Pre-print Accepted in Autonomous Robots March 03, 2014

Human-Inspired Force Compliant Grasping Primitives

Moslem Kazemi • Jean-Sebastien Valois • J. Andrew Bagnell • Nancy Pollard

The Robotics Institute Carnegie Mellon University Pittsburgh, PA, USA

Abstract We address the problem of grasping everyday objects that are small relative to an anthropomorphic hand, such as pens, screwdrivers, cellphones, and hammers from their natural poses on a support surface, e.g., a table top. In such conditions, state of the art grasp generation techniques fail to provide robust, achievable solutions due to either ignoring or trying to avoid contact with the support surface. In contrast, when people grasp small objects, they often make use of substantial contact with the support surface. In this paper we give results of human subjects grasping studies which show the extent and characteristics of environment contact under different task conditions. We develop a simple closed-loop hybrid grasping controller that mimics this interactive, contact-rich strategy by a position-force, pre-grasp and landing strategy for finger placement. The approach uses a compliant control of the hand during the grasp and release of objects in order to preserve safety. We conducted extensive robotic grasping experiments on a variety of small objects with similar shape and size. The results demonstrate that our approach is robust to localization uncertainties and applies to many everyday objects.

1 Introduction

Grasping constitutes an essential component in an autonomous robotic manipulation system operating in human environments. The wide variety of everyday objects and various environment settings requires grasping strategies that are robust to variations in shape, size, and pose, as well as uncertainties in perception and robot kinematics. Although advances in perception provide accurate object pose estimation, calibration and kinematic factors affect the accuracy at which an end-effector can be controlled for grasping. Because of such uncertainties



Fig. 1: Andy (our dual-arm robotic platform provided by ARM-S DARPA program) performs a compliant finger landing followed by a compliant grasp to grab a wire-cutter from the table.

contact information such as tactile and force/torque feedback is necessary to achieve robust grasps.

Much research in robotic manipulation is focused on geometric grasp generation and planning. Researchers study finger and object interactions and develop grasp quality metrics based on form/force closure [3]. Despite promising progress in geometric grasp planning and the efforts devoted to analyzing grasp properties, challenges still remain in real-world manipulation tasks because the existing techniques fail to consider opportunities presented by contacts with the environment. In our view, supported by our empirical studies, two factors play a crucial role in achieving more robust grasps: dealing with positioning uncertainties [10], and using compliant motions [23] to handle contacts between the robot and the environment.

Our work (Fig.1) is motivated by the task of grasping everyday objects that are small relative to an anthropomorphic hand, including a pen, screwdriver, cellphone, and hammer from their natural poses on a support surface, e.g., a table top (see Fig.2). Note that existing techniques for grasping larger objects generally tend to perform well when reliable grasp points can be found away from the support surfaces. However, that is rarely the case for small objects. In addition, because existing grasp planning tools (such as OpenRAVE [7] and GraspIt![24]) rely on precise finger-to-object contact points while avoiding the surrounding environment they are impractical for grasping small objects. Figure 3 illustrates some of the failure modes of a geometric grasp planner.

In our experience, geometric techniques are only effective when precise calibration is performed to position the robot relative to the object. However, limitations in sensing and control uncertainty make grasps that rely on precise calibration impractical. Thus, it is important to devise grasping strategies that are robust to these errors.

Most approaches to grasping attempt to model uncertainty. In contrast, we believe that the effect of uncertainty can be ignored through compliant interactions between the robot and

The Robotics Institute, Carnegie Mellon University Pittsburgh, PA 15213 USA



Fig. 2: The anthropomorphic Barrett hand grasps small objects: a cellphone, hammer, pen, and screwdriver from a table top.



Fig. 3: Failure modes of a geometric grasp planner: (a) no contact point found since the hammer is sunken into the table due to localization error, (b) computed contact points are out of collision but too close to the table, (c) the planner does not allow for the contact with the table and therefore enveloping grasps are not considered resulting in an unstable grasp, (d) releasing the object fails since the planner does not account for collision between the fingers and support surface.



Fig. 4: Compliant landing of fingers followed by compliant grasping of a hammer

the environment. Because we allow contact with the environment, compliant motions are crucial to ensure safety of the robot and successful execution of the task. Our approach is motivated by a series of human subject studies which demonstrate similar behavior in human grasping of small objects (see Section 3). The use of compliance for grasping small objects is the main contribution of this paper.

We present three simple closed-loop sensory driven manipulation primitives for robust grasping (and releasing) of small objects from support surfaces: (1) *compliant finger placement* for bringing all fingers safely in contact with the support surface, (2) *compliant object grasping* for maintaining the contact between the fingertips and the support surface during the finger closure, and (3) *compliant object release*. An example of compliant landing and grasping primitives is shown in Fig.4. A shorter version of this work has been presented in [20].

2 Related Work

The literature on robotic grasping is vast. Here we only refer to some of the related major works in the area, but we also encourage the reader to consult other reviews, e.g., [3].

First, we provide an overview of geometric grasp planning techniques where usually a grasping task is formulated as an optimization problem and sample hand pre-shapes and postures are computed based on the object geometry. Next, we talk about closed-loop control grasp schemes where sensory feedback is used to adjust the fingers or hand pose to reach a desired posture. We discuss approaches which focus directly on reducing localization uncertainty through active sensing. We also refer to works which leverage the mechanics of the task to reduce localization uncertainty. Similarly, in our work we take advantage of the contact with support surface to reduce the effect of uncertainty. Then we briefly talk about force feedback control strategies in the context of compliant manipulation. Our work incorporates compliant motion control into the finger landing and object grasping behaviors to ensure safety. Finally, we finish our review by summarizing some of the recent works on human grasping behaviors and humaninspired robotic grasping schemes and note that previous studies neglect the importance of contact with support surfaces in achieving robust grasps under uncertainty. This intuition motivated our formal human subject study which follows in Section 3.

2.1 Geometric grasp planning

Geometric grasp planning techniques rely mainly on the geometry of the object to generate grasps which are stable based on a given metric, e.g., force/form closure [3]. Ciocarlie and Allen [6] introduced the notion of of *eigengrasp* which reduces the dimensionality of the grasp control space via identifying linear subspaces which effectively represent a wide range of hand pre-shapes. The eigengrasps are generated for a known object based on its geometry and an optimization is performed to compute a form closure grasp as a combination of eigengrasps. Berenson and Sirinivasa [2] propose an optimization framework based on a cost function which encapsulates the reachablity of the object and force closure properties based on sample contact points on the surface of the object.

The requirement to have access to the object model limits the application of the above techniques which rely directly on the object geometry. Hence, people have looked at grasp planning in unstructured environments. For example, Saxena *et. al* [29] propose a supervised learning approach which computes grasp points directly based on 2-D images of the object and learned from previous experiences.

