
Aaron M. Dollar

Department of Mechanical Engineering
Yale University
15 Prospect Street
New Haven, CT 06511, USA
adollar@yale.edu

Robert D. Howe

School of Engineering and Applied Sciences
Harvard University
Pierce Hall 323
Cambridge MA 02138, USA
howe@seas.harvard.edu

The Highly Adaptive SDM Hand: Design and Performance Evaluation

Abstract

The inherent uncertainty associated with unstructured environments makes establishing a successful grasp difficult. Traditional approaches to this problem involve hands that are complex, fragile, require elaborate sensor suites, and are difficult to control. Alternatively, by carefully designing the mechanical structure of the hand to incorporate features such as compliance and adaptability, the uncertainty inherent in unstructured grasping tasks can be more easily accommodated. In this paper, we demonstrate a novel adaptive and compliant grasper that can grasp objects spanning a wide range of size, shape, mass, and position/orientation using only a single actuator. The hand is constructed using polymer-based Shape Deposition Manufacturing (SDM) and has superior robustness properties, making it able to withstand large impacts without damage. We also present the results of two experiments to demonstrate that the SDM Hand can reliably grasp objects in the presence of large positioning errors, while keeping acquisition contact forces low. In the first, we evaluate the amount of allowable manipulator positioning error that results in a successful grasp. In the second experiment, the hand autonomously grasps a wide range of spherical objects positioned randomly across the workspace, guided by only a single image from an overhead camera, using feed-forward control of the hand.

KEY WORDS—adaptive, compliant, design and control, grasping, manipulation, mechanics, mechanism design, multifingered hands, robot, underactuated, underactuated robots, unstructured environments

1. Introduction

Grasping and manipulating objects in unstructured environments, where object properties are not known *a priori* and sensing is prone to error, is one of the central challenges in robotics. The uncertainty in the relationship between the object and gripper makes it difficult to control contact forces and establish a successful grasp, particularly with traditional stiff robot end effectors. While robot hand research has been largely focused on dexterous manipulation, robots today cannot autonomously perform even simple grasping tasks in a typical home setting.

One approach to dealing with this uncertainty is through compliance and passive adaptability, so that positioning errors do not result in large forces and the grasper conforms to the object. These features have most often been implemented through control of manipulator impedance, based on active use of joint sensors for position, velocity, and force (e.g. Salisbury 1980; Cutkosky and Kao 1989; Desai and Howe 2001). However, carefully designed passive *mechanical* compliance and adaptability in the finger and hand structure can allow the gripper to conform to a wide range of objects while minimizing contact forces.

An adaptive, underactuated hand has fewer actuators than degrees of freedom. In these hands, the transmission design allows motion of other joints to continue after contact occurs

The International Journal of Robotics Research

Vol. 29, No. 5, April 2010, pp. 585–597

DOI: 10.1177/0278364909360852

© The Author(s), 2010. Reprints and permissions:

<http://www.sagepub.co.uk/journalsPermissions.nav>

Figures 1–3, 6, 7, 9, 12, 13, 15, 16 appear in color online:

<http://ijr.sagepub.com>




Fig. 1. Four-fingered, underactuated SDM hand mounted on a WAM (Barrett Technology Inc., Cambridge, MA, USA). A single motor drives all eight joints of the hand.




Fig. 2. Relative size of the SDM Hand.

on a coupled link, allowing the hand to passively adapt to the object shape. While a small number of compliant, underactuated hands have been previously proposed (e.g. Ulrich and Kumar 1988; Laliberte et al. 2002; see Dollar 2006 for a thorough review of adaptive robot hands), none have demonstrated the level of adaptability, ease of use, and reliability that we demonstrate here.

In this paper, we begin by describing the design and fabrication of our highly adaptive four-fingered grasper (Figures 1 and 2) built using Shape Deposition Manufacturing (SDM) (Merz et al. 1994; Clark et al. 2001; Dollar and Howe 2006). This process uses polymeric materials to simultaneously create the rigid links and compliant joints of the gripper, with embedded sensing and actuation components. In addition to simplifying the construction process, the result is a robust gripper, fully functional after impacts and other large loads due to unintended contact. We also describe the highly adaptive joint coupling and transmission design that reduces the need for complicated sensing and control through passive adaptation to the object physical properties, making the hand easier to operate and with greater reliability.

We then describe the results of two experimental studies in which we evaluate the ability of our grasping system to autonomously grasp a number of target objects in the presence of varying levels of positional error. These studies show that the SDM Hand can successfully grasp objects even in the presence of large positioning errors and with the simplest control.

2. SDM Hand Design

Consider a robot designed for the following unstructured grasping task: a manipulator arm is mounted on a small mobile platform for operation in a home environment. The task is perhaps to navigate through the home, pick up stray drinking glasses, and place them in a dishwasher. There are, of course, numerous complicated subtasks required in order to execute the desired functionality, such as robot navigation and planning, object recognition, and property estimation.

In terms of hand design, we are focused on the aspects of the task after the object has been identified and a target position/orientation estimate has been generated. In particular, we would like to provide a grasper such that this estimate does not have to be very accurate – the robot simply positions the hand “close enough” to the target object, issues an open-loop “close” command, and the passive mechanics of the gripper take care of the rest. To accomplish this, we carefully design the mechanical structure and transmission of the hand to appropriately incorporate compliance, adaptability, and durability to handle the uncertainty inherent with unstructured grasping tasks.

To provide both adaptability and robustness, our hand, featuring passively compliant joints, was fabricated using polymer-based SDM (Merz et al. 1994; Clark et al. 2001, Dollar and Howe 2006). 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.




