

**Merging Local and Global 3D Perception for Robotic
Grasping and Manipulation**

by

Rebecca E. Cox

B.S. Computer Engineering, Missouri University of Science and
Technology, 2014

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirements for the degree of
Master's of Science
Department of Electrical, Computer, and Energy Engineering
2017

This thesis entitled:
Merging Local and Global 3D Perception for Robotic Grasping and Manipulation
written by Rebecca E. Cox
has been approved for the Department of Electrical, Computer, and Energy Engineering

Prof. Nikolaus Correll

Prof. Christoffer Heckman

Prof. Daniel Szafir

Date _____

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Cox, Rebecca E. (M.S. ECEE)

Merging Local and Global 3D Perception for Robotic Grasping and Manipulation

Thesis directed by Prof. Nikolaus Correll

This paper presents our results towards fusing RGB-D images with data from contact and proximity sensors embedded in a robotic hand for improved object perception, recognition and manipulation. Optical depth information from multiple sensors is often inaccurate and inconsistent. These problems arise from problems with sensor calibration, but also occlusion of objects by other objects or the robot arm itself. In this paper, we propose to combine global pose information from RGB-D sensing with local proximity sensing during approach. Here, we use contact information based on a novel contact sensor and additional pose information provided by the arm's pose.

Dedication

I would like to dedicate this thesis to my dad, Jeffrey Cox. Without his guidance and support all my life, I would not be where I am.

Acknowledgements

This work has been made possible by many people, as I could not have done it alone. Foundational work that this code was built on top of includes work by Jorge Cañardo Alastuey, Andy McEvoy, Dave Coleman, and others. My teammates, Branden Romero and Radhen Patel, who competed with me in the IROS Grasping and Manipulation Challenge and continued to make me feel as though I was part of a team long after the competition.

Most importantly, I would like to thank my advisor, Nikolaus Correll, who gave me this opportunity and guided my research throughout. Being unfamiliar with robotics, I was pretty clueless how anything worked in the lab, especially robotic arms. Nikolaus was supportive, patient, and assisted me with understanding basic concepts.

Lastly, I would like to thank Sandia National Laboratories for funding and supporting me. My skills as a researcher, presenter, and writer would not be where they are without this experience and pursuing a Master's would not have been financially possible without them.

Contents

Chapter	
1 Introduction	1
2 Global Perception	4
2.1 Point Cloud Segmentation	5
2.2 Object Recognition	6
2.3 Object Tracking	7
2.4 Published Object Poses	7
3 Local Perception	8
3.1 Proximity Measurement	8
3.2 Contact Detection	9
3.3 3D Reconstruction using Proximity	11
4 Merging Perceptions	13
4.1 Scene Calibration	14
4.1.1 Experimental Validation	15
4.1.2 Future Work	16
4.2 Merging Point Clouds	17
4.3 Feature Detection	17
4.3.1 Experimental Validation	18

4.3.2 Future Work 21

4.4 Occluded Object Detection 21

4.5 Conclusion 22

Bibliography **24**

Appendix

A Algorithms **29**

Tables

Table

4.1	Scene calibration results using contact sensing to verify cube centroids found from the camera	16
4.2	We show our results for estimating the pose of the cup from an RGB-D camera alone and from using proximity measurements in the Baxter parallel grippers.	20

Figures

Figure

- | | | |
|-----|---|----|
| 1.1 | A typical system setup combining multiple sensor streams. The fingers on the Jaco arm from Kinova contain proximity and force sensors. An Asus Xtion is used for depth and RGB imaging. | 2 |
| 2.1 | RANSAC methods visualized on 2D points for finding a line. When RANSAC converges to a solution, the inliers in blue are separated from the outliers in red in the right image. | 5 |
| 2.2 | Left: Visualization of the surface normals of a cup. Right: Similarity between descriptor histograms of model cup object and the new cup object being matched. . | 6 |
| 3.1 | Shown as the black rectangles embedded in the translucent PDMS, the sensors emit infrared (red signal) and measure the reflection. This allows us to measure three important senses; proximity to an object, contact with an object, and force applied. | 9 |
| 3.2 | Sensor response (analog reading) from the SA-I (blue) and FA-I (pink) channel versus time(s) to a spoon handle when brought close to, touched and separated from the PCF-sensors on the finger of Jaco arm. | 10 |

