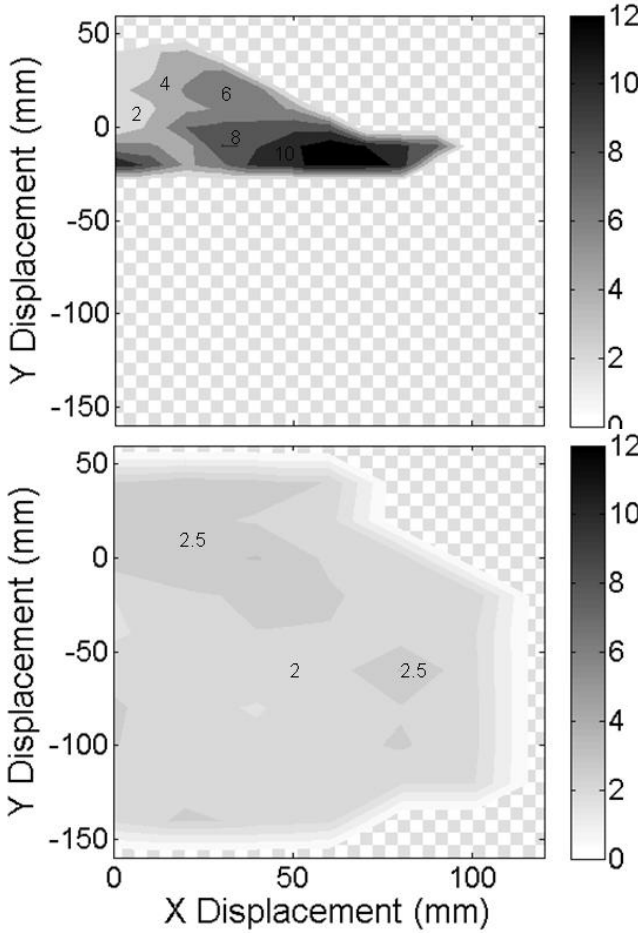

Fig. 13 Maximum force results for the rectangular object. Top: "feed forward" algorithm, bottom: "reactive control." Contours are in Newtons, with various magnitudes called out for easy interpretation

the bisector of these two lines (which contains the object's center for objects symmetric about the *y*-axis), and approaches until contact occurs a third time. At this point, the manipulator stops and attempts to grasp and lift the object, which is now more appropriately centered in the hand.

If the initial contact occurs on one of the inner segments, the manipulator is first backed up 5cm and then follows the same procedure. This is done in order to utilize the contact sensors on the distal finger links, which generated more reliable contact signals during motion in the *x*-direction due to their wider spacing left to right. For the proximal sensors, the manipulator velocity is still very low at contact on the opposing sensor (step five in Fig. 10) due to the close spacing of the proximal finger links and the manipulator control gains.

Note that abrupt contact with the target object sometimes triggered readings from multiple sensors, so a truth table was used as necessary to interpret whether these events are sharp collisions on one link of the hand or indeterminate contact with a larger region of the hand (generating an 'error' that was processed as an unsuccessful grasp).

6 Results

The results of the experimental study described above are shown in Fig. 12 for the cylindrical object, and Fig. 13 for the rectangular object. The top plot in each figure represents the results for the "feed-forward" algorithm and the bottom plot represents the results for the "reactive control" algorithm.

The horizontal and vertical axes of each plot correspond to the *x* and *y* axes as described above. Grasp success and contact force data was evaluated and recorded at 10mm increments from the zero position. Plot contours correspond to the magnitude of the force exerted during the grasp, as described by the colorbar to the right of each plot. The edges of the contoured areas correspond roughly to the edge of the effective grasp space, beyond which grasps were unsuccessful (and no force data exists). These areas are indicated by the hatched background.

Note that due to the large successful grasp range for the reactive algorithm with the rectangular object, positions were sampled at increments of 20mm, but were sampled at every 10mm for the other three cases.

