Human-Inspired Robotic Grasp Control with Tactile Sensing

Joseph M. Romano, *Student Member, IEEE,* Kaijen Hsiao, *Member, IEEE,* Günter Niemeyer, *Member, IEEE,* Sachin Chitta, *Member, IEEE,* and Katherine J. Kuchenbecker, *Member, IEEE*

***Abstract*—We present a novel robotic grasp controller that allows a sensorized parallel jaw gripper to gently pick up and set down unknown objects once a grasp location has been selected. Our approach is inspired by the control scheme that humans employ for such actions, which is known to centrally depend on tactile sensation rather than vision or proprioception. Our controller includes the six discrete states of Close, Load, Lift and Hold, Replace, Unload, and Open. During all control states, measurements from the gripper's fingertip pressure arrays and hand-mounted accelerometer are processed in real time to generate robotic tactile signals that are designed to match human SA-I, FA-I, and FA-II channels. These signals are combined into tactile event cues that drive the state transitions, enabling the controller to select an appropriate initial grasping force, detect when an object is slipping from the grasp, and judge when to release an object to set it down. We demonstrate the promise of our human-inspired approach to robotic grasp control through implementation on the PR2 robotic platform, including grasp testing on a large number of real-world objects.**

***Index Terms*—robot grasping, tactile sensing**

I. Introduction

As robots move into human environments, they will need to know how to grasp and manipulate a very wide variety of objects [1]. For example, some items may be soft and light, such as a stuffed animal or an empty cardboard box, while others may be hard and dense, such as a glass bottle or an apple. After deciding **where** such objects should be grasped (finger placement), the robot must also have a concept of **how** to execute the grasp (finger forces and reactions to changes in grasp state). A robot that operates in the real world must be able to quickly grip a wide variety of objects firmly, without dropping them, and delicately, without crushing them (Fig. 1).

Non-contact sensors such as cameras and laser scanners are essential for robots to recognize objects and plan where to grasp them, e.g., [2], [3]. Similar sensing approaches could also be used to anticipate the grasp force needed to safely pick up an object, but it is impractical to store specific information about all of the objects a robot may need to handle. Furthermore, it is dangerous to rely on such a priori models, since the grasp forces they recommend may be too strong or too weak for a specific item. Instead, we and others believe that robot grasp control should strongly rely on tactile sensing – local time-varying information about the contact between the robot's fingers and the object in question. Many important events that can be challenging for other sensory modalities to perceive, such as the slip of an object in the fingers or a glancing collision between the held object and an unseen obstacle, are easily detected using carefully crafted tactile signals. Understanding contact using tactile information and reacting in real time will be critical skills for robots to successfully interact with real-world objects, just as they are for humans.

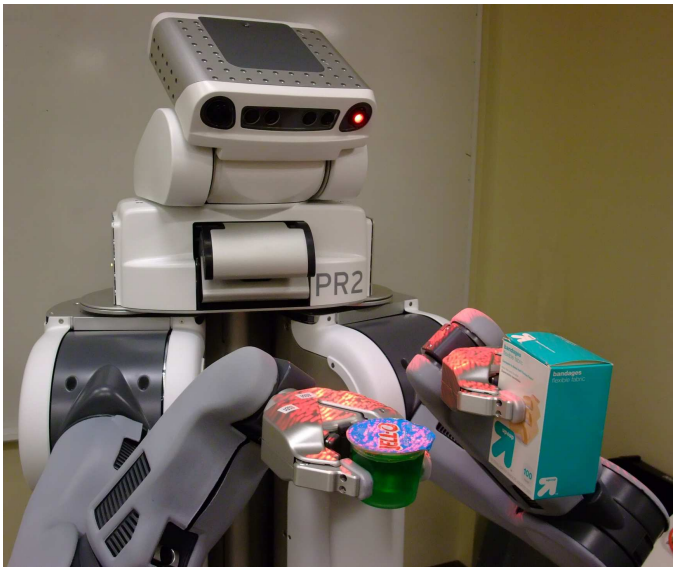

Fig. 1. The Willow Garage PR2 robot using our grasp controller to carefully handle two sensitive everyday objects.

A. Human Grasping

Neuroscientists have thoroughly studied the human talent for grasping and manipulating objects. As recently reviewed by Johansson and Flanagan [4], human manipulation makes great use of tactile signals from several different types of mechanoreceptors in the glabrous (non-hairy) skin of the hand, with vision and proprioception providing information that is less essential. Johansson and Flanagan divide the seemingly effortless action of picking up an object and setting it back down into seven distinct states: reach, load, lift, hold, replace, unload, and release. In the first phase, humans close their grasp to establish finger contact with the object. Specifically, the transition from reach to load is known to be detected through the FA-I (Meissner) and FA-II (Pacinian) afferents, which are stimulated by the initial fingertip contact. FA signifies

Manuscript received August 1, 2010; revised February 15, 2011.

J. M. Romano and K. J. Kuchenbecker are with the Department of Mechanical Engineering and Applied Mechanics, University of Pennsylvania, Philadelphia, PA, 19104 USA. e-mail: {jrom, kuchenbe}@seas.upenn.edu

K. Hsiao, G. Niemeyer and S. Chitta are with Willow Garage, Menlo Park, CA, 94025 USA. e-mail: {hsiao, gunter, sachinc}@willowgarage.com

that these mechanoreceptors are fast-adapting; they respond primarily to changes in mechanical stimuli, having small and large receptive fields, respectively. Once contact has been detected, humans increase their grasp force to the target level, using both pre-existing knowledge about the object and tactile information gathered during the interaction. This loading process is regulated largely by the response of the SA-I (Merkel) afferents, which are slowly-adapting with small receptive fields. The load phase ends when the target grasp force is reached with a stable hand posture.

Once the object is securely grasped, humans use their arm muscles to lift up the object, hold it in the air, and possibly transport it to a new location. Corrective actions (typically increases in grip force) are applied during the lifting and holding phases when the tactile feedback does not match the expected result. Srinivasan et al. [5] showed that the FA-I and FA-II signals are the primary sources of information for detecting both fingertip slip and new object contact. Slip is of critical importance to reject disturbances in the lifting and holding phases, while object contact must be detected during the replace stage to successfully transition to unloading. The SA-I afferents are again important during unload to properly set the object down before full release. These tactile sensing capabilities and corrective reactions enable humans to adeptly hold a very wide range of objects without crushing or dropping them; indeed, humans typically apply a grip force that is only 10–40% more than the minimum amount needed to avoid slippage [4], thereby achieving the dual goals of safety and efficiency.

B. Our Approach: Human-Inspired Robotic Grasp Control

Inspired by the fluidity of human grasp control, this article presents a set of methods that enable a robot to delicately and firmly grasp real-world objects once the fingertip contact locations have been selected. We describe robotic sensing methods that use finger-mounted pressure arrays and a hand-mounted accelerometer to mimic the important tactile signals provided by human FA-I, FA-II, and SA-I mechanoreceptors. These three complementary sensory channels allow us to create a high-level robotic grasp controller that emulates human tactile manipulation: in the words of Johansson and Flanagan, our controller is *"centered on mechanical events that mark transitions between consecutive action phases that represent subgoals of the overall task"* [4]. As diagrammed in Fig. 2, our approach separates robotic grasping into six discrete states:

- Close
- Load
- Lift and Hold
- Replace
- Unload
- Open

These states purposefully match those of human grasping, although we have combined Lift and Hold because their control responses are nearly identical. Each state defines a set of rules for controlling a robotic gripper to perform the specified behavior based on the tactile sensations it experiences. In addition to creating this human-inspired approach to robotic grasp control, we implemented our methods on the standardized hardware and software of the Willow Garage PR2 robot; our goal was to enable it to perform two-fingered grasps on typical household objects at human-like speeds, without crushing or dropping them.