- 3.3 Left: Scanning a staircase of blocks using the finger sensors on the custom designed Baxter gripper. Right: 3D point cloud from contact detection in the array of finger sensors compared with 3D point cloud from Asus Xtion depth sensor. The finger sensors provide precise information where the surface of the object is relative to the grippers. The offset in this badly calibrated scene can be clearly seen by the far distance ($5cm$) between these two point clouds. 11
- 3.4 Left: Best fit function for the measured sensor values for set distances away from the cup. c , a base value that is unique to each sensor, was calculated by taking the average of ten sensor readings when no object was present. Middle: Proximity sensors embedded in the Baxter parallel grippers are being rotated around the cup to capture a full scan. The proximity measurements are actively generating a new point cloud. Right: 3D Model generated from proximity measurements. Pink and cyan points are from the left and right fingers. 12
- 4.1 The above images demonstrate the typical scene calibration error and how contact sensing is used to correct it. Left: Typical offset between the robot's end-effector and the RGB-D point cloud from the camera. The blue point clouds are the cubes shown in the middle image. The pink and cyan points show the point cloud generated from proximity sensing from the left and right fingers. Middle: This image shows the real-world view that is being reconstructed in the left and right images. Right: After applying a rigid transformation calculated from touching cubes on the table, the RGB-D point cloud aligns with the robot's grippers. 13

- 4.2 All three images feature the same cup from the YCB data set with the handle on the right side. Left: Segmented point cloud from the RGB-D camera. Middle: Proximity sensors embedded in the Baxter parallel grippers are being rotated around the cup to capture a full scan. The proximity measurements are actively generating a new point cloud. Right: Merged point clouds of the cup from the depth camera and from the 3D finger sensor scan. Pink points are from the left finger and cyan points are from the right finger. The handle of the cup is clearly shown from the scan, as well as the accurate alignment of the merged point clouds. 17
- 4.3 Missing features in the RGB-D point cloud are found using the proximity sensors embedded in the robotic hand. Even when in full view of the camera, the cup's handle is often not discernible in the point cloud. 19
- 4.4 Exploring areas where no global perception information is available. 22

Chapter 1

Introduction

Despite the huge progress over the past few decades in robotics, even completing the simplest tasks remain a challenge for autonomous grasping and manipulation. In ideal sceneries where exact object poses are known, intelligent grasp planners are effective and perform nearly as well as a human hand. However, perception is a necessary component for determining precise object locations that continues to be a challenge. Both 3D perception and tactile sensing have been explored, and both might be sufficient means on their own and complement each other [44]. This is akin to grasping and manipulation in humans [47].

Accuracy in object recognition depends heavily on the accuracy of calibration, successful point cloud segmentation, lighting conditions, and whether objects are occluded or not. Tactile sensing [24] can be much more accurate as the kinematics of a robot arm are usually well known and its encoders are precise, but requires active exploration of the environment.

In this paper, we describe our results from fusing 3D perception, in-hand proximity sensing, and touch into a single representation. Here, the key ideas are that tactile sensing can provide data-rich 3D representations with accuracy and proximity sensing can seamlessly bridge between the two sensing modalities, allowing us to gather data without disturbing the object’s pose.

We are combining an Asus Xtion depth sensor with Robotic Materials’ tactile sensors for the Kinova three-fingered hand and the parallel gripper for Rethink Robotics Baxter. We use the Point Cloud Library (PCL) to perform object recognition from the raw point cloud and find the approximate location of the objects relative to the arm. Our experimental setup is shown in




Figure 1.1: A typical system setup combining multiple sensor streams. The fingers on the Jaco arm from Kinova contain proximity and force sensors. An Asus Xtion is used for depth and RGB imaging.

Figures 1.1. Small errors in the camera-to-hand transformation result in a typical object location offset around 3 cm. We show how we can use tactile sensing as a feedback mechanism to refine the camera-to-hand transformation for more accurate object location estimation.

Having accurate location information increases grasping performance, however, occluded features or objects such as an unseen handle of a cup can disrupt an attempted grasp if the gripper collides with them. We show how the proximity measurement in these sensors can be used to scan objects without disturbing the environment to fill in missing points in the RGB-D point cloud. Lastly, we show how the tactile sensors can be used to replace the RGB-D camera entirely through active scene exploration to locate objects, recognize them, and determine grasping poses.

Using tactile sensing in perception has been widely explored, for example to plan motions to learn more about the geometry of an object or match it against a known 3D model [12, 24, 32, 48]. Incorporating this closed loop feedback can correct or fill in missing information from the low-cost depth cameras commonly used in research. However, tactile sensing requires touch which often unintentionally moves the object and changes the pose, requiring a new grasping pose to

be determined. For this reason, proximity sensing using IR sensors in the fingers has also been explored for improving the pre-grasp pose prior to contact [25, 31, 54]. [35] explores using electric field sensors can be used for proximity, although this is susceptible to object composition. Even less have used the sensor data as a feedback loop to update the locations of objects [31, 50]. Little attention has been devoted on fusing proximity and depth information obtained from different means using tactile sensing to provide a common reference frame. Such refined models could then be used to improve grasp planning.

We had the opportunity to test some of the early stages of this work by competing in the 2016 IROS Robotic Arm Grasping and Manipulation held in South Korea. In preparation for the fully autonomous track, we focused on merging information from the camera with proximity sensing in the fingers to successfully complete tasks. To us, full autonomy meant that the robot should perceive its environment and make decisions based on that to complete tasks. However, due to the competition rules allowing teams to place their objects where they wanted, most competing teams did not use any perception, meaning that these teams focused only on the control side and could use simple hard coded routines to complete tasks.