Fig. 3. Details of finger parts and placement of components. Joint angle and fingerpad contact sensors are not labeled as they are not utilized in this study.

The fingers are staggered in the out-of-plane direction on the palm to allow them to completely close without interfering with one another (Figure 1). The total hand thickness is 93 mm, and the finger tip aperture is 170 mm, approximately that of a large human hand (Figure 2).

2.1. Finger Design

We begin with a short description of the fingers used on the hand (Figure 3), which are nearly identical to the fingers used in a previously developed gripper (Dollar and Howe 2006). The concave side of each link contains a soft fingerpad to maximize friction and contact area, thereby increasing grasp stability (Cutkosky et al. 1987; Shimoga and Goldenberg 1992). Links are connected via elastomer joint flexures, designed to be compliant in the plane of finger motion and stiffer out of plane.

The two links of each finger are 70 mm long (measured from the center of the joint flexures). Due to the molding process used to create them (described in detail in Dollar and Howe 2006), the SDM fingers, with embedded sensors and actuation components, are a single part weighing 39 g, with no fasteners or adhesives. This is in contrast to a similar finger design that was fabricated with conventional metal prototyping techniques used in previous work, which had 60 parts in total, including 40 fasteners, and weighed 200 g (Dollar and Howe 2005).

2.1.1. Finger Compliance and Robustness

The polyurethane used for these joints 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. When released after a large displacement of the fingertip (through the entire base joint range of motion of 0.5 radians from the rest position), joint oscillations are negligible after less than 1 s. In a conventionally fabricated grasper with metal springs (Dollar and Howe 2005), oscillations due to large step displacements were found to persist for tens of seconds after release.

The approximate tip stiffness in the x , y , and z directions (according to the convention in Figure 3) are 5.9, 7.7, and 14.2 N/m, respectively.

Due to its compliance and polymeric construction, the SDM Hand is exceptionally robust. The tip of the SDM finger can be displaced more than 3.5 cm in the out-of-plane direction (approximately 20°) without any degradation of mechanical properties. As a result, the SDM fingers, while exhibiting very low tip stiffness, 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 and displacements that can occur with unplanned contact during use of traditional research robotic hands. The mechanical behavior of the SDM joints and materials has been evaluated in depth by Dollar and Howe (2006).

To give a sense of the robustness of the finger mechanism to impacts and other potentially harmful loads, a number of more informal tests were performed. A SDM finger was repeatedly dropped from a height of over 15 m onto a stone floor, without significant damage. The fully assembled hand has been hit repetitively with a hammer, fingers jammed against objects, and even used underwater, without degradation of performance.

2.1.2. Kinematic and Stiffness Configuration

The preshape and stiffness characteristics of the hand were determined based on the results of a previously conducted optimization study (Dollar and Howe 2005). In this simulation, the joint rest angles and joint stiffness ratio of the fingers were varied and the performance analyzed to maximize the allowable uncertainty in object location and size as well as minimize contact forces.

Based on the results of this study, the preshape configuration of 25° for the proximal joint (angle with the palm of the hand) and 45° for the distal joint (angle with the proximal link) was chosen for our final finger design. In addition, the results showed that the distal joint should be much stiffer than the proximal joint, keeping the grasping surface concave and contact forces low. These angles and stiffnesses were shown to enable grasping of the widest range of object sizes with the greatest amount of uncertainty in object position, while also exhibiting low average contact force, thereby reducing the likelihood of displacing or damaging the object.




Fig. 4. Actuation schematic of the hand.

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 nylon 11 tubing to the base (Figure 3). 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 property not only makes the gripper simpler and lighter, but it also allows the gripper to be self-adapting to the target object. Figure 4 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. In addition, 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. Note that the tendon cable is fixed only to the outer link of each finger, and freely moves over all other finger components without directly exerting torque or enforcing direct motion. This actuation scheme is similar to that used by Hirose and Umetani 1978.

Figure 5 details a simplified example grasp (in a planar approximation of the hand) to further describe 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 ini-




Fig. 5. Example grasp scenario.

tial contact may produce a small contact force (Figure 5(a)). When the gripper is actuated, forces are exerted at the initial contact point while the second finger is brought into contact (Figure 5(b)). 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 (Figure 5(d)). This process is completed in a purely feed-forward manner, with the actuator simply powered at a constant torque.

A demonstration of the adaptability of the hand can be seen in Extension 1.

2.2.1. Joint Coupling Design

The joint coupling scheme employed on each finger was determined based on the results of a previously conducted optimization study (Dollar 2006). In this simulation, the joint coupling scheme (the ratio of torque applied at the distal/proximal joints divided by the stiffness ratio of the joints) was varied in order to maximize the allowable uncertainty in object location and size as well as minimize contact forces.

The results of this study suggested that, to keep unbalanced object forces low, the torque ratio (ratio of torques applied at the distal and proximal joints) should be as large as possible. However, as the torque ratio increases, the position range in which an object can be successfully grasped (the maximum allowable positioning error) is decreased. This tradeoff in force versus successful grasp range was weighed by considering the quality of the sensory information available for the grasping task and an optimum distal/proximal torque ratio of approximately 1.0 was determined.

2.3. Informal Performance Demonstration

Figure 6 exemplifies the adaptability of the hand to accommodate variations in object geometry and location. Under noisy sensing conditions, the gripper may not be able to be properly centered on the target object due to object location uncertainty, resulting in configurations such as those seen in the images on the right in Figure 6. By passively accommodating for this positioning error, the hand increases the robustness of the grasping task.