Section II summarizes previous work in the area of tactile robotic grasping and substantiates the novelty of our approach. Section III describes pertinent attributes of the PR2 platform, while Section IV defines our robotic SA-I, FA-I, and FA-II tactile channels and the low-level position and force control strategies we created for the PR2's high-impedance gripper. Section V expounds on the control diagram of Fig. 2 by carefully defining each control rule and state transition. As described in Section VI, we validated our methods through experiments with the PR2 and a large collection of everyday objects under a variety of challenging test conditions. We conclude the article and discuss our plans for future work in Section VII.

II. Background

The development of tactile sensors for robotic hands has been a very active area of research, as reviewed by Cutkosky et al. in 2008 [6]. Among the wide range of sensors one could use, Dahiya et al. [7] present a strong case for the importance of having tactile sensors capable of reproducing the rich

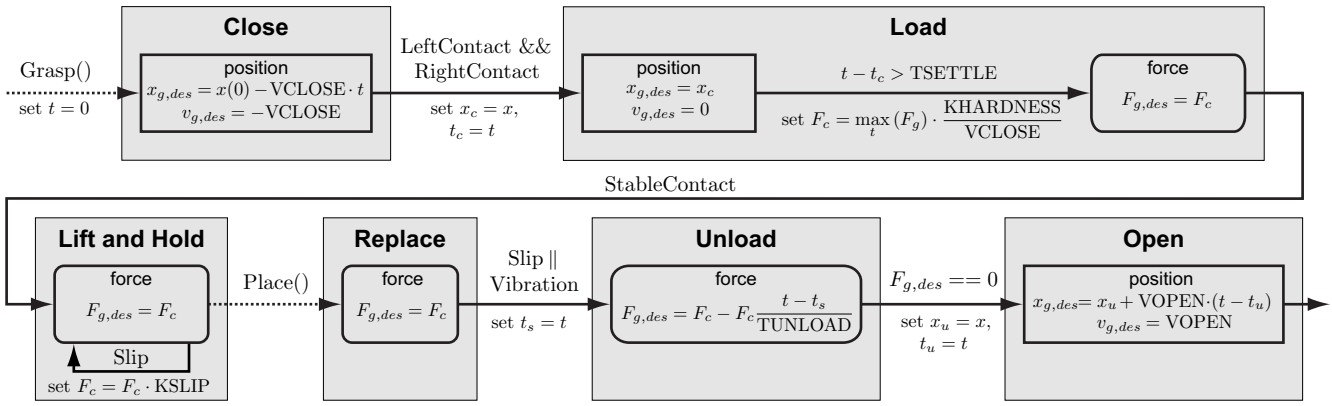

Fig. 2. The state diagram for our human-inspired robotic grasp controller. State transitions occur only after specific tactile events are detected. The details of this controller are presented in Sections IV and V. Constant-valued parameters, such as VCLOSE, are defined in Table II in the Appendix.

response of the human tactile sensing system. Along these lines, some recent sensors have even achieved dynamic response capabilities on par with the glabrous skin of the human fingertip [8]. Using tactile sensory cues with wide dynamic range as the trigger for robotic actions was first proposed by Howe et al. over two decades ago [9]. Unfortunately, most such sensors still exist only as research prototypes and are not widely available. The pressure-sensing arrays used in our work represent the state of the art in commercially available tactile sensors, and they are available on all PR2 robots. High-bandwidth acceleration sensing is far more established, though such sensors are rarely included in robotic grippers.

In this work, we present a fully autonomous PR2 robot that can perceive its environment, pick up objects from a table, and set them back down in a new location. We build on the overall system described in [10], focusing on methods for using tactile feedback (pressure and acceleration) to improve the gripper's contact interactions with grasped objects. Prior versions of this system have used pre-defined grasp forces and were limited to either crushing or dropping many objects; hence, our primary goal is to enable two-fingered grasps that are gentle but secure.

Several other research groups have recently developed autonomous pick-and-place robotic systems, such as [11]–[13]. While these other works often use tactile sensing to guide adjustments in hand position during the pre-grasp stage, they rely mainly on pre-defined grasp forces or positions and the mechanical compliance of their gripper in order to hold objects. Such strategies do not work well for the high-impedance parallel jaw grippers with which many robots are equipped. In other previous work, Yussof et al. [14] showed promising results using a custom-built tabletop manipulator with custom optical fingertip sensors that measure normal and shear forces. Takahashi et al. [15] calculated robot fingertip slip information using a pressure array similar to that used in our system. However, neither of these systems has yet been validated with a wide set of real-world objects, and we believe there is room for new approaches to processing and responding to tactile signals.

Another large but scattered body of work is devoted to understanding fingertip sensor signals at and during contact, e.g., [16]–[19]. Again, these algorithms are typically developed with custom hardware and validated against only a small and ideal set of objects, which makes it hard to draw conclusions about their general utility. In some of our own previous work, we found that sensor information such as slip [20] and contact stiffness [21] is useful in well controlled situations but can can be extremely susceptible to changes in object texture, sensor orientation, and other such factors that will naturally vary when implemented on a fully autonomous robot that needs to interact with unknown objects in human environments.

This paper aims to develop robust tactile sensory signals that will apply to as many objects as possible in order to create a tactile-event-driven model for robotic grasp control. Several researchers [7], [9] have noted the importance of such an approach, but to our knowledge this is among the first large-scale implementations of such methods in a fully autonomous robot that has not been built for the sole purpose of grasping.

III. Robot Experimental System

Robots have great potential to perform useful work in everyday settings, such as cleaning up a messy room, preparing and delivering orders at a restaurant, or setting up equipment for an outdoor event [1]. Executing such complex tasks requires hardware that is both capable and robust. Consequently, we use the Willow Garage PR2 robotic platform. As shown in Fig. 1, the PR2 is a human-sized robot designed for both navigation and manipulation. It has an omni-directional wheeled base, two seven-degree-of-freedom arms, and two one-degree-of-freedom parallel-jaw grippers. Its extensive non-contact sensor suite includes two stereo camera pairs, an LED pattern projector, a high-resolution camera, a camera on each forearm, a head-mounted tilting laser range finder, a body-mounted fixed laser range finder, and an IMU.

Figure 3 shows the parallel jaw gripper that is mounted on each of the PR2's arms. The gripper's only actuator is a brushless DC motor with a planetary gearbox and an encoder. This motor's rotary motion is converted to the parallel jaw motion through a custom internal mechanism in the body of the gripper. Unlike many of the robot hands currently being developed for grasping, e.g., [11], [13], the PR2 gripper has a high mechanical impedance due to the large gear ratio of the actuation system; note that it can be slowly back-driven by applying a large force at the fingertips. Motor output torque can be specified in low-level software, but the transmission does not include any torque or force sensors. Instead, motor effort is estimated using current sensing; the transmission's significant friction prevents this signal from corresponding well with the force that the gripper applies to an object during grasping.

A high-bandwidth Bosch BMA150 digital accelerometer is embedded in the palm of each gripper, as illustrated in Figure 3. This tiny sensor measures triaxial acceleration over a range of ± 78 m/s^2 with a nominal resolution of 0.15 m/s^2. The accelerometer has a sampling rate of 3 kHz and a bandwidth from DC to 1.5 kHz. Successive triaxial acceleration measurements are grouped together in sets of three (nine total values) and made available to the controller at a rate of 1 kHz.