Nonetheless, the aforementioned techniques along with other geometric grasping strategies either neglect or try to avoid contact with the support surface during grasp execution, and hence

often fail to successfully execute a planned grasp, particularly for small objects (see Fig. 3). In such approaches the fingers have to either maintain contact with pre-computed points on the object surface or reach a certain pre-shape to ensure a robust grasp. Such requirements limit the application of geometric grasp planning techniques in dealing with small objects which are difficult to precisely localized and also may not be reachable from every side due to support surface. In contrast, our approach does not heavily rely on the object geometry. In most cases the principal axis and center of gravity is enough to define a grasp. Moreover, our compliant grasping primitive ensure a successful grasp of the object by maintaining contact with the support surface as the finger close to reach the object. We argue that reliable grasp of small object must consider contacts with support surfaces. This intuition is supported by our 18 human grasping studies (Section 3) and our approach leverages this insight.

2.2 Closed-loop control grasping

In contrast with geometric grasp planning techniques, some formulated the grasping task as a control problem where the hand along with the fingers are actively controlled to achieve certain contact points or a stable grasp evaluated in real-time based on a given metric. Such approaches rely on the feedback from contact forces/torques, tactiles, and/or vision to close the control loop. Usually the grasp controller is initialized based on a rough estimation of the object pose.

Platt *et al.* [27] devised null space grasp controllers to achieve *frictionless equilibrium* grasp configurations by displacing the fingers over the object surface and aiming to regulate contact force/moments error residuals to zero. Tahara *et. al* [31] also propose a grasping method to realize dynamic force/torque equilibrium. The existing closed-loop control grasping strategies rely on sensory feedback (e.g., tactile) retrieved from the object surface. This implies that the object should be large enough so there will be enough data for the control scheme to converge. Moreover, the object should be reachable from different direction for the hand to assume a desired posture imposed by the underlying control scheme. Our propose primitives do not rely on object surface information but take advantage of the contact with the support surface instead to successfully cage small objects located on a support surface.

2.3 Grasping under uncertainty

We believe uncertainty is the key challenge for grasping small objects because of limitations in perception and calibration. Many techniques attempt to handle uncertainty by explicitly modeling the contact between the fingers and the object (e.g. [26]). In practice, for small objects such precise modeling is very difficult. Active sensing has been proposed to reduce uncertainties in object pose estimation. For example, Hsiao *et al.* [14] and Javdani *et al.* [15] use pre-grasp interactions to estimate the object pose without disturbing it. In a recent work by Hebert et al. ([12]), an Unscented Kalman Filter estimates the position of scene objects, including the robotic arm and the grabbed object, using inputs from various sensors such as RGB-D, image edges, and fiducials. The effectiveness of such techniques are very limited when dealing with small objects because of localization errors due to coarse sensor data retrieved from such objects.

Some have taken advantage of the mechanics of the task (e.g., as in [4, 10]) via deliberate interaction with the environment to reduce uncertainty in manipulation tasks. For example, in [8] a push-grasping motion (parallel to the support surface) is used to align and bring the object inside the capture envelope of the hand. Although such techniques are successful in grasping relatively large objects (e.g., a cup) their effectiveness in grasping small objects (e.g., a pen) from supporting surfaces is questionable. However, we note that our approach also leverages the mechanics of the task via deliberate contact with the support surface and proper response to contact forces to place the fingers on the support surface and maintain contact with the surface to achieve successful grasp of the object.

2.4 Compliant manipulation

Our work is also closely related to the notion of *active* interaction control in the context of robotic manipulation, which has been extensively studied over the past three decades under the topic of force control: *impedance* or *admittance* control (e.g., [13]), *compliance* control (e.g., [23]), *damping* control (e.g., [33]), and *hybrid force/position* control (e.g., [21]), to name a few. Please see [34] for a brief review.

The current work does not claim to introduce a new force control scheme, but incorporates existing force compliant control schemes into dexterous grasping behaviors to achieve robustness with respect to uncertainties and ensure safety. Specifically, our proposed multi-finger landing and grasping primitives employ (hybrid and explicit) compliant control strategies as proposed in [23] to control the velocity of the hand during landing and grasping behaviors based on the contact forces. The desired hand velocities are then translated to joint torques using a computed torque control method [25]. We note that the above velocity-based operational space control scheme (i.e., force \rightarrow velocity \rightarrow torque) is only one of the many ways through which one can introduce the desired compliant behavior (see [25] for a comparison). For example, one may devise a sophisticated hybrid force/position operational space formulation [21] which directly maps the desired forces in the operational space to the joint torques and offers high bandwidth control with nice decoupling properties and better performance. However, we chose to go with the velocity-based operational

2

3

4

5

7 8

9

11

12

13

14

15

16

17

18 19

20

21

22

23

24

25

26

27

28 29

30

31

32

33

34

35

36 37

38

39

40

41

42

43

44

45

46

47

48

49

50

51

52

53

54

55

56

57

58

59

60 61 4

space formulation due to its reasonable performance (please refer to [25]) and ease of its implementation on our system.

In [30] active body-environment contacts are used to generate coordinated hand and arm motions. Although they also emphasize on the importance of body-environment contacts however they demonstrate their approach on glass placing and drawer opening tasks which does not fully point out the importance of dexterous contacts with the support surfaces during finger landing and grasping. Moreover, it is not clear how the proposed motion control can be generalized to other grasping tasks. In concurrent related work, Kalakrishnan et al. [17] recently presented a learning strategy to acquire manipulation and grasping skills where an initial position control policy for the manipulation task is initialized through kinesthetic demonstration. The learned policy is then augmented with a force/torque profile that is controlled in combination with position trajectories using a force compliant strategy in a closedloop scheme that is similar to ours. The force/torque profile is learned through demonstration by optimizing a cost function which measures the task success. In contrast to their technique, our approach requires no learning and implements a simple closed-loop force feedback controller that generates the compliant motions necessary to maintain the proper contact between the fingertips and the supporting surface, and is operational across a broad range of conditions.

Finally, we should also mention the works which take advantage of (passive or active) compliance of fingers to reduce the effect of localization uncertainties. For example, Dollar and Howe [9] demonstrated that a combination of passive joint compliance and adaptive transmissions enables the hand to conform to the object shape and achieve robust grasps despite large localization uncertainties.

2.5 Human-inspired grasping

Human grasping behavior has been a extensively studied by neuroscientist, and yet there are many questions left to be addressed [5]. Studies on human grasping behaviors have inspired robotic grasping strategies. Johansson and Flanagan [16] decompose grasping tasks into a series of consecutive actions (i.e., reach, load, lift, hold, replace, and unload) delimited by contact events that represent sub-goals of the overall grasping task. Inspired by this study, Romano et al. [28] proposed a series of grasping controllers which are executed sequentially. The proposed controller mimic the human grasping actions and are triggered based on tactile signals incorporated into a robotic hand following the observations by Johansson and Flanagan from human grasping behaviors. In the above work less attention has been paid to the role of the contact between fingertips and the support surface in achieving successful grasp of small objects. In contrast, we argue that such contacts are critical to establish a proper relative position of the fingers and the object throughout the grasp. To support this belief we performed a series of formal human grasping studies. The details of this study are presented in Section 3. In summary we observed that the test subjects consistently favored contact with the support surfaces under normal conditions (as they do in everyday grasping tasks). They showed increase in contact with the support surface when they were less certain about the object pose. Please refer to Section 3 for more details and the scenarios. The results of our human grasping study inspired the compliant grasping primitives in this paper.