Fig. 6. Adaptability of the hand to variations in object location (i.e. centered and off-centered).

The images in Figure 7 demonstrate the utility of the hand in grasping everyday objects. Many of these objects are involved in challenging tasks suggested in the prosthetics literature as “practice” objects on which a recent amputee should eventually learn to grasp as they train (Klopsteg et al. 1968). Others were grasped to demonstrate the range of size, shape, and mass of objects that can be successfully grasped using the SDM Hand. These objects are grasped with only a single direct current (DC) motor for actuation, without the aid of any sensory feedback. The motor is simply run to stall, and the passive adaptability designed into the hand and transmission drives the joints to a position that conforms to the given object shape. The method of achieving this passive adaptability is described in Section 2.2.

In addition to demonstrating the ability of the hand to grasp a wide range of objects in the manner such as is shown in Figure 7, we have done numerous teleoperated grasping tasks with objects placed randomly on a table in the workspace of the robot arm. Extension 1 demonstrates a number of these tasks, in addition to showing the robustness and adaptability of the hand.

3. Experimental Evaluation

In order to determine the effectiveness of our hand at grasping objects in unstructured conditions, we wish to experimentally evaluate the performance of the hand in the presence of significant errors in the sensed target object location and a very simple control scheme.

It is intrinsically difficult to ascertain performance in an unstructured environment because so many parameters can vary, e.g. object size, shape, friction, weight, and mass distribution, as well as errors in estimating object location and pose. As far as we are aware, there have been no previous studies in which the performance of a robotic hand has been quantitatively evaluated for grasping in the presence of uncertainty. There is therefore no commonly accepted representative test to evaluate performance.

In the following sections we describe two experiments used to quantify the performance of the mechanical hand, decoupled from sensing and control considerations. In the first, we structure the environment such that we know the exact position of the target objects, and then proceed to determine the range of positions around the target object for which a successful grasp can be attained. For a robot grasper to be used in unstructured environments where the sensed object properties may not be well known, this position range should be large, corresponding to the amount of allowable position estimation error, for example from vision.

In the second experiment we create a grasping task that is meant to echo the type of uncertainty that might be seen in a typical unstructured grasping task and to demonstrate a further level of autonomy in the grasping system. We begin with an object of unknown size and location, and extract a target position estimate based on noisy sensor data (a single image from an overhead universal serial bus (USB) camera). The grasping task is then executed, the success rate recorded, and the amount of positioning error evaluated using a more accurate analysis of the object properties.

As mentioned, these two experiments are designed to demonstrate the performance of the SDM Hand decoupled from sensing, control, and other system considerations. The specific sensing and processing we choose is intentionally simple in order to utilize the least amount of information that might enable a target position estimate. Furthermore, all grasping tasks are performed in an open-loop, feed-forward manner. The hand is simply placed at the target location and closed, with the adaptive transmission (described in Section 2.2) securing the grasp. As we will show, the SDM Hand can successfully grasp objects even in the presence of large positioning errors and with the simplest control.

3.1. Grasp Range and Contact Force Experiment

In this experiment we determine the amount of allowable positioning error of the SDM Hand during grasp for two different objects placed at various locations in the robot workspace. Object contact forces were also measured at the tested locations. For unstructured grasping tasks, the range of positions for which the hand can successfully grasp the object should be large, maximizing the allowable uncertainty or error in the task, and contact forces should be small so as not to displace or damage the target object.




Fig. 7. SDM Hand grasping various household objects, a number of which are suggested in Klopsteg et al. (1968) as “practice” objects. (a) phone receiver, (b) full wine bottle, (c) full wine glass, (d) compact disc grasped on edges, (e) long rod (e.g. a broom), and (f) cordless drill.

3.1.1. Experimental Setup

The SDM Hand was mounted on a low-impedance robotic arm (Whole-arm Manipulator (WAM), Barrett Technology, Cambridge, MA, USA) for positioning (Figure 1). Only three of the four joints of the WAM were used for a total of three positioning degrees of freedom: the base roll, shoulder pitch, and elbow pitch (Figure 8). 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. Target positions were achieved to within 2 mm accuracy.

The WAM was controlled using a 1,000 Hz servo loop running on a digital signal processing (DSP) co-processor board (DS1103 PPC, dSPACE Inc., Novi, MI). The desired position was achieved using a packet identifier (PID) controller with gains chosen so that the overall stiffness was dominated by the remote environment stiffness. To increase performance and allow for the use of lower gains, the robot controller uses a feed-forward model of the forces on the arm (before contact with the object), including compensation for torque ripple, gravity, and friction.

3.1.2. Workspace

Target objects were mounted on a six-axis force/torque sensor with a resolution of 0.1 N (Gamma model, ATI Industrial




Fig. 8. Schematic diagram of the grasper mounted on the WAM robot arm, with the three arm degrees of freedom indicated.

Automation, Inc, Apex, NC, USA). Objects were mounted to the force sensor 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. Only contact forces in the plane of the workspace table were recorded, and torques were ignored. Robot inertial effects were minimized by using low accelerations during motion, reducing the task to nearly quasi-static conditions.

Two objects were tested at three positions, for a total of six conditions (Figure 9). The objects were a cylindrical




Fig. 9. Two target objects (PVC cylinder with radius 24 mm and wood block with square cross-section 90 mm side length) at three locations ((a), (b), and (c)). Note the differences in approach angle for the locations, the main factor affecting the force and grasp space results.

polyvinyl chloride (PVC) tube with a radius of 24 mm (0.3 times the grasper link length l), and a wooden block with a 84 mm \times 84 mm cross section (equivalent to 0.75 times the grasper link length l). This block was oriented such that a flat side was approximately normal to the approach direction. As shown in Figure 8, the difference in object position served to change the approach angle of the grasper with respect to the long axis of the objects, ranging from 25.6° to 42.8°.