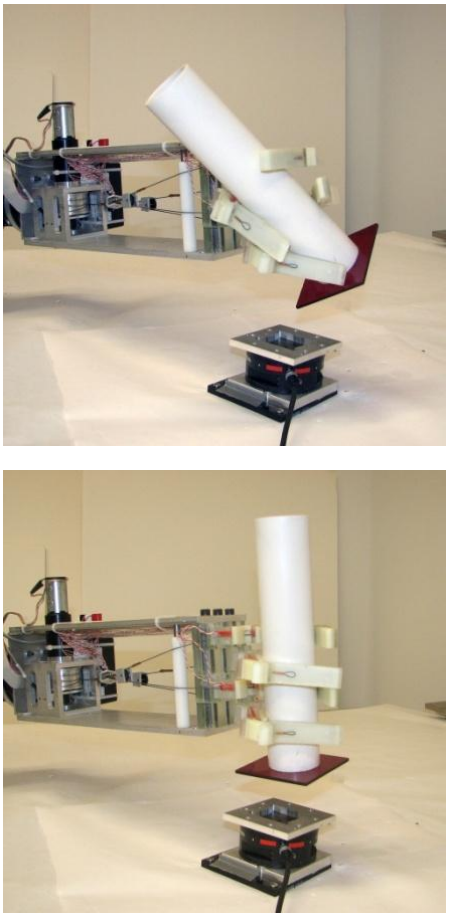

Fig. 14 Examples of poor quality (top) and good quality (bottom) grasps

7 Discussion

As expected, the addition of feedback from the contact sensors on the hand significantly decreases the forces applied to the object as it is grasped, as well as significantly increases the range of acceptable positioning offsets that still result in a successful grasp. In particular, the grasp space for the cylindrical object has been increased from approximately ± 80mm in *x* and -30 to +50mm in *y* to ± 120mm in *x* and ± 50mm in *y*. For the rectangular object, the grasp space was increased from approximately ± 90mm in *x* and -30 to +40mm in *y* to ± 120mm in *x* and -160mm to +60 mm in *y*. Put another way, the robot can cope with an initial object position estimate up to ±5cm away from its actual location in any direction (e.g. due to sensing error) for either of these objects and still get a successful grasp, utilizing only very basic sensing and control. Furthermore, unbalanced contact forces on the objects were limited to between 3-5 N for all successful grasp locations for the reactive control algorithm, whereas large regions of greater than double those values were observed under the feed forward control method.

For the "feed-forward" algorithm, the effective grasp region is bounded on the top and side (large offsets from the zero configuration) by the tendency of the object to slip out of the grasp because it is contacted by only the outer links of the fingers. On the bottom edge, the range is limited by the force exerted on the object as the arm approaches and grasps (i.e. the robot tries to push the hand through the object, dislodging it from its rest position).

For the "reactive control" algorithm, the lower edge of the effective grasp space is limited by poor sensor readings at contact with the object. The grasp space is much larger for the rectangular block due to a stronger object edge contacting the sensor. The upper edge of the range is only limited by the reach of the manipulator arm. On the side, it is simply limited by the width of the grasper (100mm). There is, however, regions of "successful grasps" beyond this due to the oblique approach caused by the fixed starting position, but this data does not add useful information since it suggests that the hand could detect objects wider than the hand itself.

Besides the performance improvements reflected in Figs. 12 and 13, the quality of the grasp for the reactive control was visably better over much of the space than for feed-forward control. An example of this effect is shown in Fig. 14. Although the object in the grasp does not drop and the grasp is thus judged "successful" in our classification, it has been, perhaps unacceptably, shifted to an awkward orientation and is less robust to disturbances during the manipulation.

During the experiments it became clear that manipulator inertia dominates the forces applied to the object during the approach phase. Contact was able to be sensed at a very low force threshold, but by the time the manipulator was able to be stopped, the applied force rose substantially. Control gains and approach strategy should be carefully considered in order to minimize manipulator velocity when contact with a target object is immenent.