Each of the gripper's two 2.3 cm $\times$ 3.7 cm $\times$ 1.1 cm fingertips is equipped with a pressure sensor array consisting of 22 individual cells. The 22 cells are divided between a five by three array on the parallel gripping surface itself, two sensor elements on the end of the fingertip, two elements on each side of the fingertip, and one on the back (see Fig. 3). These capacitive sensors (manufactured by PPS Systems) measure the perpendicular compressive force applied in each sensed region, and they have a nominal resolution of 6.25 mN. As shown in Fig. 3, the entire sensing surface is covered by a protective layer of silicone rubber that provides the fingertip compliance and friction needed for successful grasping. All pressure cells are sampled simultaneously at a rate of 24.4 Hz. Due to manufacturing imperfections and residual stresses in the deformed rubber, each sensor cell has a unique non-zero reading when the fingertips are subjected to zero force. We compensate for this non-ideality by averaging the first 0.25 seconds of each cell's pressure measurements before each

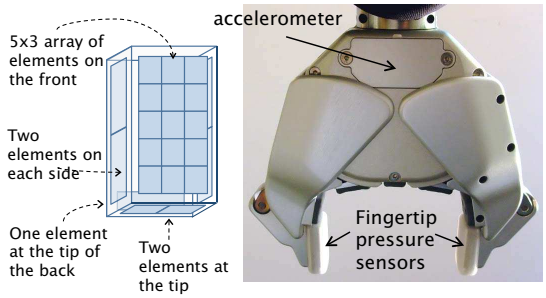

Fig. 3. The PR2 robot gripper. The accelerometer is rigidly mounted to a printed circuit board in the palm, and the pressure sensors are attached to the robot's fingertips under the silicone rubber coating.

grasp and subtracting this offset from subsequent readings.

We use the open-source ROS software system for all experiments conducted with the PR2 robot. The implemented gripper controllers are run inside a 1 kHz soft-real-time loop. Information from the tactile sensors (accelerometer and pressure cells) is available directly in this environment, as is the gripper's encoder reading.

IV. Low-Level Signals and Control

Individual sensor readings and actuator commands are far-removed from the task of delicately picking up an object and setting it back down on a table. Consequently, the high-level grasp controller diagrammed in Fig. 2 rests on an essential low-level processing layer that encompasses both sensing and acting. Here, we describe the three tactile sensory signals that we designed to match human SA-I, FA-I, and FA-II afferents, along with the position and force controllers needed for the gripper to move smoothly and interact gently with objects.

A. Fingertip Force (SA-I)

The SA-I tactile sensory channel is known to play a primary role in the human hand's sensitivity to steady-state and low-frequency skin deformations up to ∼5 Hz [4]. The cumulative response of many SA-I nerve endings gives humans the ability to both localize contact on the fingertip and discern the total amount of force applied [22]. In humans, these receptors populate the non-hairy skin of the finger with densities ranging from 40 to 100 per cm^2, with each neuron having a receptive field of ∼2–4 mm^2 [22].

The pressure arrays on the PR2 fingertips provide information that is similar to human SA-I signals, though spatially far less dense. A signal similar to the SA-I estimate of total fingertip force can be obtained by summing the readings from all fifteen elements in the pad of one finger on the gripper:

$$F_{gl} = \sum_{i=1}^{3}\sum_{j=1}^{5} f_{l(i,j)} \tag{1}$$

The value $f_{l(i,j)}$ represents the force acting on the left fingerpad cell at location i, j. The same procedure is used to calculate F_{gr} for the right finger using $f_{r(i,j)}$. The mean grip force is calculated by averaging the force experienced by the two fingers, $F_g = \frac{1}{2}(F_{gl} + F_{gr})$. Figure 4 shows an example of these fingertip force signals during object contact, along with examples of the other tactile signals described below.

We also attempted to obtain localization information about the fingertip contact using the methods described in [15], but we were not able to achieve satisfactory results. We hypothesize that differences in the fingertip shape (the PR2 has a flat fingertip that leads to multiple contact locations, as opposed to the rounded fingertips used by Takahashi et al.), and the lower number of tactile cells in our sensors were the main reasons we were unable to achieve similar results.

B. Fingertip Force Disturbance (FA-I)

Human FA-I signals are believed to be the most important indicator of force-disturbance events during bare-handed manipulation. These force disturbances occur at many instants including the initial object contact, object slippage, impacts between a hand-held object and the environment, and the end of object contact. FA-I afferents respond to skin deformations in the 5–50 Hz frequency range [4]. The FA-I receptors populate glabrous human skin with densities ranging from 70 to 140 per cm^2, and each receptor has a receptive field of ∼3–5 mm^2 [23].

We process the data from the PR2's pressure arrays to create a signal similar to the human FA-I channel. Our chosen calculation is to sum a high-pass-filtered version of the forces detected in all fifteen fingertip cells:

$$\widetilde{F}_{gl}(z) = \sum_{i=1}^{3}\sum_{j=1}^{5} H_F(z) f_{l(i,j)}(z) \tag{2}$$

The force measured in each cell $f_{l(i,j)}$ is subjected to a discrete-time first-order Butterworth high-pass filter $H_F(z)$ with a cutoff frequency of 5 Hz, designed for the 24.4 Hz sampling rate of the pressure signals. The resulting filtered signals are then summed to obtain an estimate of the > 5 Hz force disturbances $\widetilde{F}_{gl}$ acting on the entire left fingerpad. The process is repeated for the right finger to obtain $\widetilde{F}_{gr}$.

C. Hand Vibration (FA-II)

FA-II afferents are known to be the primary tactile channel by which humans sense interactions between a handheld tool and the items it is touching. During object grasping and manipulation, these receptors are particularly useful for detecting contact between handheld objects and other things in the environment, such as a table surface. The FA-II mechanoreceptors respond to high-frequency vibrations in a range from 40 to 1000 Hz. They are relatively rare in the hand, and they have a wide receptive field of at least 20 mm^2 [22], [23].

We create a robotic analog to the FA-II sensory channel by processing data from the PR2's hand-mounted accelerometer, as follows:

$$\widetilde{a}_h(z) = \sqrt{(H_a(z)a_{h,x})^2 + (H_a(z)a_{h,y})^2 + (H_a(z)a_{h,z})^2} \tag{3}$$

The hand vibration signal $\widetilde{a}_h$ is calculated by taking the magnitude of the high-pass-filtered three-dimensional acceleration vector. The filter applied to each of the three Cartesian acceleration components ($a_{h,x}$ $a_{h,y}$ $a_{h,z}$) is a discrete-time

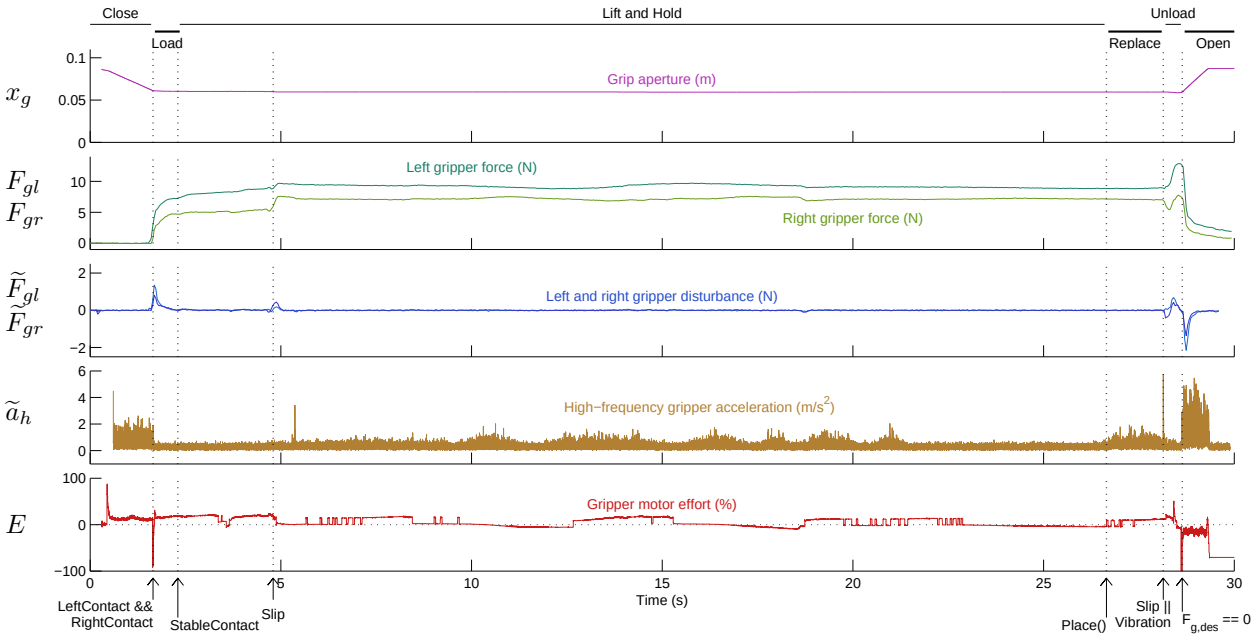

Fig. 4. Time history data for an interaction between the gripper and an Odwalla juice bottle. Controller states are marked with bars at the top, and important signals (state transitions and tactile events) are indicated with arrows at the bottom.