We propose a compliant grasping strategy which performs pinching as well as enveloping grasps [32]. In contrast to pinching grasps, where the object is restrained by the fingertips at certain contact points, enveloping grasps are formed by wrapping the fingers and the palm around the object. Theoretical analysis (e.g. as in [32] and [11]) can be used for pulling objects from a surface into an enveloping grasp when identifying grasps for new objects. The choice of the grasp is affected by various parameters including the task and the size, shape and weight of the object. Enveloping grasps are shown to be superior in terms of restraining objects as expressed in [3]. Although preferable, they often are more challenging to perform in scenarios where the object is lying down on a supporting surface, e.g., grabbing a hammer lying on a table (Fig. 4). In fact, in such scenarios, the fingers need to come in full contact with the support surface and then slip underneath the object while the hand is pushed downwards to maintain the proximity to the support surface. Proper control of the hand to achieve such motion without breaking or stalling the fingers is the motivation for our work.

3 Human Grasping Studies

This work attempts to learn from human hand and fingers motion behaviors for grasping small objects from a support surface, and to transfer those to robotics manipulation techniques. The motivation for mimicking human grasping method is to take advantage, to the extend allowed by the robotics hardware, of effective grasping strategies for manipulating objects that have been difficult up until now for robots, such as thin objects resting on flat surfaces. In a recent study, Johansson and Flanagan [16] decompose grasping tasks into a series of consecutive actions (i.e., reach, load, lift, hold, replace, and unload) delimited by contact events that represent sub-goals of the overall grasping task. For example, the goal of the initial reach phase is marked by the fingertips contact with the object and the goal of the subsequent load phase is marked by breaking the contact between the object and the support surface. Although the above study show the role of contact in triggering actions, less attention has been paid to the role of the contact between fingertips and the support surface in achieving successful grasp of small objects. In contrast, our belief is that such contacts are critical to establish a proper relative position of the fingers and the object throughout the grasp. We

performed a series of human subject studies to validate our belief. First, we explain the experiments setups and scenarios under which the experiments were performed. Then we report and analyze the results obtained from our human subject study and conclude the section by explaining how the results are applied to robotics grasping behaviors.

3.1 Experimental setup

We set out to register the contacts trajectories (if any) between a subject fingertips and the support surface while grasping objects. To that end, we developed a custom iPad application to detect and log the location of each fingertip contact points (up to 10 separate tracks were allowed at 60Hz sampling rate) anywhere on the iPad surface. The fingertip contacts are time-indexed and stored as trajectories and then aggregated as contact density histograms. We obtained the approval for our study from our internal Institutional Review Board (IRB) through a formal application process. A group of ten right-handed volunteers were selected at random in our institution, and we ensured that they had no prior knowledge of our research. During the experiments, we provided them with short and straightforward instructions regarding the intent of the experiment. For example, we specified whether they should observe the object during the grasps, if the surface had special properties (e.g. hot to the touch), and finally if the purpose of the grasp was to pick-and-place or rather work with the object. We avoided providing additional comments to preserve impartiality to the way objects were grasped. The subjects were asked to sit normally behind a table with the iPad in front of them. The objects that were tested include a pen, a medium-sized screwdriver, a flip cellphone, and a regular hammer (see Fig.2). In addition, the subjects were asked to wear contact sensitive gloves ¹ since the iPad surface could not register finger nail contacts.

Figure 5 shows the experimental setup along with a sample finger trajectory recorded during a typical grasping task. Figure 5(b) shows a 2-D histogram of fingertip contacts trajectories, together with a representation of where the object was located on the iPad. The center of the figure shows the highest density of contacts due to subjects dragging their fingers towards the center of the object and coming to a pinch underneath the object.

3.2 Study scenarios

We conducted the grasping experiments under four different scenarios:

- 1. *Normal* scenario: the subjects were asked to grasp the object as they normally do everyday for a *pick-and-place*
- ¹ Agloves, http://www.agloves.com/



(a) Left: fingertip trajectories, right: examples of robot and human grasps



Fig. 5: Experiments studying human grasping strategies

with no restrictions involved. This scenario defines the baseline experiment in our comparison with other scenarios.

- 2. *Blind* scenario: where the subjects were asked to localize the object first and then perform the grasp after first looking away. Our objective in performing this experiment was to evaluate how uncertainties in the object location would affect the amount of contacts with the support surface. This resembles the *sense*, *plan*, then *act* approach used in most robotic systems. In grasping scenarios this would be implemented as 1) localization the object, 2) plan to the target, and 3) open-loop (i.e without visual feedback) grasp execution.
- 3. *Hot surface* scenario: the subjects were asked to grasp the objects while pretending that the iPad surface was extremely hot, so it should be avoided as much as possible while grasping the object. The intend for this test was to determine if contacts with the support surface could be completely avoided, and also qualitatively observe slippages or other apparent difficulties compared to other scenarios.
- 4. *With intention* scenario: the subjects were asked to grasp the object for the purpose of performing a relevant task afterward, e.g., writing with the pen, using the screwdriver to tighten a screw, answering a call with the cellphone, and hammering a nail. Our objective is to observe how grasping with intention of manipulating the object afterward (as

opposed to just picking and placing it in the *normal* scenario above) affected the surface contact pattern.

Fig. 6: Experiments studying human grasping strategies

For each object we asked each subject to perform the grasp action twice. In the end we had 10 subjects, 2 grasps per object, for a total of 20 sampled grasps for each object. The subjects fingertips contact with the iPad surface were recorded for each grasping action as described above.

3.3 Results Analysis

Figure 6 shows the average contact time between the fingertips and the support surface over 20 grasps. The results show that subjects consistently favor making contact with the support surface for grasping all objects except when they were explicitly told to avoid the surface, i.e., the *hot surface* tests. We observed that in the *hot surface* scenario, the participants performed each grasp with more hesitation, and occasionally dropped the pen (2 clear instances).

Figure 7 shows the percentage difference in average contact time between the *normal* (used as the baseline) and other test scenarios (i.e., *blind*, *hot surface*, *with intention*). In the *blind* scenario, an increase in the contact duration is observed for all objects due the initial finger positions with respect to the object. The difference is less pronounced for larger objects (e.g., cellphone versus hammer) since bigger objects can be localized with less contact. We also noticed that the grasps *with intention* had relatively shorter grasping time when compared to the *normal* tests, which seems to indicate that the subjects payed more attention in planning the finger placement during those grasps. This was true for all objects, except for the hammer for which similar grasps were used in both scenarios. Finally, the data from the *hot surface* tests shows the relative ease for successfully grasping object while avoiding surface

With intention Hammer Hot surface Blind Cellphone Screwdriver Pen -150 -100 -50 50 100 150 200 0 Contact time percentage difference compared to normal scenario

Fig. 7: We compared the duration of finger contacts with the iPad surface in each scenario and the normal scenario (Fig.6) as our baseline. This plot shows the percentage difference from the normal scenario for each object-scenario experiment.

contacts, indicating a *preference* by the test subjects for such contacts in most cases.