7.1 Future Work

An immediate direction for future work is the enhancement of the sensory suite to facilitate better contact detection. Extension of the piezofilm to cover the fingertips is an obvious improvement. This would enable distinguishing between head-on contact with the center of an object and sideways contact with the edge of it. The addition of some type of joint-angle sensing would also be adventageous. This would allow the reactive control algorithm to detect missed contacts by monitoring the deflection of the finger joints. Sensors of this type would be especially helpful in cases where contact forces are not large enough to register on the piezofilm sensors, as can be the case when the velocity of the manipulator is particularly slow. As an alternative, we will also investigate contact sensors with a static response that are less likely to miss the contact transients during slow manipulator movements.

More generally, we would like to determine the tradeoffs of different combinations of sensory suites for robotic grasping. By categorizing and evaluating sensor types according to the nature of the information they provide about the environment

and target object properties, we hope to gain insight into the fundamental object and contact properties that are required to for different levels of functionality.

8 Conclusions

These results demonstrate that, because the hand is compliant, even the most basic form of tactile sensing can minimize forces and maximize grasp range in unstructured environments. Our previous work showed that a compliant hand can successfully grasp despite significant object location errors, but contact forces can be relatively large because the joint springs have to deflect to account for the error (Dollar and Howe 2007). The addition of contact sensing in effect allows the controller to refine the estimate of the object location. Simple binary contact sensing does not provide a very precise estimate of the location – that would require more elaborate sensors and signal processing (e.g. Son et al. 1996, Fearing and Binford, 1991). Such sensing is expensive and difficult to implement in a form that is sufficiently robust for unstructured environments. But because the hand's compliance passively adapts to the object location and shape, the small residual errors in object location after recentering on the object generate only small forces.

These benefits from simple contact sensing would not accrue to a stiff grasper. Small errors in object position would generate large forces unless the controller precisely adjusted the joint configuration. This would be problematic due to the finite force sensing threshold and the various time delays associated with sensing and control (sensor readout and processing time, deceleration of the arm inertia, etc.).

For grasping on mobile platforms (e.g. Khatib 1999, Saxena et al. 2008, Kemp et al. 2008), object model estimates from imperfect sensing and imprecise knowledge of the mobile base and arm positions often lead to large positioning errors of the robot and end-effector. The resulting grasping process is therefore typically unreliable and/or exceedingly slow. The combination of hand compliance with simple contact sensors as described in this paper can address these performance limitations of mobile grasping systems and speed their implementation in domestic and workplace environments.

Acknowledgement This work was supported with funding from the Honda Research Institute USA.

References

Breazeal, C. et al. (2008). Mobile, dexterous, social robots for mobile manipulation and human-robot interaction. *International Conference on Computer Graphics and Interactive Techniques, ACM SIGGRAPH 2008, new tech demos*.

Brock, D. L. (1993). A sensor-based strategy for automatic robotic grasping. Doctoral Dissertation, Massachusetts Institute of Technology.

Butterfass, J., Grebenstein, M., Liu, H., Hirzinger, G. (2001). DLR-Hand II: Next Generation of a Dextrous Robot Hand. *Proceedings of the 2001 IEEE International Conference on Robotics and Automation (ICRA 2001)*, pp. 109-114

Clark, J. E. et al. (2001). Biomimetic design and fabrication of a hexapedal running robot. *Proceedings of the 2001 International Conference on Robotics and Automation (ICRA 2001)*

Dollar, A. M. & Howe, R. D. (2005). Towards grasping in unstructured environments: Grasper compliance and configuration optimization. *Advanced Robotics*, 19(5), 523-544.

Dollar, A. M., & Howe., R. D. (2006). Joint Coupling Design of Underactuated Grippers. *Proceedings the 30th Annual ASME Mechanisms and Robotics Conference, 2006 International Design Engineering Technical Conferences (IDETC)*.