3.1.3. Experimental Procedure

The experiment begins by manually 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; this represents the positioning of the hand under perfect visual sensing (the hand is centered on the object) and perfect contact sensing (stopping the manipulator at the instant of initial contact). The y direction was taken along the line lying between the robot origin and the center of the object, normal to the direction of gravity. The x direction is normal to the y direction, and also normal to the direction of gravity (the z direction).

In order to examine the behavior of the grasping system for a range of “error” in positioning, a grid of positions from the zero position was calculated. The performance of the hand was tested at 10 mm 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 (the grasping behavior is not symmetric in y).

The manipulator joint angles were calculated using the inverse kinematics of the robot and rounded to the nearest tenth of a degree. For each position on the grid, the robot moves to within a tenth of a degree of the target joint configuration at each joint. The robot then initiates the grasp by driving the grasping motor to a preset torque (stall) and thus closing all fingers. When an encoder indicates motor stall, the motor current is reduced to the small amount required to prevent back-driving of the motor due to the tendon force. The arm then attempts to lift the object vertically out of the force sensor mount. Forces on the object and whether the grasp was successful were recorded for each position. The vertical position of the hand was kept constant across the object’s position at approximately 19 cm above the table (Figure 9). The sensors on the hand are not used in this study. This simple, strictly feed-forward hand control mode is used to evaluate the benefits of the optimized passive compliance and adaptive coupling approach to hand design.

Each location on the (x, y) grid of positions was tested three times, and the force results averaged. Force was recorded at 1,000 Hz during the experiment. Data from the force sensor was filtered by taking the median of the previous 20 force samples (0.02 s).

A grasp was deemed successful if the object was lifted vertically out of the force sensor mount a distance of 150 mm, and the grasp appeared to be stable (i.e. no slippage of the object was visually observed). Grasps could fail at a given position for a number of reasons: passive contact force pushes the object out of the sensor mount or pushes the sensor out of the table mount, too few fingers make contact with the object, or an imbalance of forces on the object due to undesirable positioning leads to it being ejected from the grasp.

3.1.4. Results

Figures 10 and 11 show the results of the force and successful grasp space study for the two objects at three configurations each. The left column (F_{approach}) indicates the magnitude of the maximum force applied to the object during the approach phase of the grasp (the hand has not yet been actuated). The right column (F_{grasp}) indicates the magnitude of the maximum force applied to the object during the grasp phase (fingers are closing in on the object, before the motion of the arm to lift the object out of the sensor mount).

The various points on the plots that are labeled correspond to interesting or demonstrative configurations. A description

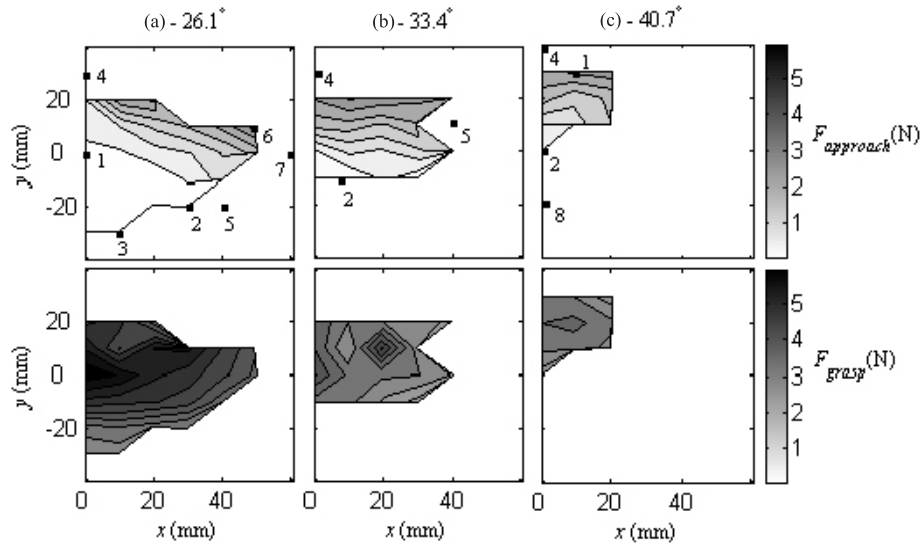


Fig. 10. Forces on the PVC cylinder object during the approach (top row) and grasp (bottom row) phases for the three object locations (columns). Labeled configurations correspond to the behavior indicated in Table 1.