Figure 8 shows the 2-D contact histograms for each object/scenario obtained by superimposing the corresponding contact histograms of all subjects. The contact histograms graphically illustrates the distribution of contacts in the vicinity of an object place at the center of the figure and aligned with the vertical axis. The surfaces colored with shades closest to red denote areas with higher number of contacts. As shown in Fig. 8, the *blind* scenario histogram cover the widest area around the object, which demonstrates the greatest need to use the contact surface to reduce uncertainty before grasping the object. In contrast, the hot surface experiment histograms showed the lowest amount of surface contacts due the test subjects' ability to directly capture the objects. Finally, we want to point out slight differences in the contact histogram between the normal and with intention scenarios, due to the differences in chosen finger pre-shapes for both cases. For example, Fig. 9 shows the hand pre-shapes in two different scenarios: pick and place task (left) and with intention of writing (right). Although the objective of this experiment is not to study the relation between the finger pre-shapes and a given task, the results clearly show the subjects' preference in maintaining contact with the support surface when grasping objects with the intention of manipulating them afterwards.

3.4 Final remarks

Unless explicitly forbidden, the test subjects favored to contact the support surface when grasping objects. The advantage of such approach was to reduce finger position uncertainty with respect to the objects, and also to completely cage them early in the grasp. As a result of those strategies, the grasps observed are more robust than pinch grasps. These observa-



1

2 3

4

5

6

7 8

9 10

11

12 13

14

15 16

17

18

19

20 21 22

23

24

25

26

27

28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43

44

45

46

47

48

49

50

51

52

53

54

55

56

57

58

59



Fig. 8: 2-D contact histograms obtained by superimposing the fingertip contacts of all subjects for each object-scenario. We did our best to precisely pose the object at pre-defined locations (shown as overlay bounding boxes) on the iPad surface.

tions inspired the development of the grasping primitives for dexterous grasping of small objects. The crux of our approach is to establish contact with the support surface and maintain it while moving the fingers toward the object to cage the object. To ensure safety and accuracy we employ simple sensorydriven force compliant strategies similar to humans avoiding too much finger pressure during grasping tasks.

4 Hardware and Control Requirements

In this section we explain the necessary sensory feedback and control required for the implementation of the proposed prim-

itives. We believe that such capabilities are essential for any manipulation systems, and hence should be readily available on existing systems.

4.1 Hardware Requirements

The compliant motion primitives presented in this work use two sensing modalities: (1) a 3-axis force/torque sensor positioned at the wrist between the manipulator arm and the hand, and (2) strain gauges between the proximal and distal finger segments. Moreover, the proposed primitives assume knowledge of fingertip positions from forward kinematics. Force feed-



Fig. 9: Hand pre-shapes chosen by a subject to grasp a marker under two scenarios: pick and place task (left), and with intention of writing (right). In our human subject studies we observed that in both scenarios the subjects maintain contact with the support surface. However, we note differences in the contact patterns as shown in Fig.8.

back from the wrist closes the loop for performing compliant motions as described below. Finger strain gauges are solely used to detect when fingers are in contact with the support surface.

Ideally, either of the above sensing modalities can provide the necessary feedback to implement the proposed primitives. However, we found the finger strain gauges hard to calibrate and very noisy, so we only used them as a binary sensor to detect individual finger contacts. At the same time, the 3-axis force/torque was noisy enough that looking at the direction of the force vector was impractical to determine which finger had made contact, especially in the presence of kinematic modeling errors. Other sensory hardware can be used as long as the above feedback is provided. For example, tactile sensors at fingertips can be used to detect the finger contacts.

4.2 Velocity-based operational space control of hand

We employ a velocity-based operational space formulation to generate compliant motion of the hand in response to forces seen at the wrist. This formulation is fairly straightforwrd to implement and demonstrates an overall good performance compared to other variations of operational space control [25]. Desired joint velocities to track a given hand velocity $\dot{\mathbf{x}}_d$, determined based on the task primitive, are calculated using Liegeois' resolved motion rate control approach as: [22]

$$\dot{\mathbf{q}}_d = \mathbf{J}^+ \dot{\mathbf{x}}_d + \lambda (\mathbf{I} - \mathbf{J}^+ \mathbf{J}) \nabla \mathbf{H}(\mathbf{q}), \tag{1}$$

where \mathbf{J} is the robot Jacobian with its pseudo-inverse denoted J^+ , λ is a gain value, and H(q) is a null-space cost/utility function. Different criteria can be used to define H(q) depending on the objective, e.g., avoiding joint limits or kinematic singularities.

The desired motors torque command is calculated using the computed torque control method with a PD controller [25] to track the desired joint velocities in $\dot{\mathbf{q}}_d$,

• `

$$= \mathbf{M}(\mathbf{q}_d)\ddot{\mathbf{q}}_d + \mathbf{C}(\mathbf{q}_d, \dot{\mathbf{q}}_d) + \mathbf{g}(\mathbf{q}_d) + \mathbf{K}_d(\dot{\mathbf{q}}_d - \dot{\mathbf{q}}) + \mathbf{K}_p(\mathbf{q}_d - \mathbf{q})$$
(2)

where M is the inertia matrix, C is the Coriolis/centrifugal vector, \mathbf{g} is the gravity vector, \mathbf{K}_d and \mathbf{K}_p are gain matrices, and τ is the joints torque vector. \mathbf{q}_d and $\ddot{\mathbf{q}}_d$ are obtained through time integration and differentiation of $\dot{\mathbf{q}}_d$, respectively, i.e.

$$\mathbf{q}_d = \mathbf{q}_d(t - \Delta t) + \dot{\mathbf{q}}_d \Delta t \tag{3}$$

and

$$\ddot{\mathbf{q}}_d = \frac{\dot{\mathbf{q}}_d(t) - \dot{\mathbf{q}}_d(t - \Delta t)}{\Delta t} \tag{4}$$

In the following sections, we describe how the desired hand velocity $\dot{\mathbf{x}}_d$ is computed to achieve a compliant motion behavior in response to contact forces.

4.3 Coordinated position control of fingers

The fingers are controlled along their pre-defined trajectories using a position-based method. The trajectory is defined as a sequence of waypoints based on the given task primitives. This controller is used to coordinate the position of fingers along their desired trajectories during the grasping or releasing of objects.

In our implementation we used the Open WAM Driver (OWD) maintained by the Personal Robotics Lab at CMU ² which provides the computed torque method for velocitybased operational space control along with the finger position control for a WAM robotic arm.