Although we were disappointed that any work done towards true autonomy was overshadowed by these hard-coded routines, it demonstrates how challenging even the simplest perception problems continues to be. We explore in this research how simple, low-cost sensors can complement traditional RGB-D cameras to provide more robust perception.

Chapter 2

Global Perception

The system setup allows the RGB-D camera to be placed in an arbitrary location around the robot. To find the camera-to-robot transform, an augmented reality (AR) tag is rigidly mounted to the wrist of the robot. When visible to the camera, the system can estimate the transformation using the known location of the robot's wrist. Locating the AR tag is done through the widely used AR tracker ROS package. This leaves a typical scene calibration error around 3cm that will be corrected in a later section. Objects that are visible to the camera can now be related to a location reachable by the arm.

To provide simple object detection, we built a perception pipeline using PCL that takes a raw point cloud and RGB image as input and publishes found objects and their poses. Using PCL's built in feature detection libraries, we are able to successfully perform object labeling on individual objects, but don't always have the information available in a single frame to be able to estimate object pose accurately, especially in a cluttered environment with occluded views. In such cases, the center-of-mass of the point cloud is not congruent with the actual center, making grasping difficult. Worse, key geometries that are important for grasping and manipulation are hidden and might not be inferred from partial models.

While the approach described here is basic, we believe its challenges and limitations to be ubiquitous in manipulation, even when using more advanced methods for perception.

2.1 Point Cloud Segmentation

The first part of the pipeline involves removing noise and narrowing down the point cloud to regions that contain the objects. We can assume that the objects will always be placed on a table or flat surface that will fill a large portion of the field of view in the camera. For this reason, the initial step is to find the flat surface and extract it from the point cloud using Random Sample Consensus (RANSAC), a simple method of separating inliers of the 2D plane from everything else in the cloud.

In Figure 2.1, we show how RANSAC finds a best fit line from 2D points. As the name implies, RANSAC takes a random set of points to represent the shape being fit (2 points for a line, 3 points for a plane), draws the shape through them, and counts the inliers and outliers. It can be assumed that points that are not along the shape will vote inconsistently. We include inliers to be within 2cm of our plane. We can also take advantage of the fact that the table is static in the point cloud to decrease computation time between frames. Once the table has been found once, future RANSAC algorithms will be provided with a hypothesis of this found orientation.




Figure 2.1: RANSAC methods visualized on 2D points for finding a line. When RANSAC converges to a solution, the inliers in blue are separated from the outliers in red in the right image.

The remaining point cloud is clustered into smaller objects using a nearest neighbor approach. This is calculated by taking a sphere of radius r_s around every point in the cloud and grouping points contained in the same sphere. Objects that are placed further than the distance r_s from each other will be in different clusters and will get processed separately. Objects that are stacked or

within the spherical radius become clustered into a single object and require further segmentation techniques such as color-based segmentation.

2.2 Object Recognition

After the objects have been segmented, we individually analyze each one and compare it against the known data set. Due to the redundancy in point cloud information of a single object, the objects are down-sampled to their keypoints. These keypoints are sparse enough to make processing much more efficient while still containing enough information for surface reconstruction. 3D descriptors from [53] are calculated from the surface normals of these keypoints and used in searching for correspondence with an object from the data set, shown in Figure 2.2. When enough matching descriptors are found, the object is labeled and a pose is estimated based on the translations from matching descriptors. This approach is described in more detail in [44].