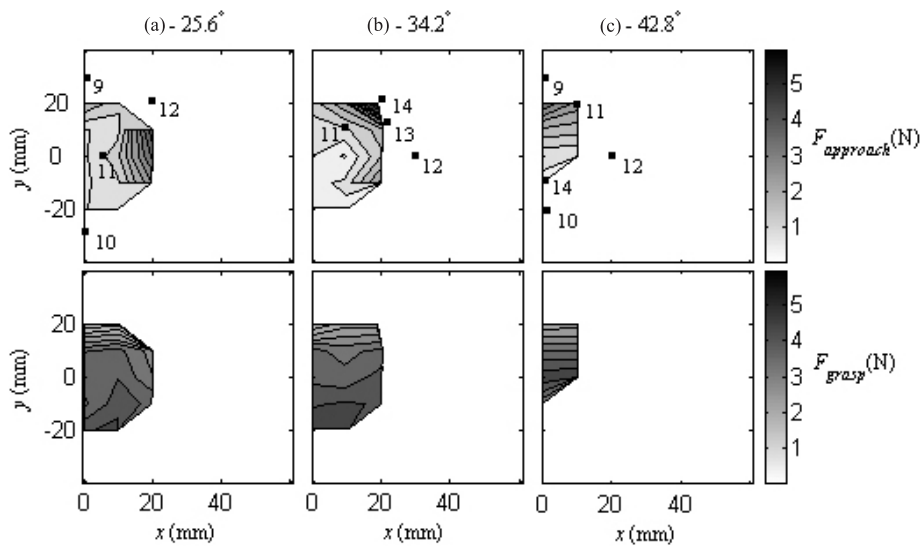


Fig. 11. Forces on the wooden block during the approach (top row) and grasp (bottom row) phases for the three object locations (columns). Labeled configurations correspond to the behavior indicated in Table 2.

of the grasping behavior at these points is given in Tables 1 and 2.

The boundary of these plots is a rough approximation of the successful grasp range (the amount of allowable positioning error resulting in a successful grasp) for the particular object and position. Note that the successful grasp range is significantly affected by the approach angle of the hand. The steeper the approach angle, the less likely enough fingers

will be in contact with the object to create a stable grasp (Figure 9).

The results show that the PVC cylinder (48 mm diameter) could be successfully grasped at positions up to 50 mm from the center in x , and +20 mm, -30 mm in y , for a total allowable positioning error of over 100% of the object size in each direction. Not surprisingly, shallow (more horizontal) hand orientations lead to larger successful grasp ranges. For the wooden

Table 1. Cylindrical Object

#	Grasp behavior
1	Four-fingered grasp
2	Three-fingered grasp
3	Two-fingered grasp
4	Object knocked from mount due to palm hitting object
5	Object twists out of grasp
6	Left fingertip sticks, then slides into place
7	Miss object completely
8	Two fingers make contact – no grasp

Table 2. Rectangular Block

#	Grasp behavior
9	Object knocked from mount due to palm hitting object
10	Two fingers make contact – no grasp
11	Four-fingered grasp
12	Object knocked from mount due to finger jamming against object
13	Left fingertip sticks, then slides into place
14	Three-fingered grasp

block (84 mm × 84 mm cross section), positioning errors of up to 20 mm from the center in *x*, and ±20 mm in *y* resulted in a successful grasp, for a total allowable positioning error of over 45% of the object size.

In general, the shape and orientation of these objects lend themselves better to a shallow or horizontal hand orientation, aligning the axis of the power grasp configuration with the major axis of the object. For this reason, additional manipulator or wrist degrees of freedom can greatly expand the amount of allowable positioning uncertainty across the manipulator workspace, particularly if the orientation of the major axis of the object can be estimated and the hand orientation controlled to match.

It can be seen from the contours that, in general, F_{approach} increases with increasing *y*. This is expected since motion forward increases the passive deflection of the joints due to contact, increasing the force. With decreasing *y*, the force goes to zero, as passive contact with the object is lost. The apparent discrepancy with this trend, seen in Figure 11(a), is simply an artifact of the sampling and contour generation.

As *x* increases, F_{approach} increases as well. This is particularly significant in the wooden block cases, where the forward-most finger first “jams” against the face of the block, eventually slipping to the side, enabling a successful grasp. As *x* in-

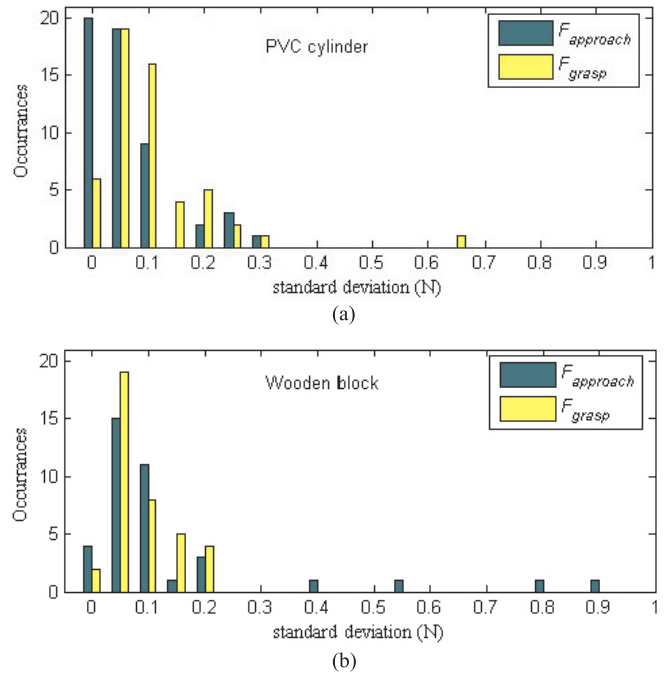


Fig. 12. Histograms of the standard deviation of the force measurements for the PVC cylinder (a) and wooden block (b).

creases, the amount of “slip” of this finger necessary for a successful grasp increases, thereby increasing the passive force. Note that, as in this example, the maximum passive force often occurs before the hand has reached the target position.

The trends in the F_{grasp} plots can be largely explained in terms of the object size relative to the fingers. For each object there is a “grasp equilibrium” position, located approximately with the object centered in the closed hand in the *y* direction, where the forces on the object would balance even without friction. Since the zero position for each object was based on the location of the front of the object and not the center, the size of the object affects the grasp equilibrium position. This position is in negative *y* for smaller objects (i.e. the object is “too close” to the base of the hand at the zero position) and positive *y* for larger objects (i.e. the object is “too far” from the base at the zero position). In general, positions far away from the equilibrium position will result in high forces.