5 Force Compliant Grasping Primitives

A grasping task can be decomposed into four distinct, sequentially executed steps as illustrated in Fig. 10. (1) compliant landing: fingers are placed in a pre-defined grasp pre-shape and the hand is maneuvered downward until all fingers fully rest on the support surface, (2) compliant grasping: force feedback is used to maintain a desired contact force at the fingertips while the fingers joints are synchronously closed to capture the object, (3) lift and transportation: the object is lifted away from the surface and carried to the destination, (4) compliant release: the object is gently deposited on the support surface using a method similar to step 2. A similar sequence of grasping controls has been suggested in [28] inspired by neuroscience studies, but it lacks the finger landing step and the compliant primitives proposed in this paper.

² Open WAM Driver: http://personalrobotics.ri.cmu. edu/pr-ros-pkg/owd/html/index.html





Fig. 10: A grasping task decomposed into four distinct steps



Fig. 11: An example of compliant finger landing: the circles indicate the finger which just made contact with the support surface as the hand is servo controlled around pre-defined control axes. The appropriate control axis is selected based on the finger-surface contact information. To ensure safety of the fingers while maintaining contact with the support surface the hand is servo controlled in compliance with contact forces.

5.1 Force compliant fingers landing

Landing begins by positioning the palm at a safe hovering distance above the surface, in the direction of the plane's normal. To pose the hand we assume that the location of the support surface and its normal are roughly known from perception. However, please note that our control strategy does not explicitly rely on this information.

We achieve safe finger landing by controlling the hand towards the support surface using a compliant controller. We use a velocity-based controller (Eq. 2) to generate the compliant motions. This controller brings all fingers to contact with the support surface, while preserving them from damage. Landing is achieved as soon as all fingertips establish contact with the surface. Because the fingers do not necessarily contact the surface simultaneously, we continuously update the servo control point and axis to correct the hand's motion. This process terminates when landing is achieved (i.e., all fingers are in contact with the support surface). Figure 11 illustrates a finger landing sequence. Our landing strategy can be adapted to support various grasp pre-shapes.

Our approach relies on feedback from the finger strain gauges to determine whether a finger is in touch with the support surface. It sets contact flags C_i for fingers i = 1, 2, 3. These contact flags are updated in real-time during landing, and are used to compute the axis around which the hand needs to be rotated. The control axes are predefined based on the selection of control points at the fingertips and their location with respect to the hand's (or the end-effector's) frame. We choose the control points at the fingertips, i.e., CP_{100} , CP_{010} or CP_{001} ,



Fig. 12: Control points and axes specified based on the the fingertips location and contact information. In the compliant landing primitive, the contact between the fingertips and the support surface is sensed through the finger strain gauges. The contact information is used to select the appropriate control axis to servo the hand and land the fingers which are not in contact with the surface.

as shown in Fig.12. For a single finger contact, the control axis passes through the fingertip and is defined parallel with the line passing through the other two fingertips. When two fingers contact the surface, the mid-point between fingertips defines the control point, i.e., CP_{110} , CP_{101} or CP_{011} , and the control axis is specified by the two control points at the fingertips. The landing primitives continuously observe the contact status of the fingers and calculate the appropriate control axis around which an angular velocity ${}^{c}\omega_{d}$ (with a constant magnitude) is applied to land the non-touching fingers (see Algorithm 1). Figure 13 shows the landing control diagram.

In our experiments, uncertainties in localizing the surface normal and/or modeling errors in the robot kinematics resulted in an offset in the computation of the control axes. As a result, rotation is performed around a slightly different control axes, causing the contacting fingers to either lose contact with the surface or apply significant force onto the surface. The former leads to failure, and the latter endangers the fingers.

We rely on force compliance to avoid these risks. We move the hand in compliance with the forces exerted onto the fingers, which ensures a proper contact between the fingertips and the surface. The compliant motion introduces a linear ve-



Fig. 13: Compliant landing control diagram (the compliant grasping control loop follows a similar design but with the angular velocity set to zero, i.e., ${}^{c}\omega_{d}$)

locity only at the control point and along the palm's normal $\vec{\mathbf{n}}_h$ (shown in Fig. 12), in response to forces measured at the wrist.

The desired force to servo the hand is calculated as

$$f_d = f(t) - f_r - f_t, \tag{5}$$

where f(t) is the current force seen at the wrist along the palm normal, f_r denotes the reference force recorded *a priori* before the fingers touch the surface. The parameter f_t is an attractive force value used to ensure downward motion of the hand (along its palm normal) when none of the fingers is in contact with the support surface, i.e. when $f \approx f_r$. The attractive force determines the amount of force that the hand will exert to the surface before it starts to respond and comply with contact force (in our implementation we used an attractive force of 2N).

Finally, the linear velocity at the control point is given as

$${}^{c}\boldsymbol{v}_{d} = k_{f}f_{d}\,\overrightarrow{\mathbf{n}}_{h} \tag{6}$$

where k_f is a positive gain. In practice we observed that passing f_d through a dead-band filter helps to decrease the oscillation effect when f_d approaches zero, i.e.

$$f_d := \begin{cases} 0 & \text{if } ||f_d|| < f_{db}, \\ f_d - f_{db} & \text{if } f_d > f_{db}, \\ f_d + f_{db} & \text{if } f_d < -f_{db}. \end{cases}$$
(7)

where f_{db} is the width of the dead-band filter (in our implementation we used a dead-band width of 0.5N).

The velocity screw at the control point ${}^{c}\dot{\mathbf{x}}_{d}$ is composed of a linear force compliant component ${}^{c}\boldsymbol{v}_{d}$ and an angular velocity ${}^{c}\omega_{d}$ around the corresponding control axis as explained above, i.e.,

$${}^{c}\dot{\mathbf{x}}_{d} = \begin{pmatrix} {}^{c}\boldsymbol{\upsilon}_{d} \\ {}^{c}\boldsymbol{\omega}_{d} \end{pmatrix}$$
(8)

Algorithm 1: Force Compliant Finger Landing

Agorithm 1.1 ofce Compnant I inger Landing					
1	begin				
2	Record the current finger strains, S_{ir} ;				
3	Set the finger contact flags to zero, $C_i \leftarrow 0$;				
4	repeat				
5	Get the current finger strains, S_i ;				
6	Get the current force along palm's normal, f ;				
7	Update the fingers contact flags, C_i based on strains;				
8	Determine the control point and axis, and the control point				
	velocity screw $^{c}\dot{\mathbf{x}}_{d}$;				
9	Compute the desired hand velocity screw $\dot{\mathbf{x}}_d$; Apply the				
	hand velocity $\dot{\mathbf{x}}_d$ to joints using Eq.2;				
10	until $C_i == 1$ for $i = 1, 2, 3;$				
11 end					

The control point velocity screw ${}^{c}\dot{\mathbf{x}}_{d}$ is transformed to the corresponding hand velocity screw $\dot{\mathbf{x}}_d$ given kinematics of the fingers and the current position of control point. Finally, the hand velocity screw is then applied using the computed torque method as (2). Algorithm 1 summarizes the proposed landing strategy.