first-order Butterworth high-pass filter $H_a(z)$ with a 50 Hz cutoff frequency, designed for the 3 kHz sampling rate of our acceleration data stream.

D. Position and Force Control

In addition to the rich tactile sensations described above, humans excel at manipulation because they can move competently through free space but quickly transition to regulating grasp force during object contact [4]. Replicating the fluidity of human grasping with a parallel jaw gripper thus requires well designed position and force controllers. Both of these controllers appear several times in the high-level state diagram of Fig. 2; each controller block is labeled with its type, along with the desired motion or force output.

The PR2 gripper is a geared mechanism, so it lends itself well to position control. We define its position x_g in meters and its velocity v_g in meters per second. The position is zero when the fingers touch and positive otherwise, so that the position value corresponds to the grip aperture. The gripper velocity follows the same sign as position, with positive values indicating that the hand is opening. We found that we could achieve good position tracking via a simple proportional-derivative controller with an additional velocity-dependent term to overcome friction:

$$\begin{aligned} E &= \text{KP} \cdot (x_g - x_{g,des}) + \text{KD} \cdot (v_g - v_{g,des}) \\ &\quad -\text{sign}(v_{g,des}) \cdot \text{EFRICTION} \end{aligned} \tag{4}$$

Here, E is the motor effort (N), KP is the proportional error gain (N/m), KD is the derivative error gain (Ns/m),and $x_{g,des}$ and $v_{g,des}$ are the desired gripper position (m) and velocity (m/s) respectively. EFRICTION is a scalar constant for feed-forward friction compensation, applied to encourage motion in the direction of $v_{g,des}$. Note that motor effort is defined to be positive in the direction that closes the gripper, which is opposite the sign convention for the motion variables. Table II in the Appendix lists values and units for all of the constants used in our controllers, including KP, KD, and EFRICTION.

We created a force controller on top of this position controller to enable the PR2 to better interact with delicate objects. This controller requires access to the fingertip force signals F_g described above in Section IV-A. Forces that compress the fingertips are defined to be positive, so that positive motor effort has the tendency of creating positive fingertip forces. As is commonly done, the force controller drives the desired position and velocity terms based on the error between the desired force and the actual force:

$$F_{g,min} = \min\left(F_{gl}, F_{gr}\right) \tag{5}$$

$$v_{g,des} = \text{KF} \cdot \left(F_{g,min} - F_{g,des}\right) \tag{6}$$

$$\text{KF} = \begin{cases} \text{KFCLOSE} & \text{if } F_{g,min} - F_{g,des} < 0, \\ \text{KFOPEN} & \text{otherwise} \end{cases} \tag{7}$$

The force we servo the grasp on, $F_{g,min}$, is the smaller of the two sensed fingertip forces (5), which helps to ensure that a strong dual-finger contact is made with the object. Errors in tracking the desired force $F_{g,des}$ are multiplied by the constant gain KF to yield the desired velocity for the position controller (6). This desired velocity is integrated over time to provide the position controller with a desired grip aperture $x_{g,des}$. Experimental testing revealed that high values of the gain KF improved force tracking but also caused the commonly encountered force-controller effect of chattering. We found that an asymmetric gain definition (7), where KFCLOSE is greater than KFOPEN, allows for the best balance of stability and responsiveness during grasping. We believe this asymmetry is most likely due to mechanical features of the

gripper's complex drive train, such as compliance, backlash, and slight backdriveability.

V. Robotic Grasp Control

This section describes the high-level grasp controller we developed to enable a robotic gripper to mimic the human capacity for picking up and setting down objects without crushing or dropping them. This control architecture is diagrammed in Fig. 2 and explained sequentially below. It makes use of all three of the low-level tactile signals and both of the low-level controllers defined in the previous section.

A. Close

Closing is the starting point for any grasp. In this phase the hand goes from having no contact with the object to contacting the object with both fingers. The Close state begins when a client software module requests a grasp, denoted by the Grasp() signal in Fig. 2. It is assumed that prior to this request the gripper has been maneuvered so that its fingers are on opposing sides of an object, approximately equidistant to both surfaces, and that closing the gripper's jaws will result in stable two-fingered contact with the object. Selection of such poses without requiring a model of the object is an ongoing area of research, e.g., [3], [10], as is correcting for errors in execution of the selected pose, e.g., [10], [13]. Here, we assume a reasonable starting arm pose and focus solely on controlling the gripper's one degree of freedom.

After the Grasp() request has been received, we use the position controller defined in (4) to close the gripper with a constant desired velocity $v_{g,des} = -\text{VCLOSE}$. This closing movement continues until contact is detected on both the left and right fingers, which we define as the logical *and* of the following two signals:

$$\text{LeftContact} = (F_{gl} > \text{FLIMIT}) \parallel (\widetilde{F}_{gl} > \text{DLIMIT}) \quad (8)$$

$$\text{RightContact} = (F_{gr} > \text{FLIMIT}) \parallel (\widetilde{F}_{gr} > \text{DLIMIT}) \quad (9)$$

Each fingertip force signal is compared with the constant FLIMIT, and each fingertip force disturbance signal is compared with the constant DLIMIT. Note that a finger's contact signal is true if the threshold on either its F_g or its $\widetilde{F}_g$ has been exceeded, and recall that the second of these signals is merely a high-pass-filtered version of the first. Through implementation on the PR2, we found that DLIMIT can be much smaller than FLIMIT because $\widetilde{F}_{gl}$ is not sensitive to the pressure cells' low-frequency fluctuations; consequently, the force disturbance condition is almost always first to trigger a contact event. A sample instance of LeftContact && RightContact is indicated in Fig. 4.

While several groups [5], [9] have noted that higher frequency vibrations are useful in detecting contact, we could not find a reliable indication of such events in our $\widetilde{a}_h$ signal. This difficulty is likely due to the specific hardware of the PR2. A large distance separates the accelerometer from the fingertips, and the compliant silicone fingertip coverings soften impacts considerably. Furthermore, significant high-frequency vibration noise occurs whenever the gripper motor is turning, which masks out any small tactile cues that might occur (see the $\widetilde{a}_h$ signal during the Close and Open phases in Fig. 4). Fortunately, the tactile signals derived from the pressure cells are very reliable at detecting the start of contact, providing a consistent cue to transition to the Load phase.

B. Load

The goal of the Load phase is to apply a grip force that is appropriate for the target object, so that the object can be lifted and manipulation can begin. Selecting this grip force level is challenging when the robot does not have a detailed mechanical model of the object or prior experience in gripping it. Consequently, we designed a novel method for choosing a reasonable starting grasp force; many alternatives were tested, and we report here only the most reliable one.