Figure 12 shows histograms of the standard deviation of the force measurements (three samples at each configuration) for the two objects. Note that the total number of samples is different for the two objects: 38 for the wooden block and 54 for the PVC cylinder. While the values of standard deviation are typically less than the sensor resolution (0.1 N), there are a number of instances of large variation in the force measurements between trials, particularly during the approach phase for the wooden block. These instances occur at positions close to transition points between general grasp behaviors. For instance,

if the tip of a finger is very close to one of the edges of the wooden block, slight changes in hand or robot configuration can lead to drastically different behaviors (jamming against the object face *versus* gently slipping to the side).

3.2. Visual-guided Grasping Experiment

In this second experiment we seek to analyze grasper performance in conditions more representative of an unstructured environment. Target object properties are not known prior to the start of each trial and are estimated based on available basic sensing information. Note that the intent in these studies is to quantitatively evaluate the performance of the hand under uncertainty and therefore the sensing schemes are left intentionally simple.

To enable determination of three-dimensional object location using a single image from a fixed camera, spherical target objects were chosen. In this way, the object distance from the camera can be estimated by measuring the diameter of the object in the camera image combined with prior knowledge of the table location.

After each task was performed using the rough target estimate, accurate object data was then used to calculate the amount of positioning error present in each trial. This data provides a second quantification of the hand performance under uncertainty.

3.2.1. Experimental Setup

The details of the manipulator and hand are the same as those used in the previous experiment (Section 3.1.1). Again, the hand was used without a wrist, for a total of three positioning degrees of freedom, and with no control of hand orientation.

The chosen objects span a wide range in size, from approximately the minimum to the maximum size sphere that the hand can reliably grasp: a tennis ball ($r = 32$ mm), a softball ($r = 44$ mm), a small soccer ball ($r = 74$ mm), and a volleyball ($r = 105$ mm). Figure 13 shows the four objects as seen from an overhead camera. A small empty tape roll (3 cm inner diameter \times 1.5 cm height) was placed under the spheres to prevent them from rolling away on the uneven breakaway table before the grasp was initiated, but the objects were otherwise not fixed to the table and required only small forces (on the order of 1 N) to move.

A total of 12 trials for each of the four objects were conducted. The objects were placed on the workspace table in a manner such that all regions of the workspace were covered over the 12 trials, but without any additional structuring as to their specific arrangement. Only one object was placed on the table per trial. The workspace table is positioned approximately 22 cm below the origin of the robot. As in the experiment presented in Section 3.1, variations in target object



Fig. 13. Overhead image of the workspace showing the four target objects and the robot arm. Images of the workspace taken from this camera were used to find the target object location and size.

position result in different approach angles of the robot hand due to the absence of a wrist and only three positioning degrees of freedom. Objects closer to the base are approached from above, while objects far from the base are approached from the front.

3.2.2. Experimental Procedure

The target location of the robot manipulator was determined based on a single overhead image of the workspace taken from a low-resolution USB camera (640 \times 480 pixels, QuickCam Pro 3000, Logitech Inc., Fremont, CA, USA). The camera was positioned at a height of 1.63 m above the workspace, viewing a 1.26 m \times 0.94 m portion of the workspace table. The lens distortion of the camera was accounted for by calibrating using a Matlab-based camera calibration toolbox (Bouguet 2006). The calibration was achieved to a mean pixel error of 0.40, corresponding to 0.79 mm.

To register the camera to the robot workspace, a small black sphere was mounted to the end of the manipulator. The sphere was positioned within 3 cm (\pm 2 cm) of the workspace table and a total of 32 images were taken spanning the robot workspace. The two spaces were registered using a linear least-squares fit, with a root-mean-square (RMS) error of 1.98 mm. The mapping was found with a combination of the forward kinematics of the manipulator and the centroid of the sphere in the camera image. The resulting resolution of the camera is 1.97 mm/pixel of the workspace table.

During experimental trials, the target object was located in the red-green-blue (RGB) image by detecting the “colored”




Fig. 14. Image from Figure 13 after processing to find the “color” pixels.

pixels. Pixels with a ratio of the red/green and red/blue channels between 0.9 and 1.1 were interpreted as “gray”, and part of either the table or the robot. An example result of this method, performed on the image of Figure 13, is shown in Figure 14.

A bounding box was fit to the “color” blobs, and the largest taken as the target object. The largest side of this box was taken as the object diameter. This value, in conjunction with knowledge of the height of the workspace table, was used to locate the center of the object normal to the table. The centroid of the object pixels was taken as the object location in the plane of the workspace table.

Once the target position is determined based on the camera image, the robot first moves to a position 15 cm away from the target, normal to the sphere. This “approach” point ensures a consistent approach phase on the object regardless of the initial manipulator configuration, and that the hand makes contact with the target with the opening of the hand facing the object. After reaching the approach point, the robot then moves in to the object, initiating the grasp once the target position has been reached (within approximately one tenth of a degree in all joints). The arm then lifts the object upwards 15 cm, with the grasp deemed successful if the object appeared to be stable (i.e. no slippage of the object was visually observed). See Section 2.2 for further details on the actuation method of the hand. Figure 15 shows the hand in the rest configuration and grasping the smallest and largest objects.