5.2 Force Compliant Grasping

Our compliant grasping strategy begins with the assumption that all fingertips are in contact with the support surface. This strategy consists of two closed-loop controllers that run independently and in parallel (see Algorithm 2). The first servos the hand using a velocity-based operational space controller. The second is a position-based controller that moves the fingers along their pre-defined trajectories (see Section 4.2).

The velocity-based operational space controller is similar to (2), which is described in the previous section. This controller generates the compliant motion of the hand in response to forces exerted by the support surface to the fingertips. The goal of this controller is to maintain contact between the fin-



Fig. 14: Force compliant grasping primitive: the hand is controlled in compliance with contact forces exerted from the support surface to the fingertips while the fingers follow their pre-defined trajectory to reach the object.

Algorithm 2: Force Compliant Grasping						
1	begin					
2	Record the current force along palm's normal, f_r ;					
3	Start fingers trajectory tracking by sending the first					
	waypoints to fingers controller;					
4	repeat					
5	Get the current force along palm's normal, <i>f</i> ;					
6	Compute the desired hand velocity $\dot{\mathbf{x}}_d$ using Eq.9;					
7	Apply the hand velocity $\dot{\mathbf{x}}_d$ to joints using Eq.2;					
8	if all fingers reached their waypoints then					
9	Send the next fingers waypoint positions;					
10	end					
11	until all fingers reached their desired final positions;					
12	12 end					

gers tips and the support surface, while closing the fingers towards the desired object-caging configuration (see Fig. 14)

The compliant motion applied to the hand is composed of only a linear motion calculated, similar to (6), as

$$\boldsymbol{v}_d = k_f f_d \, \vec{\mathbf{n}}_h \tag{9}$$

where f_d is defined in (5). The hand velocity screw is then given as $\dot{\mathbf{x}}_d = [\boldsymbol{v}_d^T, \mathbf{0}^T]^T$ which is applied to the joints using the computed torque in (2).

While the fingers move along their predefined trajectories, our compliant velocity controller responds to forces that are due to the fingers' contact with the support surface. These forces are measured at the wrist. The fingers' positions are coordinated to ensure proper caging of the object, without missing it. To coordinate the fingers' positions, as soon as all fingers reach their waypoints, new waypoints are provided to the position-based controller (our second controller). Due to uneven contacts, some fingers may be lagging behind. This coordination strategy is key for successful grasping of small objects.

Our compliant grasping strategy can be used for performing both pinching and enveloping grasps. In a pinching grasp, the object is restrained by the fingertips only. We achieve this behavior by stopping the fingers as soon as contact with the object is detected via strain feedback. The enveloping grasp continues after pinching the object. It applies additional torque to the fingers while pushing hand downward, which in practice encourages the object to slip towards the palm. At the same time, we close the fingers to fully capture the object.



Fig. 15: Force compliant release/placement primitive

5.3 Force compliant object release

The ability to accurately place an object and release it from grasp is as important as the ability to grasp and lift it, particularly for tasks such as stacking objects, assembly, or exchanging objects between hands in bi-manual manipulation. We devise a compliant strategy for gentle release and placement of a grasped object on a support surface. The proposed approach effectively avoids abrupt release of the object from the grasp and ensures gentle placement on the support surface from both precision and enveloping grasps. Release from enveloping grasps is more challenging due to the inevitable, extensive contact between the fingers and support surface during the release. The proposed release/placement primitive is inspired by human release skills and utilizes the same methodology we employed to develop the compliant grasping primitive. The main idea is to servo control the hand in compliance to forces exerted on the fingers as the they open to release the object (see Fig.15).

Assuming the hand (with the object grasped) is located above the support surface, the release primitive begins by servoing the hand downward until contact with the surface is detected via continuous thresholding of the force seen at the wrist. Relying on the contact as a signal to stop the hand motion is inspired by human release strategy and has been used previously to trigger releasing objects (e.g., [28]).

Once contact between the hand/fingers and the support surface is established, we proceed by opening the fingers while compliantly servo-controlling the hand from the support surface. This is achieved by running two concurrent controllers: a velocity-based operational space controller to control the hand, and a position-based controller to open the fingers along their pre-defined trajectories. The proposed strategy follows the same scheme presented in Algorithm 2 with a minor difference in the desired force f_d which is calculated as $f_d =$

 $f(t) - f_r$, where the reference force f_r is recorded at the time of the initial contact between the hand/fingers and the support surface.

6 Experiments and Results

To validate and demonstrate the robustness and effectiveness of the proposed grasping primitives we have performed extensive experiments on a fully integrated manipulation system: a compliant 7-DOF Barrett Whole-Arm Manipulator (WAM) equipped with a Barrett 3-finger dexterous hand BH-280, and an integrated perception system. The positioning accuracy of the perception system varies depending on the object of interest. Overall, our system is capable of providing object pose with an accuracy of about 1-2 centimeters and 5-10 degrees error in position and orientation, respectively. In Section 4.1 we described the minimum hardware/Control requirements.

All objects in our experiments were located in their natural poses on a table top localized by the perception system. For each of the objects in our grasping experiments the vision system provides for an object a "launch" pose for the hand where the grasp is to be initiated. The launch pose is calculated in a way to restrict the plane specified by the fingertips parallel to the support surface with the hand centered above the object (or the target location for the release of object).

To grasp the object the hand assumes a predefined preshape weakly dependent on object geometry. Based on our empirical observations and given the flexibility of our proposed grasping approach, one could choose the same grasp pre-shape for objects with similar geometries. For example, for all objects with cylindrical shape (e.g., pen, screwdriver, hammer, pipe, etc.) we used a cup-like grasp pre-shape (as shown in Fig.4). In fact, we have observed that the cup-like pre-shape works remarkably well for grasping many of the small objects we used in our experiments. This also highly conforms with our anecdotal human behavior observations. The width of the finger opening to form the cup pre-shape can be chosen arbitrarily wide as long as it satisfies the localization accuracy of the perception system to make sure that the object can be caged.

We present the experimental results and our empirical observations in three categories: fingers placement/landing, object grasping, and object release/placement experiments as follows. A video of the grasp sequence is available online at http: //youtu.be/gxaXCYY87Z0.

6.1 Fingers landing experiments

An example of finger landing/placement experiment for grasping a hammer from a table top is shown in Fig.16. Initially the hand is at its launch pose centered above the hammer, see Fig.16(a). As it is seen the fingertips plane (or the hand palm) is not parallel to the table due to misalignment of the hand caused by uncertainties combined in perception and the robot kinematics. Clearly approaching the hammer along this orientation will not place all fingers in contact with the support surface (see Fig.16(b)), and hence will not yield a stable and robust grasp if executed. To fix the hand orientation and ensure all fingers are in contact with the support surface, the proposed finger landing approach servo controls the hand around appropriate control axes (as described in Section 5.1) based on the current finger touch states determined from the finger strain feedbacks. For example, in this experiment, the hand is initially servo controlled along the normal to the fingertips plane (Fig.16(a)) until the touch between finger F_2 and the support surface is detected (Fig.16(b)). As indicated in the plot the strain threshold to identify the touch is 50 units reported by the OWD and is about 0.5N force. Next the hand is servo controlled around the control axis at finger F_2 (Fig.16(c)) until finger F_3 reaches the support surface. Finally, the hand rotates around the control axis between fingers F_2 and F_3 until finger F_1 contacts the surface.