After contact is detected, the robot pauses for a short period of time (TSETTLE) while trying to hold the position of the first contact x_c. This pause serves two purposes: first, many objects take a short time to mechanically respond to the initial contact, and second, the pressure sensor cells update at a rate that is much slower than the controller rate (24.4 Hz vs. 1000 Hz). *We have found that the force response during this contact settling time is a very useful indicator of how hard an object should be grasped.* Thus, the robot records the maximum average force seen by the gripper fingers during this time and calculates the target force to hold the object as:

$$F_c = \max_t(F_g) \cdot \frac{\text{KHARDNESS}}{\text{VCLOSE}} \quad (10)$$

As one would expect, the force response of an object strongly depends on the speed of the finger impact. We remove this dependence on velocity by dividing the gain KHARDNESS by VCLOSE, the speed at which the fingers are commanded to impact the object. This approach enables the controller to calculate a contact force F_c that is relatively independent of the closing speed. At the end of the brief pause, the robot transitions to force control, using the computed F_c value as the desired force $F_{g,des}$.

The force control mode continues until the gripper achieves StableContact, which we define as:

$$\begin{aligned}\text{StableContact} = {} & (|F_{g,min} - F_{g,des}| < \text{FTHRESH}) \\ & \&\& \ (|v_g| < \text{VTHRESH}) \end{aligned} \quad (11)$$

This condition requires that the smaller of the two fingertip forces $F_{g,min}$ be within the tolerance FTHRESH of the command force $F_{g,des}$, and that the gripper speed be below the tolerance VTHRESH. Through testing, we have found that grasps become stable very quickly on most everyday objects. Slower stabilizations occur with very hard objects, which require additional effort after contact to ramp up the steady-state force command, and extremely soft objects, which need some time before the velocity settles to zero.

C. Lift and Hold

After StableContact is achieved, the controller transitions to the next grasp phase: Lift and Hold. In this phase the robot holds the object between its fingertips and moves its arm to accomplish higher-level tasks. It is desirable for the grasp

controller to hold the object firmly enough to avoid slipping, but gently enough to avoid crushing. As with FA-I signals in human grasping, the high-pass filtered force signal $\widetilde{F}_g$ is a strong indicator of slip. This signal is more reliable than F_g itself, since it does not vary significantly with low-frequency robot motion and reorientation of the object with respect to gravity. We calculate the Slip condition as follows:

$$\begin{aligned} \text{Slip} \quad = \quad & \left(|\widetilde{F}_g| > F_g \cdot \text{SLIPTHRESH}\right) \\ & \&\& \ \left(F_g^{BP} < \text{FBPTHRESH}\right) \end{aligned} \tag{12}$$

Our Slip condition is met only when both subsidiary sensory comparisons evaluate to true. First, the magnitude of the force disturbance signal $\widetilde{F}_g$ must exceed a threshold defined by the product of the total force F_g and a constant SLIPTHRESH; using this product rather than a constant value makes the robot less sensitive to force variations as the grip force increases. This approach was again inspired by human capabilities: human force perception is known to follow Weber's Law, where a stimulus must vary by a certain percent of its magnitude to have the change be detectable [24].

Second, slips are considered only when the grasp force is not slowly varying. We evaluate this force stability condition by treating the average fingerpad force F_g with a first-order Chebyshev discrete-time band-pass filter with a pass-band from 1 to 5 Hz and comparing with the empirically tuned constant FBPTHRESH. The lower frequency cutoff of 1 Hz removes the mean value from the grip force signal, while the higher cutoff of 5 Hz removes the slip effects that are seen in $\widetilde{F}_g$. This condition prevents the formation of a feedback loop when the controller increases its grip force to stop slip events, as discussed below.

If Slip does occur, the controller increases the desired grip force F_c by a small percentage of its current value, such that $F_c = F_c \cdot \text{KSLIP}$. Several alternative methods of responding to slip were tested, such as increasing the desired grip force proportional to the magnitude of the slip event, as done by Takahashi et al. [15]. However, we found that our system does not exhibit a strong correlation between the amount of slip (in either speed or distance) and the $\widetilde{F}_g$ signal; instead, this signal depends somewhat on the properties of the grasped object, so we use it to detect only that the grasp has been disturbed.

D. Replace

The transition from Lift and Hold to Replace occurs when a client software module sends the Place() signal. The robot should enter this mode only when it is ready to put the held object down, which is typically after it has moved the gripper to within a few centimeters of the desired placement location. After issuing the Place() command, the higher-level robot controller moves the gripper toward the target surface at a moderate speed. During this time, the low-level force controller hold the final target force from the previous phase, F_c, and the Replace controller monitors the Slip and Vibration signals to detect contact between the object and environment. We define the Vibration condition as a threshold on the high-frequency hand acceleration signal:

$$\text{Vibration} = (\widetilde{a}_h > \text{ATHRESH}) \tag{13}$$

When either Slip or Vibration becomes true, the robot assumes that contact has occurred between the object and the target surface, and it moves into the Unload phase.

It is necessary to observe both the Slip and Vibration conditions to account for the variety of contact scenarios that can occur. In the case of very light objects, which are appropriately held lightly during the Lift and Hold phase, the object slips between the robot's fingers, so the hand does not experience significant vibrations. In the case of heavy objects, which are held firmly, no Slip signal is detected since the object is held firmly enough to prevent slip, but the impact vibrations are easily apparent in the $\widetilde{a}_h$ signal. For many objects between these two extremes, both Slip and Vibration conditions are often true at contact.

E. Unload

The Unload phase is entered automatically after the held object contacts the target surface. The goal of this phase is simply to let go, but our controller performs this unloading gradually to avoid abrupt transitions. The desired grip force is linearly reduced to zero over a set period of time using:

$$F_{g,des} = F_c - F_c \frac{t - t_s}{\text{TUNLOAD}} \tag{14}$$

Here, t represents the present time, and t_s represents the starting time for the Unload state. TUNLOAD is a constant that determines the unloading duration. The state is exited when $F_{g,des}$ reaches zero, which occurs when $t - t_s == \text{TUNLOAD}$.

F. Open

Once the robot has released the object on the surface, it proceeds to open the gripper. This movement is accomplished by the position controller, using a constant positive desired velocity of VOPEN.

VI. Experimental Validation

We carried out focused experiments to test the performance of the two most novel aspects of the robotic grasp controller described in this paper: the grip force chosen during Load and responses to Slip during Lift and Hold. To understand how our approach compares to more simplistic grasping solutions, we then conducted a more general test of the PR2's capabilities using a collection of fifty everyday household objects.

A. Grip Force Estimation During Loading

As described in Section V-B, the controller in the Load phase selects the target grip force based on the maximum force measured during the gripper's initial contact with the object, normalized by contact speed. We evaluated this technique through grasp testing on eight different objects: a paper cup, a paperboard tea box, a ripe banana, an empty soda can, a raw chicken egg, a tennis ball, a glass bottle, and a full soda can.

We began the experiment by obtaining ground-truth measurements of the minimum grip force necessary for the PR2 to lift each object. These tests were done by placing each object in a known location and orientation on a table. The robot then

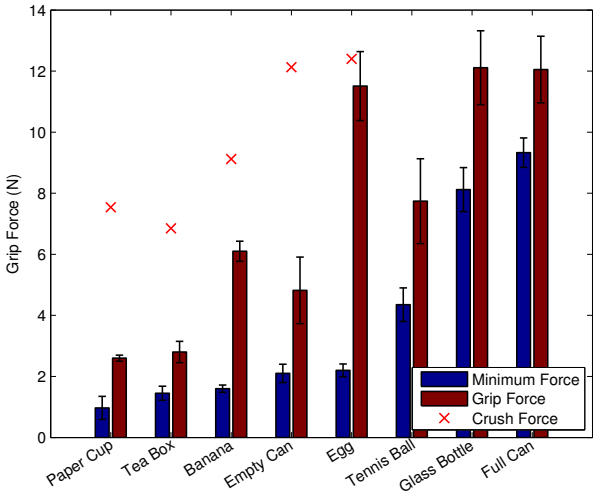

Fig. 5. The target grip force chosen by our Load method when grasping eight everyday objects. The gain KHARDNESS was empirically tuned to yield a grip force (red bar) that is consistently above the minimum grip force necessary to lift the object (blue bar). This calculation provides a good estimate for a large range of objects, but one can see its tendency to overestimate the force necessary to hold objects that are both hard and light, such as the egg. The red $\times$ symbol marks the crushing force for all objects that can be crushed by the robot gripper.