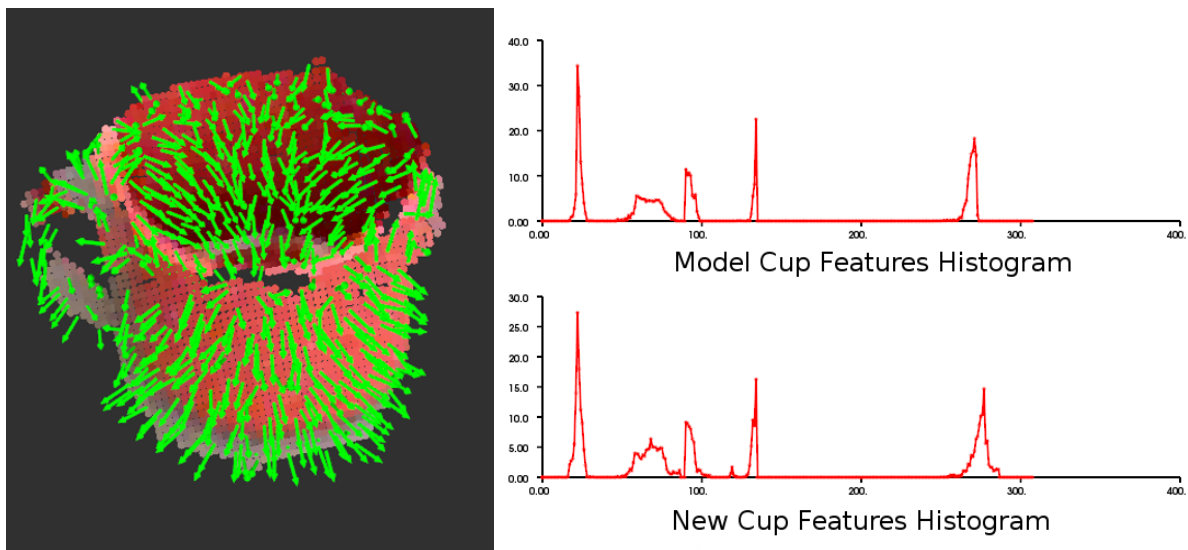


Figure 2.2: Left: Visualization of the surface normals of a cup. Right: Similarity between descriptor histograms of model cup object and the new cup object being matched.

2.3 Object Tracking

The final step in the perception pipeline aims to improve object recognition during run time by decreasing the likelihood that an object found previously will be lost due to two main reasons: objects being moved or occluded by a robotic arm during manipulation and noise in a point cloud that results in not enough correspondences to be found that were observed previously. When an object is found, we record a geometric centroid and a timestamp. In the likely case that in the next few seconds an unidentified object is found to have a near centroid, but failed correspondence matching, then it is assumed to be the same object as identified before. We found that correctly identified objects often lost their label every few frames, even without moving the object. Small fluctuations in the received point cloud and down-sampling reduce the number of matched descriptors found. Similarly, descriptor matching can fail due to the robotic arm partially occluding the view of an object during manipulation. Keeping track of found objects in this way greatly reduced false negatives seen.

2.4 Published Object Poses

Lastly, object poses and their corresponding labels are published to ROS for grasping and manipulation. The pose consists of the centroid of the object, calculated from the center of mass of the points, and the orientation, giving a default orientation of the identity matrix (no orientation).

Chapter 3

Local Perception

Local perception occurs within a short distance from the hand, hence the name local. We use combined proximity, contact and force sensors described in [43] to get proximity measurements from the robotic hand to an object with a range of up to 5 cm, visualized in Figure 3.1. The sensors are available commercially¹ and easy to integrate with existing hardware. They consist of a digital infrared distance sensor that emits infrared and measures proximity. For protection from abrasion and to allow light to pass through, they are embedded in a soft transparent polymer, which doubles as a spring for force measurements. Due to an inflection point in the sensor signal at contact, contact can be robustly estimated independently of the surface reflectivity. This information is used to improve gripper alignment during grasping and determine when contact with objects is made, which we describe in [41].

3.1 Proximity Measurement

The sensors consist of an infrared emitter and ambient light sensor chip that allows a variety of parameters to be tuned, such as the emitter current and the carrier frequency. The higher the emitter current, the more resolution achieved for objects farther away. However, as objects move closer, the ambient light sensor may become over saturated, so current values over 100mA were avoided. Configuring the emitters to different carrier frequencies helped prevent cross-talk. The ability to select the frequency of each sensor enables arranging sensors in opposite pairs, such as

¹ www.roboticmaterials.com




Figure 3.1: Shown as the black rectangles embedded in the translucent PDMS, the sensors emit infrared (red signal) and measure the reflection. This allows us to measure three important senses; proximity to an object, contact with an object, and force applied.

required on a robotic gripper, without interference.

Reflected light is strongly non-linear as proximity decreases and is highly dependent on the surface properties such as color, texture, and angle of incidence. Previous research discussed in [43], recorded noisy sensor values. We will show in Section 3.2 that calibrating for surface properties and subtracting out a base sensor value, the sensors are in fact not noisy and can be used for more accurate local perception.

3.2 Contact Detection

Unlike proximity measurements, detecting contact using these sensors is not affected by surface properties due to an inflection point that is observed in the signal [43]. Regardless of the initial amount of reflected signal, as proximity distance between the sensors and object decreases, the nonlinear increase in reflection experiences a peak rate of change, see Figure 3.2. This real-time feedback has already shown to be beneficial in our experiments during grasping by providing immediate feedback of the object’s location relative to the fingers when contact occurs. In previous work, this contact event was used for object detection for improving grasping and 3D reconstruction of where contact was observed. For example, thin objects in a cup such as a straw or spoon




Figure 3.2: Sensor response (analog reading) from the SA-I (blue) and FA-I (pink) channel versus time(s) to a spoon handle when brought close to, touched and separated from the PCF-sensors on the finger of Jaco arm.

are detected from Jaco’s three fingered hand using a search motion above the cup. This is often necessary in grasping to find these narrow objects for the reason that they are typically not detectable by the low-resolution depth cameras due to them not reflecting enough depth information. When contact is observed, the fingers are aligned and closed for the ideal grasp on the object for manipulation. However, this search motion requires that the sensors be moved forward slowly to prevent accidentally bumping that object before it can be grasped.