The plots in Fig.16 show the fingers' current contact status and the fingers which are expected to contact the table next at each instant of time. The fingers strain values along with the hand linear (along the normal of fingertips plane) and angular velocities during the entire landing process are also shown. The change in angular velocity direction denote the change of control point and axis when a finger contact the surface. Please note that we apply a fixed magnitude angular velocity.

Due to kinematic errors in calculating the control axis extra force might be exerted to a finger as the hand rotates to land other fingers. For example, the strain in finger F_2 increases even after its first contact with the support surface as shown in Fig. 16. However, the force compliant motion incorporated in the proposed landing technique prevents excessive force from being applied to the fingers by moving the hand away from the surface to decrease the strain on fingers and avoid damaging them.

The proposed landing primitive can be used for landing fingers from different pre-shapes. For example, Fig.17 shows landing experiments two different grasp pre-shapes. For every grasping experiment finger landing is first executed to ensure contact between all the fingers and the supporting surface. This is a key prerequisite to the robustness and success of our proposed compliant grasping approach as shown in the experiments which follow. This is also a unique and novel strategy which may find applications in other robotic manipulation tasks to place fingers on the surface of objects in a compliant scheme.

6.2 Object grasping experiments

Each grasping experiment starts off with all the fingers initially in contact with the supporting surface. This is ensured





Fig. 16: A typical compliant fingers landing experiment: the fingers (current and expected) touch states are continuously updated based on the strain gauges feedbacks as the hand is servo controlled to land the fingers.

using the finger landing strategy presented above. Figure 18 shows snapshots of a grasping experiment to grab a pipe from a table top. The plot shows a number of parameters including: status of fingers being stuck or reached their waypoints, fingers trajectories, compliant linear velocity of the hand along its palm normal, and the force applied to the hand along its palm normal. As it is shown, the hand responds compliantly to the wrist forces while the fingers close along their trajecto-



Fig. 17: Compliant fingers landing/placement experiments from different hand pre-shapes

ries. For example, between seconds 6 to 8 all the fingers have reached the body of the pipe (Fig.18(c)) and apply forces to the pipe body as well as the table to manage to go underneath the pipe (about second 8). At this moment the forces applied by the fingers pushes the hand upward while lifting the pipe and managing to fully cage it. The compliant motion of the hand plays a crucial role to prevent damage to the fingers. For example, there are other instances (about seconds 2 and 4) where the fingers are stuck but the compliant motion prevents the robot to exert excessive pressure to the fingertips. Many of the state of the art geometric grasp planning techniques either try to avoid the contact between the robot and support surface or are not capable of accounting for such contact, and hence fully ignore it.

The experiments presented here show that the coordination between the fingers and the hand movement plays a crucial role in achieving robust grasps of small objects. The simple yet effective compliant grasping primitive presented in this work successfully achieves this goal as shown through numerous experiments. Figure 19 shows representative examples of experiments we have performed to grasp a variety of small objects, a screwdriver, a pen, and a cellphone, etc. As shown the grasp used is a precision grasp using the fingertips to restrain the object. These successful results were not achievable without maintaining contact between the fingertips and the support surface and coordinating finger trajectories, behaviors which are fully integrated into our compliant grasping approach. Moreover, to compensate for uncertainties in placing the fingertips on the support surface, the finger landing primitive was used prior to grasp execution.

6.3 Object release experiments

The last set of our experiments demonstrate the compliant object release and placement. The proposed release strategy has



Fig. 18: Force compliant grasping of a pipe from a table top

been successfully applied to release objects from both enveloping as well as precision (fingertip) grasps. Due to space limitations we do not provide the plots explaining the details of the release strategy. However, as we noted before, the release strategy heavily borrows from the compliant grasping technique and one can view it as grasp execution but in reverse order. A representative example from numerous experiments which we performed is shown in Fig.20. The proposed release technique is highly robust to the uncertainties in localizing the height and orientation of the support surface and manages to gently release and replace objects on the support surface without damaging the fingers.

We conducted a series of grasp repeatability experiments on common objects found in our lab; namely a D-battery Maglite, a foot-long screwdriver, and a standard hammer. The objects



Fig. 19: Examples of force compliant grasping experiments



Fig. 20: An example of compliant object release/placement

Table 1: Compliant grasp performance on sample objects

Statistics	Maglite	Screwdriver	Hammer
Num. Trials	24	30	35
Success Rate	92%	93%	97%

were randomly positioned on a table, within the robot's reach. The perception system was first used to localize the object, followed by the sequence of motions to position the hand at the grasp launch pose. From that point, the sequence of *landing*, grasping, transportation and release was executed, and success rate measured. The results are shown in Table1. Success rate of 92% and above show that the approach is robust. Note that the vast majority of the failures were due to early testing failures, where faults in our finger reset procedures caused fingers to jam when contacting the table. The aforementioned compliant grasping primitives were devised and successfully implemented to address grasping tasks as a part of the DARPA Autonomous Robotic Manipulation challenge [1]. A number of grasping tests (hammer, screwdriver, maglight, and shovel) were performed independently by DARPA on a different robot using the software we provided. For those tests, careful attention to finger calibration was given and our success rate for the objects mentioned was 100%, even on objects not previously encountered.



Fig. 21: We have extensively employed the proposed primitives in developing and execution of a number of complex everyday manipulation tasks in the context of DARPA ARM-S program: for example unzipping a gym bag, grabbing a wire cutter from table and cutting a wire, and clearing a pile of unknown objects.

Since their development, we have extensively employed the proposed primitives in developing and execution of a number of complex everyday manipulation tasks in the context of DARPA ARM-S program: for example, grabbing a gym bag strap and unzipping it (see Fig. 21(a)) and grabbing a wire cutter from table and cutting a wire (see Fig. 21(b)). Also recently, the proposed primitives were used in the task of manipulating and clearing a pile of unknown objects (see Fig. 21(c)). Extensive experiments demonstrated the effectiveness and safety of our compliant grasping primitives in cluttered environments (please see [18, 19] for more results and details).

7 Conclusion

We addressed the problem of grasping objects that are small relative to an anthropomorphic hand, including a pen, screwdriver, cellphone, and hammer from their natural poses on a supporting surface, e.g., a table top. We argue that contact with support surfaces is critical for grasping small objects. We devised three simple, yet effective, manipulation primitives for robust grasping (and releasing) of small objects from support surfaces: (1) compliant finger landing for bringing all fingers safely in contact with the support surface, (2) compliant object grasping for maintaining the contact between the fingertips and the support surface during the finger closure, and (3)compliant object release. We conducted extensive grasping experiments on a variety of small objects with similar shape and size. The results demonstrate that our approach is robust to localization uncertainties and highlights the benefits of compliant, contact driven control strategies for grasping tasks.