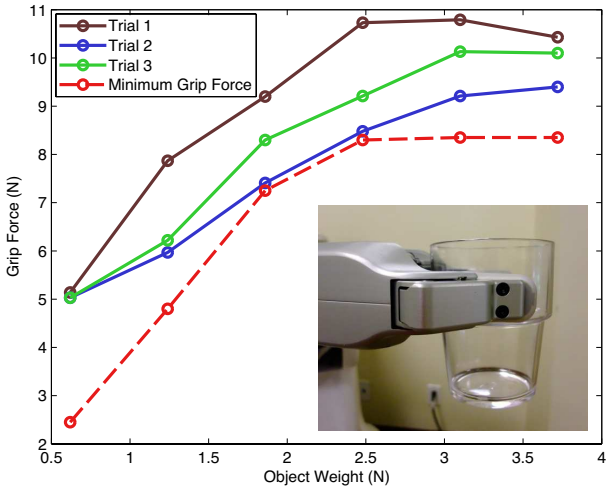

Fig. 6. Slip test results for three different trials. The glass cup was repeatedly filled with marbles to promote slip at a variety of grip force levels. The ground truth data (red dashed line) indicates the minimum grip force needed to prevent slip. As seen here, our Lift and Hold controller has been designed to grasp objects more tightly when it detects slip. This behavior reduces the likelihood of dropped an object without requiring unnecessarily high grasp forces.

closed its gripper on the object using only the force controller described in equations (5)–(7), with the desired force $F_{g,des}$ set to a small value (starting between 0.5 N and 6 N, depending on the object.) The robot then used its arm joints to move the gripper up by 10 cm. The experimenter visually monitored the translational slip between the object and the gripper during the lift. If more than 5 mm of slip occurred, the trial was repeated with the grasp force incremented by 0.1 N. If the object did not slip, the desired grip force was recorded as the minimum grip force needed for lifting. This entire process was repeated eight times per object. The blue "Minimum Force" bars in Fig. 5 show the mean and standard deviation of the eight ground truth measurements for each of the eight objects.

The experiment was then repeated using the grasp controller described in this paper. We performed eight trials with each of the same eight objects, located in the same position and orientation as before. For each trial, the desired loading force F_c was recorded, as calculated with (10). The red "Grip Force" bars in Fig. 5 show the mean and standard deviation of the eight grip force levels that the robot chose during the Load phase for each of the eight test objects.

Lastly, we determined the force necessary to crush each object (if crushing was possible) by successively incrementing the force controller's desired grip force by 0.1 N until the object began to deform significantly. Only a single recording was done for the crushing force because this operation damages the object. These crush force measurements appear as red X's with the other results in Fig. 5. In all cases, our controller chose a grip force above the minimum level needed to avoid slip. For crushable objects, it chose grip forces well below the crush limit for all objects except the egg, which it crushed in three of the eight trials. Subsequent informal testing revealed that a slight reduction in the KHARDNESS gain allows the system to pick up eggs without crushing them, as demonstrated in the movie that accompanies this paper.

B. Slip Response During Lift and Hold

We conducted a separate experiment to test slip compensation in the Lift and Hold phase. As described in Section V-C, Lift and Hold uses the force controller to try to maintain a constant target grasp force; it watches for Slip events, which are derived from the pressure transducer data, and it responds by increasing the target grasp force by a small percentage.

We sought to understand our system's slip response by having the gripper hold a smooth straight-sided object that periodically increased in weight. At the start of this experiment, the cylindrical section of a glass cup was placed in the robot gripper, as seen in the inset of Fig. 6. The weight of the cup was measured to be 0.6 N, and it was oriented vertically. The experimenter began a trial by activating the Lift and Hold controller with an initial desired grip force of 5 N. Batches of fifteen marbles (about 0.6 N per batch) were then added to the cup at intervals of three seconds. The gripper was lightly shaken for two seconds after each batch of marbles was added, during which time the controller reacted to any detected Slip events. The final selected grip force value was recorded in software before the experimenter added another batch of marbles. The marbles were added five times to give the cup a final weight of approximately 3.7 N. This procedure was repeated three times to produce the data shown in the solid traces of Fig. 6.

This test's ground truth data was obtained for each of the six cup weights using the force controller of Section IV-D. The controller's desired grip force was started at 1.0 N. After the cup was grasped by the robot, the experimenter lightly shook the gripper to emulate the slight disturbances that occur during arm motion. If the cup fell out of the gripper or slipped more than 5 mm after two seconds of shaking, the trial was repeated with a grasp force incremented by 0.1 N. The grasp force needed to hold the cup at each of the six weights is shown by the red dashed "Minimum Grip Force" line in Fig. 6. One can see that this value increases up to an object weight of about 2.5 N and then levels off at approximately 8.3 N.

Fig. 7. The 50 objects used in our robustness test. These objects were chosen for no specific reason except that they have a wide range of properties including; hard/soft, heavy/light, sticky/slippery, brittle/elastic.

The controller always chose a grip force value above the level needed to prevent slip, which helps validate our human-inspired approach to gripping unknown objects. The variation between the three trials is primarily due to differences in how the experimenter shook the gripper; stronger external disturbances cause more corrective actions and higher grip force levels.

C. Grasping Robustness

Beyond testing specific aspects of our controller's performance, we wanted to understand how the methods proposed in this paper would work on their intended subject, everyday real-world objects. We thus gathered the collection of 50 common objects shown in Fig. 7, purposefully seeking out items that could be challenging for a robot to grasp. The only requirement on these objects is that they are all within the robot's perception and manipulation capabilities (not too large to grasp, not too heavy to pick up, etc.). The objects included in the collection are as follows: apple, banana (rotten), Band-Aid box, beer bottle (empty), beer bottle (full), can of Spam, can of peas, candle, cereal box (empty), Coffee-mate bottle, duct tape roll, foam ball, gum container, Jello cup, juice box, large plastic bowl, magic marker, masking tape roll, medicine bottle, milk carton (empty), office tape dispenser (heavy), ointment tube (full), peach (soft), plastic juice bottle (empty), plastic juice bottle (full), plum (rotten), rubber football, Solo plastic cup, Saran wrap box, ShiKai shampoo bottle (empty), small wooden bowl, soap bottle (empty), soap box (full), soda can (empty), soda can (full), soup can (full), stress ball, stuffed bear, stuffed elephant, Suave shampoo bottle (empty), sunglasses case, tea box (metal), tea box (paperboard), tennis ball, thin plastic cup, Tide bottle (full), towel, Vasoline container (full), water bottle (empty), and wood plank.

The robot's task for this experiment was to pick up each object from a table and set it down in a different location. We used the object perception and motion planning code of Hsiao et al. to enable the PR2 to accomplish this task autonomously [10]. Starting object poses on the table were hand-chosen to ensure grasp feasibility, and the grasp selection and reactive grasping components of [10] were used to set up stable grasps, with the object centered within the grasp before starting, as per our assumptions for the Close phase. After grasping the object, the robot lifted the object off the table, then moved it from above the table to the side of the robot, so that the arm would not interfere with perception for object placement. The object was then moved to the opposite side of the table for placement. This motion was planned using a randomized joint-angle planner and typically involved a great deal of object rotation, as is typical of many complex pick-and-place operations. During such motions, the robot would ideally prevent the object from rotating or slipping out of the grasp, while continuing to avoid crushing the object.