3.2.3. Results

Figure 16 shows the location of the center of the target objects in the workspace as measured using the algorithm pre-

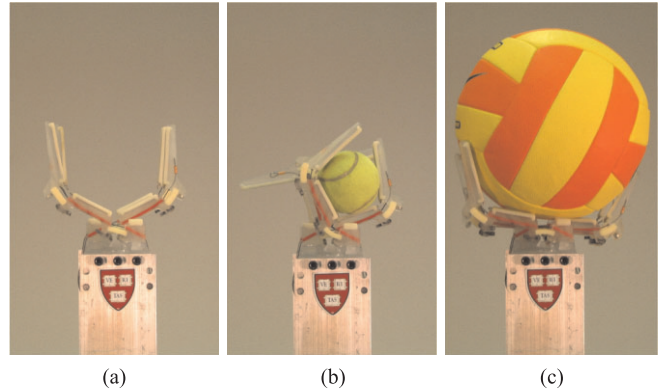


Fig. 15. Images of the SDM Hand: (a) in the approximate “rest” configuration during this task; (b) grasping the tennis ball (only three fingers make contact due to its small size); and (c) grasping the volleyball (the fingers must be pressed open by contact forces on the object). Note that due to the hand compliance, gravity tends to slightly open or close the hand from its normal “rest” configuration, depending on manipulator orientation.

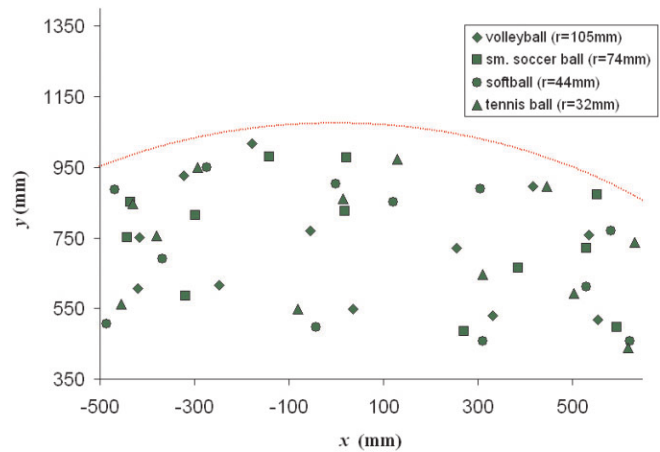


Fig. 16. Placement of the four objects on the workspace table. The arc is the approximate workspace reach limit of the robot.

sented above, with the axes corresponding to the Cartesian robot space (with zero at the robot origin). Due to their diameter, larger objects can be grasped further from the base than smaller objects, since the actual grasp target position is closer to the robot than the object center. Similarly, smaller objects can be grasped closer to the base and further to the sides since the entire object is more likely to be in the image space of the camera. The arc on the outer edge is the approximate limit of the robot workspace (i.e. arm fully extended) for the largest object (the volleyball, $r = 105$ mm).

Table 3. Object Radius Measurements

<i>Object</i>	<i>True r (mm)</i>	<i>Est. r (mm)</i>	<i>Error</i>	<i>Std. dev</i>
Volleyball	105	116.3	11.3	2.1
Sm. Soccer ball	74	74.3	0.3	2.2
Softball	44	46.1	2.1	1.6
Tennis ball	32	34.0	2.0	1.3

Table 4. Target Position Errors

<i>Object</i>	<i>Radius (mm)</i>	<i>Average position error (mm)</i>	<i>Std. dev</i>
Volleyball	105	25.2	4.5
Sm. Soccer ball	74	12.9	2.5
Softball	44	14.4	2
Tennis ball	32	15.5	2.1

The target objects were able to be successfully grasped for every trial as shown in Figure 16, despite fairly large positioning errors. Errors in each trial could have come from a number of sources. The calculation of the radius of the object (which was used in determining the distance of the center of the object from the table, as well as the target location on the perimeter of the object) was subject to large errors (Table 3). Other factors that likely contributed to errors were camera resolution (1.97 mm/pixel at the table surface) and calibration error, shadows, errors in identifying “object” pixels in the workspace image, and hysteresis in the viscoelastic joints of the hand.

In order to evaluate this error, we segmented each of the camera images by hand after the experiment in order find the true object center in the camera image and determine a more accurate target location. By doing so, we quantified the uncertainty inherent with the simplistic image processing – errors due to misinterpretation of the target object pixels due to shadows, glare, and lack of focus. Combining this true center with an accurate measure of the object radius, we generated a measure of the positioning error (the linear distance between the original target position and the “true” target position in three dimensions) for each trial. These results, summed up in Table 4, show between 7.5 and 33.5 mm positioning error for each trial, with an overall average of 17.0 mm. Note that this measure of positioning error does not account for the calibration error between the camera and robot and imperfect lens distortion correction.

4. Conclusions and Future Work

One of the main goals of this project is to simplify the amount of processing and control necessary to perform robust grasp-

ing. Indeed, we empirically demonstrated that a hand with optimized passively compliant joints and adaptive coupling can allow the grasping system to adapt to the large positioning errors that can occur in unstructured grasping tasks. Even with simplified positioning and control (three degrees of freedom arm with no wrist, a single actuator for the eight joints of the hand, and feed-forward hand control), we are able to grasp 5 cm-scale objects in the presence of positioning error of up to 100% of the object size and 10 cm-scale objects in the presence of positioning error of up to 33% of the object size (Experiment 1). We also demonstrated a greater level of autonomy and lent further weight to the argument that the SDM Hand might perform well in unstructured environments by grasping a wide range of spherical objects using real, noisy sensor data that resulted in large positioning errors (Experiment 2).