ACKNOWLEDGMENTS

The authors gratefully acknowledge funding under the DARPA Autonomous Robotic Manipulation Software Track (ARM-S) program.

References

- J. A. Bagnell, F. Cavalcanti, L. Cui, T. Galluzzo, M. Hebert, M. Kazemi, M. Klingensmith, J. Libby, T. Yu Liu, N. Pollard, M. Pivtoraiko, J-S. Valois, and R. Zhu. System design and implementation for autonomous robotic manipulation. In *IROS*, 2012.
- Dmitry Berenson and Siddhartha Srinivasa. Grasp synthesis in cluttered environments for dexterous hands. In *IEEE-RAS International Conference on Humanoid Robots*, December 2008.
- A. Bicchi and V. Kumar. Robotic grasping and contact: a review. In ICRA, pages 348–353, 2000.
- R.C. Brost. Automatic grasp planning in the presence of uncertainty. *Int. J. Robot. Res.*, 7(1):3 – 17, 1988.
- 5. Umberto Castiello. The neuroscience of grasping. *Nature Reviews Neuroscience*, 9(6):726736, 2005.
- Matei T. Ciocarlie and Peter K. Allen. Hand posture subspaces for dexterous robotic grasping. *IJRR*, 28(7):851 – 867, 2009.
- Rosen Diankov. Automated Construction of Robotic Manipulation Programs. PhD thesis, Carnegie Mellon University, Robotics Institute, August 2010.
- M.R. Dogar and S.S. Srinivasa. Push-grasping with dexterous hands: Mechanics and a method. In *IROS*, pages 2123 –2130, 2010.
- A. M. Dollar and R. D. Howe. Towards grasping in unstructured environments: Grasper compliance and conguration optimization. *Advanced Robotics*, 19(5):523–544, 2005.
- M.A. Erdmann and M.T. Mason. An exploration of sensorless manipulation. *IEEE J. Robot. Autom.*, 4(4):369 – 79, 1988.
- K. Harada, M. Kaneko, and T. Tsujii. Rolling-based manipulation for multiple objects. *IEEE Trans. Robot. Autom.*, 16(5):457 – 68, 2000.
- Paul Hebert, Nicolas Hudson, Jeremy Ma, Thomas Howard, Thomas Fuchs, Max Bajracharya, and Joel W. Burdick. Combined shape, appearance and silhouette for simultaneous manipulator and object tracking. In *ICRA*, pages 2405–2412, May 2012.
- N. Hogan. Impedance Control: An Approach to Manipulation: Part ITheory. J. Dyn. Sys., Meas., Control, 107(1):1–7, 1985.
- K. Hsiao, L. Kaelbling, and T. Lozano-Perez. Robust grasping under object pose uncertainty. *Autonomous Robots*, 31(2-3):1–16, 2011.
- Shervin Javdani, Matthew Klingensmith, J. Andrew (Drew) Bagnell, Nancy Pollard, and Siddhartha Srinivasa. Efficient touch based localization through submodularity. In *ICRA*, May 2013.
- Roland S. Johansson and Randall J. Flanagan. Coding and use of tactile signals from the fingertips in object manipulation tasks. *Nature Reviews Neuroscience*, 10(5):345–359, 2009.
- Mrinal Kalakrishnan, Ludovic Righetti, Peter Pastor, and Stefan Schaal. Learning force control policies for compliant manipulation. In *IROS*, pages 4639–4644, 2011.
- Dov Katz, Moslem Kazemi, J. Andrew Bagnell, and Anthony Stentz. Clearing a pile of unknown objects using interactive perception. In Proceedings of IEEE International Conference on Robotics and Automation, March 2013.
- Dov Katz, Arun Venkatraman, Moslem Kazemi, J. Andrew Bagnell, and Anthony Stentz. Perceiving, learning, and exploiting object affordances for autonomous pile manipulation. In *Robotics: Science* and Systems, June 2013.
- Moslem Kazemi, Jean-Sebastien Valois, J. Andrew Bagnell, and Nancy Pollard. Robust object grasping using force compliant motion primitives. In *Robotics: Science and Systems*, July 2012.
- O. Khatib. A unified approach for motion and force control of robot manipulators: The operational space formulation. *IEEE Journal of Robotics and Automation*, 3(1):43–53, 1987.
- A. Liegeois. Automatic supervisory control of the configuration and behavior of multibody mechanisms. *IEEE Trans. Syst. Man Cybern.*, SMC-7(12):868 – 71, 1977.
- Matthew T. Mason. Compliance and force control for computer controlled manipulators. *IEEE Trans. Syst. Man Cybern.*, 11(6):418 – 432, 1981.

- 24. A.T. Miller and P.K. Allen. Graspit! a versatile simulator for robotic grasping. *IEEE Robot. Autom. Mag.*, 11(4):110 22, 2004.
- J. Nakanishi, R. Cory, M. Mistry, J. Peters, and S. Schaal. Operational space control: a theoretical and empirical comparison. *Int. J. Robot. Res.*, 27(6):737 57, 2008.
- 26. Van-Duc Nguyen. Constructing stable grasps. Int. J. Robot. Res., 8 (1):26 37, 1989.
- Jr. Platt, R., A.H. Fagg, and R.A. Grupen. Null-space grasp control: Theory and experiments. *IEEE Trans. Robot.*, 26(2):282 – 95, 2010.
- J.M. Romano, K. Hsiao, G. Niemeyer, S. Chitta, and K.J. Kuchenbecker. Human-inspired robotic grasp control with tactile sensing. *IEEE Trans. Robot.*, 27(6):1067 –1079, 2011.
- A. Saxena, J. Driemeyer, and A. Ng. Robotic grasping of novel objects using vision. *IJRR*, 27(2):157 173, 2008.
- T. Sugaiwa, Y. Yamaguchi, H. Iwata, and S. Sugano. Dexterous hand-arm coordinated manipulation using active body-environment contact. In *IROS*, pages 4168–4173, 2009.
- K. Tahara, S. Arimoto, and M. Yoshida. Dynamic force/torque equilibrium for stable grasping by a triple robotic fingers system. In *IROS*, pages 2257–2263, 2009.
- J.C. Trinkle, J. M. Abel, and R. P. Paul. An investigation of frictionless enveloping grasping in the plane. *Int. J. Robot. Res.*, 7(3):33–51, June 1988.
- D.E. Whitney. Force feedback control of manipulator fine motions. J. Dyn. Sys., Meas., Control, 99(2):91–97, 1977.
- T. Yoshikawa. Force control of robot manipulators. In *ICRA*, volume 1, pages 220–226, 2000.