TABLE I
OUTCOMES OF GRASP TESTING WITH FIFTY EVERYDAY OBJECTS.

	100% Motor Effort	Our Methods
Crushed	100%	3.3%
Rotated Within Grasp	10%	18%
Slipped Within Grasp	4%	8%
Dropped	4%	8%

Each object was tested under two grasp control conditions. The first condition was the original manipulation code described in [10], which takes a naive approach by always closing the gripper with 100% motor effort. The second condition was a portion of the human-inspired robotic grasp controller described in this paper, which included our Close, Load, and Lift and Hold phases. Unfortunately, our Replace, Unload, and Open phases could not be included in this experiment due to code integration difficulties. Our controller was tested first because the naive controller tends to damage crushable objects.

During testing, the experimenter presented the objects to the robot one by one and manually recorded the outcome of each trial. As tabulated in Table I, four different types of errors occurred: the robot might crush a crushable object (30 of the 50 objects were crushable), it might let the object rotate significantly within its grasp (more than $\sim 10°$), it might let the object slip significantly within its grasp (more than ~ 3 cm translation), and it might drop the object either by failing to raise it off the table at the start of the Lift and Hold phase, or by allowing it to slip from its grasp later in Lift and Hold.

The naive controller was found to crush all thirty of the crushable objects, while our controller crushed only one (the rubber football). This drastic improvement in handling delicate objects is balanced by a higher incidence of objects that rotate, slip, and/or drop; our controller commits all three of these errors about twice as often as the naive controller. The two most challenging items were the rubber football and the full Tide bottle; both controllers allowed them to slip within the grasp and fall on the table. This was due to the large size of both objects: the football was larger than the maximum gripper diameter, and the Tide bottle was within 1 cm of the maximum diameter.

Quantitative data was also recorded for our controller's grasp of each of the 50 objects. Fig. 8 presents a histogram of the grip force chosen during the Load phase for the 49 objects that were successfully lifted. These values range from 2.5 N to 27.5 N with most objects below 7.5 N. The objects with the lowest initial grasp force were the towel, the Coffee-mate bottle, the large plastic bowl, and the stress ball, in ascending order. The objects with the highest initial grasp force were the wood plank, the duct tape roll, and the sunglasses case, in descending order. From this we observe that soft objects

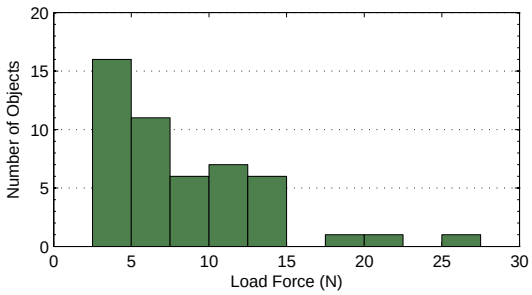

Fig. 8. Histogram for load force in object marathon.

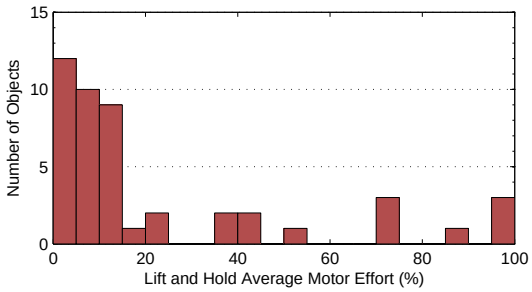

Fig. 9. Histogram for motor effort in object marathon.

generally receive lower initial grip forces than hard objects, as one would expect from the design of our controller.

Fig. 9 provides a histogram of the average motor effort required during Lift and Hold for the 46 objects that were not dropped. One can quickly see that our controller is much more efficient than the naive controller, which always uses 100% motor effort. Because the PR2 gripper has high friction in its mechanism and also has compliant fingertips, it can securely hold objects such as the foam ball and the Band-Aid box with zero motor effort. At the other end of the spectrum, the wood plank, the duct tape roll, and the full beer bottle all had an average motor effort of 100%.

VII. Conclusion

This article introduced a set of sensory signals and control approaches that we hope will help to standardize future approaches to robotic grasping by virtue of their demonstrated usefulness, biologically-inspired origins, and simplicity. We presented a framework that highlights how tactile feedback can be used as a primary means of completing a manipulation task, with encouraging results. While it is clear to the authors that not all tasks can completed solely via tactile information, we feel that it is a promising and underutilized tool in mobile manipulation systems.

In future work we hope to add additional sensing modalities into our object handling framework, including estimates of the necessary object grip force from visual and laser recognition, audio feedback about crushing and damaging objects, weight sensing after an object has been lifted off the table, grasp disturbance prediction from arm motion data, and grasp quality information based on the fingerpad contacts. Our current estimation of the initial grip force necessary to lift an object is solely dependent on the hardness information gleaned during contact. While it has shown to be a strong indicator for many everyday objects, it does have certain failure cases where hardness information is deceptive, such as soft but heavy objects (e.g., a heavy trash bag or stuffed animal), and light but hard objects (e.g., an egg or thin wine glass). We believe that supplementing this information with additional object data will lead to a superior grip force estimator.

Our signal for detecting slip information was successful, but the force response to slip events could be improved by using a gripper with superior dynamic response. The PR2 gripper is admittedly inferior to several of the compliant and well-modeled designs existing in the literature; we hypothesize that these methods would be even more successful if implemented with these alternative research systems. Furthermore, Takahashi et al. [15] have shown that it is possible to obtain useful centroid-of-contact information with spherical fingertips. This is a feature we were unable to reproduce with the PR2's flat fingertips, but we may attempt to redesign the fingertip shape in the future if it proves highly beneficial. The addition of shear-force sensing capabilities to the fingertip may also prove an important indicator of slip information, and it is a close parallel to an important mechanoreceptor of the hand we do not discuss, the SA-II channel, which detects skin stretch.

As we continue to refine this system and increase the range of objects it can handle, all relevant code is freely available at https://code.ros.org/.

TABLE II
Values chosen for controller constants

ATHRESH	4.2 m/s^2	KHARDNESS	0.027 m/s
*DLIMIT	0.02 N	*KP	20,000 N/m
*EFRICTION	7.0 N	KSLIP	1.08
FBPTHRESH	0.25 N	SLIPTHRESH	0.01
*FLIMIT	0.75 N	*TSETTLE	0.05 s
FTHRESH	0.15 N	TUNLOAD	0.20 s
*KD	5,000 Ns/m	VCLOSE	0.04 m/s
*KFCLOSE	0.0013 m/Ns	VOPEN	0.05 m/s
*KFOPEN	0.0008 m/Ns	*VTHRESH	0.001 m/s

Appendix
Controller Constants

To facilitate a generic presentation of our grasp controller, the mathematical constants used in this paper are designated with an all-capitalized naming convention. Table II shows the values and units that our controller actually employs for these constants. The * symbol indicates values that are robot specific.

Acknowledgments

The authors thank Matei Ciocarlie for his help in adding the described grasp controller to the PR2 object manipulation framework, and they thank Derek King for his assistance in troubleshooting the PR2 accelerometer signals.

References

[1] C. C. Kemp, A. Edsinger, and E. Torres-Jara, "Challenges for robot manipulation in human environments: Developing robots that perform useful work in everyday settings," *IEEE Robotics and Automation Magazine*, pp. 20–29, March 2007.