Unintentionally bumping objects is a common problem in tactile sensing that can be prevented by the combined use of proximity to detect the oncoming contact to slow the end-effector. Initially, 3D models were constructed using only contact sensing in the array of sensors embedded Baxter parallel grippers. This configuration was used for 3D modeling over the Jaco fingers due to the well known location of all seven internal sensors, as well as there being more sensors available for higher resolution models from less motions. The Jaco fingers each have one embedded sensor, where the Baxter grippers have seven internal sensors and one tip sensor that was used to detect the surface of the table. These seven sensors were touched along stacks of blocks shown in Figure 3.3. The 3D model clearly shows the staircase, as well as the offset in the scene calibration. However




Figure 3.3: Left: Scanning a staircase of blocks using the finger sensors on the custom designed Baxter gripper. Right: 3D point cloud from contact detection in the array of finger sensors compared with 3D point cloud from Asus Xtion depth sensor. The finger sensors provide precise information where the surface of the object is relative to the grippers. The offset in this badly calibrated scene can be clearly seen by the far distance ($5cm$) between these two point clouds.

useful, making enough contact points on an object for reconstruction or object recognition is not always practical. For this reason, we explored using proximity measurements for 3D reconstruction and contact detection for optimizing scene calibration discussed below.

3.3 3D Reconstruction using Proximity

The proximity sensors were previously used to generate crude 3D models of objects discussed in [43]. Albeit noisy, surface features of objects could be clearly seen. We hypothesized that the models were noisy due to the sensors being calibrated for properties of a different surface and the non-orthogonal incidence angles at irregular surface features. To test this, the sensors were calibrated for the surface of the cup shown in Figure 3.4 by moving the cup to known distances from the sensors and recording the measured reflection in the sensors. A best fit curve was found and is shown in Figure 3.4. It was found to be $y = a * b^x + c$, where x is the proximity distance in cm , y is the sensor measurement, and c is the base value that is unique to each sensor. The

function can be inverted to:

$$x = \frac{\log(y - c) - \log(a)}{\log(b)} \quad (3.1)$$

In the Baxter parallel gripper shown in the middle image in Figure 3.4, there are seven sensors embedded in the side of each finger directing inwards and a sensor embedded in each fingertip that points downward. The kinematics of Baxter’s fingers are well known, as well as the exact location of each sensor in the finger. This information, as well as the proximity measurements to each sensor, allows us to construct 3D point clouds shown in Figure 3.4. When a sensor measurement is received above the noise floor, Equation 3.1 is used to calculate the location for the new point at an x perpendicular distance from the sensor and add it to the object’s point cloud. In an accurately calibrated scene, we show that we can construct a full 3D model from the sensors that aligns with the RGB-D image. In the experimental section, we demonstrate that the sensors can be used to increase grasping performance by using these 360° scans to find hidden features such as the handle of the cup.






Figure 3.4: Left: Best fit function for the measured sensor values for set distances away from the cup. c , a base value that is unique to each sensor, was calculated by taking the average of ten sensor readings when no object was present. Middle: Proximity sensors embedded in the Baxter parallel grippers are being rotated around the cup to capture a full scan. The proximity measurements are actively generating a new point cloud. Right: 3D Model generated from proximity measurements. Pink and cyan points are from the left and right fingers.

Chapter 4

Merging Perceptions

Discussed in Chapter 3, the point cloud generated from proximity and contact sensing is a similar point cloud to ones produced from an RGB-D camera, minus the color information. When scanning continuous surfaces, we show that the point clouds are nearly identical, like the cylindrical cup shown in Figure 4.2. In this section we demonstrate that we can use these sensors to improve global perception by optimizing the scene calibration transform, accurately estimating an object's pose, and performing active scene exploration in regions not visible to the camera.




Figure 4.1: The above images demonstrate the typical scene calibration error and how contact sensing is used to correct it. Left: Typical offset between the robot's end-effector and the RGB-D point cloud from the camera. The blue point clouds are the cubes shown in the middle image. The pink and cyan points show the point cloud generated from proximity sensing from the left and right fingers. Middle: This image shows the real-world view that is being reconstructed in the left and right images. Right: After applying a rigid transformation calculated from touching cubes on the table, the RGB-D point cloud aligns with the robot's grippers.

4.1 Scene Calibration

As camera calibration is a recurrent problem in robotics, we investigated whether tactile information can be used to correct calibration error of a RGB-D camera. In addition to calibrating the intrinsic non-linear parameters of a camera, camera calibration requires calculating a transformation to find the location of the objects from the camera's viewpoint to a location relative to the robot arm. We use the widely used `AR_tracker` library available as a ROS package to locate the AR tag used for calculating the camera-to-robot transformation. The low resolution of the RGB-D camera and slight errors in AR tag placement result in a typical scene offset of 3cm.

To mitigate these offset errors, we investigate how contact at the end-effector can be used to correct the transformation. We chose to focus on updating the camera pose over the poses of objects due to most of the offset errors propagating from the estimated camera pose and being similar across all detected objects. It can be assumed that the estimated pose of the objects may also have an offset from the camera due to the centroid being derived from the center of mass (COM) which is calculated from the point cloud that will contain at most three sides of the cubes, biasing the object's location towards the more visible sides. However, this offset from using the COM as the object centroid is minor in small objects such as a cube ($0.5cm$) compared to the scene calibration offset ($3cm$).