Dollar, A. M., & Howe., R. D.(2007). Simple, Robust Autonomous Grasping in Unstructured Environments. *Proceedings of the 2007 International Conference on Robotics and Automation (ICRA 2007)*

Fearing, R. S., & Binford, T. O. (1991). Using a cylindrical tactile sensor for determining curvature. *IEEE Transactions on Robotics and Automation*, 7(6), 806-817

Fearing, R.S. (1987). Some Experiments with Tactile Sensing during Grasping. *Proc. of the 1987 Intl. Conf. on Robotics and Automation (ICRA 1987)*, pp. 1637-1643.

Hirose, S., & Umetani, Y. (1978). The Development of Soft Gripper for the Versatile Robot Hand. *Mechanism and Machine Theory*, 13, 351-359.

Howe, R. D. & Cutkosky, M. R. (1992) "Touch sensing for robotic manipulation and recognition," in O. Khatib et al., (Ed.), *Robotics Review 2* (pp. 55-112) Cambridge: MIT Press.

Howe, R. D., & Cutkosky, M. R. (1993). Dynamic tactile sensing: Perception of fine surface features with stress rate sensing. *IEEE Trans. on Robotics and Automation*, 9(2):140-151

Jacobsen, S. C., Iversen, E. K., Knutti, D. F., Johnson, R. T., Biggers, K. B. (1986). Design of the Utah/MIT dextrous hand. *Proceedings of the 1986 IEEE Int. Conf. on Robotics and Automation (ICRA 1986)*, pp. 1520-1532.

Jacobsen, S. C., McCammon, I. D., Biggers, K. B., Phillips., R. P. (1988). Design of Tactile Sensing Systems for Dextrous Manipulators. *IEEE Control Systems Magazine*, 8(1), 3-13.

Kemp, C. C., Anderson, C. D., Nguyen, H., Trevor, A. J., Xu, Z. (2008). A Point-and-Click Interface for the Real World: Laser Designation of Objects for Mobile Manipulation. *Proceedings of the 3rd ACM/IEEE international conference on Human robot interaction*, 241-248

Khatib, O. (1999). Mobile Manipulation: The Robotic Assistant. *Robotics and Autonomous Systems*, 26, 175-183.

Lee, M. H. & Nicholls., H. R. (1999). Tactile Sensing for Mechatronics: A State of the Art Survey. *Mechatronics*, 9(1), 1-31

Merz, R., Prinz, F. B., Ramaswami, K., Terk, M., Weiss, L. (1994). Shape Deposition Manufacturing. *Proceedings of the Solid Freeform Fabrication Symposium, Univ. of Texas at Austin*

Natale, L. & Torres-Jara, E. (2006). A sensitive approach to grasping. *Sixth International Conference on Epigenetic Robotics*

Saxena, A., Driemeyer, J., Ng, A. Y., (2008). Robotic Grasping of Novel Objects using Vision. *International Journal of Robotics Research*, 27(2), 157-173.

Son, J. S., Cutkosky, M. R., & Howe, R. D. (1996). Comparison of contact sensor localization abilities during manipulation. *Robotics and Autonomous Systems*, 17(4), 217-233

Tegin, J., & Wikander., J. (2005). Tactile Sensing in Intelligent Robotic Manipulation – a review *The Industrial Robot*, 32(1), 64-70.

Tomovic, R., Bekey, G., & Karplus, W. (1987). A strategy for grasp synthesis with multifingered robot hands. *Proceeding of the 1987 IEEE International Conference on Robotics and Automation. (ICRA 1987)*, 4, 83-89

Tremblay, M. R., Hyde, J. M., Cutkosky, M. R. (1995). An Object-Oriented Framework or Event-Driven Dextrous Manipulation. *Proceedings of the 4th Intl. Symposium on Experimental Robotics*