[2] R. B. Rusu, A. Holzbach, R. Diankov, G. Bradski, and M. Beetz, "Perception for mobile manipulation and grasping using active stereo," in *Proc. IEEE-RAS International Conference on Humanoid Robots*, Paris, Dec. 2009.

[3] A. Saxena, J. Driemeyer, and A. Y. Ng, "Robotic grasping of novel objects using vision," *International Journal of Robotics Research*, vol. 27, no. 2, pp. 157–173, Feb. 2008.

[4] R. S. Johansson and J. R. Flanagan, "Coding and use of tactile signals from the fingertips in object manipulation tasks," *Nature Reviews Neuroscience*, vol. 10, pp. 345–359, May 2009.

[5] M. A. Srinivasan, J. M. Whitehouse, and R. H. LaMotte, "Tactile detection of slip: Surface microgeometry and peripheral neural codes," *Journal of Neurophysiology*, vol. 63, pp. 1323–1332, 1990.

[6] M. R. Cutkosky, robert D. Howe, and W. R. Provancher, *Springer Handbook of Robotics*, B. Siciliano and O. Khatib, Eds. Springer, 2008.

[7] R. S. Dahiya, M. Gori, G. Metta, and G. Sandini, "Better manipulation with human inspired tactile sensing," in *Proc. of Robotics Science and Systems (RSS 2009) Workshop on Understanding the Human Hand for Advancing Robotic Manipulation*, 2009, pp. 36–37.

[8] A. Schmitz, M. Maggiali, M. Randazzo, L. Natale, and G. Metta, "A prototype fingertip with high spatial resolution pressure sensing for the robot iCub," in *Proc. IEEE-RAS International Conference on Humanoid Robots*, December 2008, pp. 423–428.

[9] R. D. Howe, N. Popp, P. Akella, I. Kao, and M. R. Cutkosky, "Grasping, manipulation, and control with tactile sensing," in *International Conference on Robotics and Automation*, 1990.

[10] K. Hsiao, S. Chitta, M. Ciocarlie, and E. G. Jones, "Contact-reactive grasping of objects with partial shape information," in *IROS*, 2010.

[11] L. Natale and E. Torres-Jara, "A sensitive approach to grasping," in *Proc. International Conference on Epigenetic Robots*, Paris, France, September 2006.

[12] A. Jain and C. C. Kemp, "El-e: an assistive mobile manipulator that autonomously fetches objects from flat surfaces," *Autonomous Robots*, vol. 28, pp. 45–64, 2010.

[13] A. M. Dollar, L. P. Jentoft, J. H. Gao, and R. D. Howe, "Contact sensing and grasping performance of compliant hands," *Autonomous Robots*, vol. 28, pp. 65–75, 2010.

[14] H. Yussof, M. Ohka, H. Suzuki, N. Morisawa, and J. Takata, "Tactile sensing-based control architecture in multi-fingered arm for object manipulation," *Engineering Letters*, vol. 16, no. 2, 2009.

[15] T. Takahashi, T. Tsuboi, T. Kishida, Y. Kawanami, S. Shimizu, M. Iribe, T. Fukushima, and M. Fujita, "Adaptive grasping by mulit fingered hand with tactile sensor based on robust force and position control," in *International Conference on Robotics and Automation*, 2008.

[16] B. Eberman and J. K. Salisbury, "Application of change detection to dynamic contact sensing," *The International Journal of Robotics Research*, vol. 13, p. 369, 1994.

[17] D. Dornfeld and C. Handy, "Slip detection using acoustic emission signal analysis," in *International Conference on Robotics and Automation*, 1987.

[18] N. Wettels, V. J. Santos, R. S. Johansson, and G. E. Loeb, "Biomimetic tactile sensor array," *Advanced Robotics*, vol. 22, pp. 829–849, 2008.

[19] A. Kis, F. Kovács, and P. Szolgay, "Grasp planning based on fingertip contact forces and torques," in *Proceedings of Eurohaptics*, 2006.

[20] J. M. Romano, S. R. Gray, N. T. Jacobs, and K. J. Kuchenbecker, "Toward tactilely transparent glove: Collocated slip sensing and vibrotactile actuation," in *World Haptics*, 2009.

[21] S. Chitta, M. Piccoli, and J. Sturm, "Tactile object class and internal state recognition," in *International Conference on Robotics and Automation*, 2010.

[22] G. A. Gescheider, J. H. Wright, and R. T. Verrillo, *Information-processing channels in the tactile sensory system: a psychophysical and physiological analysis*. New York: Psychology Press: Taylor & Francis Group, 2009.

[23] K. O. Johnson, "The roles and functions of cutaneous mechanoreceptors," *Current Opinion on Neurobiology*, vol. 11, pp. 455–461, 2001.

[24] S. Allin, Y. Matsuoka, and R. Klatzky, "Measuring just noticeable differences for haptic force feedback: Implications for rehabilitation," in *Proc. IEEE Haptics Symposium*, March 2002, pp. 299–302.

Joseph M. Romano is a Ph.D. student in the department of Mechanical Engineering and Applied Mechanics (MEAM) at the University of Pennsylvania. Joe works in Dr. Kuchenbecker's Haptics Research Group, a part of the GRASP Laboratory, investigating the relationship between kinesthetic large-scale motion information and localized tactile and cutaneous signals to improve human and machine haptic systems. He completed a Masters in MEAM at the University of Pennsylvania in May of 2010, and B.S. degrees in both Computer Science and Engineering Mechanics at the Johns Hopkins University in 2007.

Kaijen Hsiao is a research scientist at Willow Garage, Inc. Her research interests lie in grasping and manipulation, particularly with respect to dealing with uncertainty, incorporating information from diverse sensors, and adding robustness to robot capabilities. Dr. Hsiao received her Ph.D. in Computer Science from MIT in 2009, and she has been at Willow Garage since then.

Günter Niemeyer is a senior research scientist at Willow Garage Inc. and a consulting professor of Mechanical Engineering at Stanford University. His research examines physical human-robotic interactions, force sensitivity and feedback, teleoperation with and without communication delays, and haptic interfaces. This involves efforts ranging from real-time motor and robot control to user interface design. Dr. Niemeyer received his M.S. and Ph.D. from MIT in the areas of adaptive robot control and bilateral teleoperation, introducing the concept of wave variables. He also held a postdoctoral research position at MIT developing surgical robotics. In 1997 he joined Intuitive Surgical Inc., where he helped create the da Vinci Minimally Invasive Surgical System. He joined the Stanford faculty in the Fall of 2001, directing the Telerobotics Lab and teaching dynamics, controls, and telerobotics. He has been a member of the Willow Garage research group since 2009.

Sachin Chitta is a research scientist at Willow Garage, Inc. His research interests are in motion planning, control, and sensing for mobile manipulation with a particular interest in making robots useful in unstructured indoor environments. Prior to coming to Willow Garage, Dr. Chitta received his Ph.D. in Mechanical Engineering from the University of Pennsylvania in 2005, and he was a postdoctoral researcher working on learning for quadruped locomotion and modular robotics in the GRASP Lab at Penn from 2005 to 2007. At Willow Garage, he has worked on motion planning and control for mobile manipulation tasks like door opening, cart pushing and object retrieval and tactile sensing for reactive grasping and object state recognition with the PR2 mobile manipulation platform.

Katherine J. Kuchenbecker is the Skirkanich Assistant Professor of Innovation in Mechanical Engineering and Applied Mechanics at the University of Pennsylvania. Her research centers on the design and control of haptic interfaces, and she directs the Penn Haptics Group, which is part of the GRASP Robotics Lab. Dr. Kuchenbecker serves on the program committee for the IEEE Haptics Symposium, and she has won several awards for her research, including an NSF CAREER award in 2009. Prior to becoming a professor, she was a postdoctoral fellow at the Johns Hopkins University, and she earned her Ph.D. in Mechanical Engineering from Stanford University in 2006.