To conduct scene calibration, centroids of objects from the vision system are compared against centroids found using touch discussed Algorithm 1. The centroids consist of an (X, Y, Z) coordinate location for both the vision and the contact system allowing us to use simple, widely used algorithms to calculate the transformation correction. For example, after finding the actual centroid of an object relative to the hand, there exists a translation from the vision system centroid of the object that can be easily calculated using subtraction. When more centroids are found using touch, an algorithm based on a linear least squares approach using singular value decomposition (SVD) is applied to find a rigid transformation.

4.1.1 Experimental Validation

We showed how touch can be used as a feedback mechanism to improve the camera-to-robot transformation. Conducting scene calibration on either system, Jaco or Baxter, we often experience a 3cm or greater error offset from the centroid calculated by the point cloud COM from vision to the actual centroid relative to the end-effector. This error can be visually seen in the point cloud overlay with the robot model in RVIZ. However, to quantify the offset and measure improvement, we used the Baxter parallel grippers to scan many cubes on a table and find their actual centroid relative to the robot.

Shown in 4.1, wooden cubes $2.5cm^3$ are distributed on a table with enough space in between to allow for the grippers to scan without bumping into other cubes. To keep the scanning trajectory consistent, the cubes were given identical orientations, where the sides align with the X,Y,Z axis of the robot base. Although the cubes' orientations could be estimated using methods available in PCL, slight errors in estimation or noise in the depth image may result in a miscalculated orientation. An incorrect orientation will cause the grippers to move the cube during touch, invalidating the experiment. For this reason, only the orientation shown was used.

To compensate for the scene calibration error, a radius of $3cm$ is scanned around each cube until the cube is found, stepped through in Scene Calibration Algorithm in the appendix. To measure incremental improvements from verifying more cube centroids, we first measured the average offset of every cube recognized by scanning each one. Shown in Table 4.1, the average centroid offset was $3.39cm$, which was on target to the expected offset.

We now have a set of object centroids from vision and a set from contact. When choosing one centroid at random to derive a translation matrix, we saw an average improvement in offset error of $1.48cm$. The average of two translations from two cube centroids saw roughly the same improvement, within a millimeter difference. These improvements were better than expected despite the camera-to-robot transformation often having a large rotational error, as much as 15° around an axis. Shifting the camera's pose could cause only the selected cube to achieve a more accurate

position while simply shifting the other cubes to different offset errors. However, these experiments repeatedly showed improvement from translation alone even when there existed a large rotational error. This may be due to the cubes being placed closely together on the table.

Cube Centroids Verified	Transformation Update Performed	Average Offset of Centroid (cm)
0	N/A	3.39
1	Translation	1.48
2	Averaged Translation	1.36
3	Rigid Transformation	0.54
4	Rigid Transformation	0.54
5	Rigid Transformation	0.49

Table 4.1: Scene calibration results using contact sensing to verify cube centroids found from the camera

To correct a rotational offset, at least three cube centroids are required to solve the linear least squares method based on singular value decomposition (SVD). As shown visually in Figure 4.1, the derived rigid transformation corrected the centroid offsets to within 0.5cm . The remaining offset is likely due to the vision centroids being calculated from the COM of the point clouds which will always be biased towards the sides that are more visible to the camera. This accurate scene calibration allowed us to align the point clouds of the cup shown in Figures 4.2 and 4.3.

4.1.2 Future Work

The current methods required contact with the cubes. If not aligned properly, the cube would be moved. A better way to correct the scene calibration would be to use proximity sensing so that contact is not required. One way, might be do to a scan of an object and use an algorithm like Iterative Closest Point (ICP) to find the best transform to merge the clouds. Although scans

of unknown objects are likely to have an inaccurate 3D model with smoothed edges constructed from proximity sensing, the general size and shape may be enough to use ICP to derive the rigid transformation.




Figure 4.2: All three images feature the same cup from the YCB data set with the handle on the right side. Left: Segmented point cloud from the RGB-D camera. Middle: Proximity sensors embedded in the Baxter parallel grippers are being rotated around the cup to capture a full scan. The proximity measurements are actively generating a new point cloud. Right: Merged point clouds of the cup from the depth camera and from the 3D finger sensor scan. Pink points are from the left finger and cyan points are from the right finger. The handle of the cup is clearly shown from the scan, as well as the accurate alignment of the merged point clouds.

4.2 Merging Point Clouds

We have shown how we generate point clouds for global and local perception. In this section, we work with a camera transform that was found using the previously described methods. This accurate scene calibration enabled us to easily merged the two points clouds visualized together in Figure 4.2.