The use of the hand and the processing of any sensory information in these experiments were intentionally simple. This is not intended as a prescription for a procedure to grasp objects in an unstructured environment or the best way to analyze the available sensory information. Instead, our aim was to test the mechanical hardware performance of the hand, particularly under large uncertainties due to poor sensing.

There are a number of logical extensions to this work. The degree of autonomy demonstrated here can easily be expanded upon by utilizing the sensory information available from the joint angle and contact sensors already included in the hardware of the hand. This information, used in conjunction with an approximate model of object size and location from basic visual sensing, will make the grasping task even more robust to variations in object shape and position. Additional orientation degrees of freedom will also improve the performance by better relating hand and object geometry.

The ability of the hand to perform complicated grasping tasks can be further evaluated by operating the manipulator in teleoperation mode, allowing for more precise and dexterous positioning in order to perform more sensitive tasks. A preliminary study of the use of this mode indicates that a broad range of difficult tasks can be performed even with simple kinematics and hand control.

A major drawback of the current design is the lack of pinch-grip functionality of the hand, preventing small objects from being grasped. Essentially any object that can make contact with at least three fingers can be grasped (at least 5.5 cm in one dimension), but the hand has trouble acquiring objects such as a pen laying on a table, as the fingertips have not been adequately designed for this task. Being able to grasp small objects must be addressed in order to create general-purpose robot graspers for unstructured environments.

Ultimately, robot hands will require more complexity to accomplish dexterous manipulation tasks. When extrapolating the concepts demonstrated with the SDM Hand to a more dexterous robot end effector, a number of research directions emerge. One path is to explore the added functionality as actuators are added. How can the functionality of the hand be

increased if a second or third actuator is added, including the ability to dexterously manipulate objects within the hand? How should those successive degrees of actuation be implemented and how should the hand morphology change to take advantage of those? Underactuation can also be leveraged for low-dimensional dexterous hands if done carefully to take advantage of passive compliance and environmental affordances. However, it is clear that greater numbers of degrees of actuation will be required as desired levels of dexterity increase. In-depth investigation of the roles of compliance and underactuation in dexterous manipulation is an important area of future research, as there will always be a desire to reduce mechanical, sensing, and control complexity while retaining functionality.

Acknowledgments

The authors would like to thank Chris Johnson and Francisco Isenberg for their assistance in setting up, testing and debugging the SDM process. This work was supported in part by the Office of Naval Research grant number N00014-98-1-0669.

Appendix: Index to Multimedia Extensions

The multimedia extension page is found at <http://www.ijrr.org>

Extension	Type	Description
1	Video	Design and performance characteristics of the SDM hand

References

- Bouguet, J.-Y. (2006). http://www.vision.caltech.edu/bouguetj/calib_doc/.
- Clark, J. E., Cham, J. G., Bailey, S. A., Froehlich, E. M., Nahata, P. K., Full, R. J., and Cutkosky, M.R. (2001). Biomimetic design and fabrication of a hexapedal running robot. *Proceedings of the 2001 International Conference on Robotics and Automation*, Seoul, Korea.
- Cutkosky, M. R., Jourdain, J. M., and Wright, P. K. (1987). Skin materials for robotic fingers. *Proceedings of the 1987 IEEE International Conference on Robotics and Automation*, pp. 1649–1654.
- Cutkosky, M. R. and Kao, I. (1989). Computing and controlling the compliance of a robotic hand. *IEEE Transactions on Robotics and Automation*, **5**(2): 151–165.
- Desai, J. P. and Howe, R. D. (2001). Towards the development of a humanoid arm by minimizing interaction forces through minimum impedance control. *Proceedings of the 2001 IEEE International Conference on Robotics and Automation*, pp. 4214–4219.
- Dollar, A. M. (2006). Design principles for robust grasping in unstructured environments. *Ph.D. Dissertation*, Harvard University, Division of Engineering and Applied Sciences, October.
- Dollar, A. M. and Howe, R. D. (2005). Towards grasping in unstructured environments: Grasper compliance and configuration optimization. *Advanced Robotics*, **19**(5): 523–544.
- Dollar, A. M. and Howe, R. D. (2006). A robust compliant grasper via shape deposition manufacturing. *ASME/IEEE Transactions on Mechatronics*, **11**(2): 154–161.
- Hirose, S. and Umetani, Y. (1978). The development of soft gripper for the versatile robot hand. *Mechanism and Machine Theory*, **13**: 351–359.
- Klopsteg, P. E. et al. (1968). *Human Limbs and their Substitutes*, New York, Hafner Publishing Co.
- Laliberte, T., Birglen, L., and Gosselin, C. (2002). Underactuation in robotic grasping hands. *Machine Intelligence & Robotic Control*, **4**(3): 1–11.
- Merz, R., Prinz, F. B., Ramaswami, K., Terk, M., and Weiss, L. (1994). Shape deposition manufacturing. *Proceedings of the Solid Freeform Fabrication Symposium*, University of Texas at Austin, 8–10 August.
- Salisbury, K. J. (1980). Active stiffness control of a manipulator in Cartesian coordinates. *Proceedings of the 19th IEEE Conference on Decision and Control*, pp. 95–100.
- Shimoga, K. B. and Goldenberg, A.A. (1992). Soft materials for robotic fingers. *Proceedings of the 1992 IEEE International Conference on Robotics and Automation*, pp. 1300–1305.
- Ulrich, N. and Kumar, V. (1988). Grasping using fingers with coupled joints. *Proceedings of the 20th Biennial ASME Mechanisms Conference*, Orlando, FL.