4.3 Feature Detection

In order for robots to operate in unconstrained environments, they must be able to cope with occluded objects and features of objects that are not visible from global perception. Proximity sensing has a great advantage over tactile sensing in that it does not require contact to observe surface properties. Proximity in the hand can be used to find the most ideal grasp without having to disturb and potentially dislocate the object.

When the surface of an object is parallel to the array of IR sensors in the Baxter grippers, or perpendicular to the sensor signal, an accurate calculation on the sensor measurement can be used to generate a 3D model of the object as shown in 4.2. The cup is an ideal object to construct a 3D model of to search for the handle due to the minimal irregular surface features present. The single irregularity, the handle, is not often identifiable in the point cloud from the camera, even when clearly visible in the RGB image due to the low resolution of the depth camera and the handle being so narrow.

To scan the cup, Baxter grippers are lowered around the centroid until the table is sensed from the tip sensors. The wide grippers are able to freely rotate around the cup without catching on the handle. We use the known location of the sensors as well as the calculated proximity to the cup to estimate the diameter of the object in between the sensors on each finger stepped through in the Handle Search Algorithm in the appendix. When the diameter measured is greater than the cylindrical diameter, we know that we have found the handle between the fingers. For this object, the handle is always closest to the sensors that measure the least proximity due the cup centroid from vision being within $1cm$ of the true centroid. For other objects where distance to the sensor is not enough information, irregularities can be detected from the rate of change of proximity measurements during scanning. As shown in Figure 4.3, sensor measurements scanning over the handle experience a high rate of change.

4.3.1 Experimental Validation

To verify that we do in fact see improvement in grasping when we utilize proximity sensing, we tested the system’s ability to detect the handle of the cup shown in Figure 4.2 when vision fails to do so. The wide Baxter grippers shown in the middle image are just wide enough to perform a full rotation around the cup’s centroid without bumping the handle. However, the wider grippers, even when closed fully, are too wide to grasp the cup around its cylindrical center. This leaves the only successful grasping pose to be around the handle, so if the handle is not found, it can be assumed that these grippers will almost always fail to grasp the cup.




Figure 4.3: Missing features in the RGB-D point cloud are found using the proximity sensors embedded in the robotic hand. Even when in full view of the camera, the cup’s handle is often not discernible in the point cloud.

Finding the handle through vision largely depends on the noise in the point cloud and the point of view from the camera. When it is clearly visible, we can accurately estimate the location of the handle. We iterate through the point cloud and measure the distance from each point to the centroid. When we encounter distances that are greater than the radius of the cylindrical center, we know that we have found the handle and publish its location. To provide ground truth, a human is used to verify the results between each run. It was found that vision was always accurate when the handle is visible to the camera and only failed when the cup occluded the handle. Table 4.2 shows that the probability of the handle being found from this was 78.13%. We were expecting lower success due to the handle not always being visible in the point cloud, even when it is clearly visible to the camera. However, over the course of a few frames the likelihood of one frame containing

points on the handle is dramatically increased. These vision methods could easily be substituted for a more complex orientation estimation algorithm.

Sensor Data	Handle Found
RGB-D Image	78.13%
Proximity Sensing	100%
RGB-D Image and Proximity Sensing	100%

Table 4.2: We show our results for estimating the pose of the cup from an RGB-D camera alone and from using proximity measurements in the Baxter parallel grippers.

After a handle position is received, we test the sensors ability to find it in the same place, or if not found, be able to accurately estimate the position. The recognized cup is assumed to be in a well calibrated scene so that no search motion is required. The grippers are lowered to the cup and rotated until the handle is scene or 180 degrees mentioned previously, the diameter of the slice of the object between the grippers is measured using proximity sensing on both fingers. When a diameter greater than the threshold is encountered, we know that the handle is between the grippers. Whichever sensors measured the least distance, meaning that the object surface is much closer to that finger, was assumed to be the side that the handle is on. When the handle is found, again the human in the loop is asked to verify the results to provide ground truth. We found that this was successful 100% of the time. This was expected since the sensor values were calibrated for this object and the surface remained parallel to the fingers throughout scanning.

Finally, to test the handle location found, the grippers would grasp the cup and rotate a random angle to put the handle in a new location. The arm would then move to a home location to prevent occluding the cup from the camera and the experiment is rerun.

4.3.2 Future Work

Although the proximity measurements were specific to the cup, they have shown to produce non-noisy 3D models of other objects like the cubes in 4.1. This can be used for a more general feature detection algorithm discussed next.

4.4 Occluded Object Detection

In an unconstrained environment, there will often be situations where no vision information is available to the camera such as the basket shown in Figure 4.4. Another example is the Amazon picking challenge [9] where many contestants mounted the RGB-D camera on the arm so it could be moved to a position that enables it to capture images directly inside the storage bins. However, this limited perception is lost when the end-effector is moved inside the storage bin, blocking the camera’s view. In both situations, perception at the end-effector is necessary for detecting objects relative to the hand and determining an optimal grasp.

Using similar scanning trajectories to those used to find objects on a table, occluded objects in an unseen region can be detected by moving the sensors slowly in a defined manner until objects or the boundaries are detected. Shown in Figure 4.4, to detect the salt shaker in the basket, the grippers were slowly lowered down into the basket repeatedly until the inner finger sensors detected an object surface or the tip sensors detected the bottom of the basket. If one of the tip sensors was lowered directly onto the top of the salt shaker, a "false positive" for the bottom of the basket was identified from the large difference between the locations of the known bottom and where the sensor detects it. The grippers could then be moved horizontally until they can lower without crushing the shaker.

The point clouds generated from the in-hand sensing contain the same type of points as those obtained from a depth camera, although at a the lower resolution. Compared to a depth camera which captures 640x480 points per frame, the sensors obtain at most 16 points per frame (1 point per sensor). However, many sensor frames can be combined to generate 3D models comparable

and more complete than those from the camera as shown in Figure 4.3. Although not investigated in this paper, the 3D models from the sensors are compatible with the perception pipeline from Section IV. When object recognition is not successful due situations such as overlapping objects, methods such as [51] that use the point cloud alone to find grasping poses can still be applied due to the point cloud compatibility.




Figure 4.4: Exploring areas where no global perception information is available.

4.5 Conclusion

The presented techniques for fusing sensor data from the camera and in-hand sensing are simple yet effective for improving perception and grasping performance. We have shown that such sensors embedded in the hand greatly improved object pose estimation from the camera by providing a feedback mechanism to correct scene calibration errors. We also showed how proximity sensing can be used to detect irregular features and generate a 3D point cloud that can be merged with the depth camera point cloud and is compatible with widely used point cloud processing libraries.

Now that we have demonstrated the improvement from using proximity sensors that are susceptible to surface properties and sharp features of objects, we would like to continue this work using less susceptible sensors that measure time-of-flight rather than reflection. Finally, we will test the object recognition methods such as correspondence matching on the point clouds from

the sensors to determine if we can design a comparable system that can perform grasping and manipulation without the need for global perception in the loop.

Bibliography

- [1] DJ Balek and RB Kelley. Using gripper mounted infrared proximity sensors for robot feedback control. In Robotics and Automation. Proceedings. 1985 IEEE International Conference on, volume 2, pages 282–287. IEEE, 1985.
- [2] Stefan Begej. Planar and finger-shaped optical tactile sensors for robotic applications. Robotics and Automation, IEEE Journal of, 4(5):472–484, 1988.
- [3] Gines Benet, Francisco Blanes, José E Simó, and Pascual Pérez. Using infrared sensors for distance measurement in mobile robots. Robotics and Autonomous systems, 40(4):255–266, 2002.
- [4] Jeannette Bohg, Aythami Morales, Tamim Asfour, and Danica Kragic. Data-driven grasp synthesis—a survey. Robotics, IEEE Transactions on, 30(2):289–309, 2014.
- [5] Max Born and Emil Wolf. Principles of optics: electromagnetic theory of propagation, interference and diffraction of light. Cambridge University press, 1999.
- [6] Dengke Cai, Andreas Neyer, Rüdiger Kuckuk, and H Michael Heise. Raman, mid-infrared, near-infrared and ultraviolet–visible spectroscopy of pdms silicone rubber for characterization of polymer optical waveguide materials. Journal of Molecular Structure, 976(1):274–281, 2010.
- [7] Ziliang Cai, Weiping Qiu, Guocheng Shao, and Wanjun Wang. A new fabrication method for all-PDMS waveguides. Sensors and Actuators A: Physical, 204:44–47, 2013.
- [8] Berk Calli, Aaron Walsman, Arjun Singh, Siddhartha Srinivasa, Pieter Abbeel, and Aaron M Dollar. Benchmarking in manipulation research: The YCB object and model set and benchmarking protocols. arXiv preprint arXiv:1502.03143, 2015.
- [9] Nikolaus Correll, Kostas E Bekris, Dmitry Berenson, Oliver Brock, Albert Causo, Kris Hauser, Kei Okada, Alberto Rodriguez, Joseph M Romano, and Peter R Wurman. Analysis and observations from the first amazon picking challenge. IEEE Transactions on Automation Science and Engineering, 2016.
- [10] R. Cox and N. Correll. Merging local and global 3d perception using contact sensing. In AAAI Spring Symposium on Interactive Multi-Sensory Object Perception for Embodied Agents, Stanford, CA, 2017.
- [11] Ravinder S Dahiya, Giorgio Metta, Maurizio Valle, and Giulio Sandini. Tactile sensing: from humans to humanoids. Robotics, IEEE Transactions on, 26(1):1–20, 2010.

- [12] Hao Dang and Peter K Allen. Stable grasping under pose uncertainty using tactile feedback. Autonomous Robots, 36(4):309–330, 2014.
- [13] Raphael Deimel and Oliver Brock. A novel type of compliant, underactuated robotic hand for dexterous grasping. Robotics: Science and Systems, Berkeley, CA, pages 1687–1692, 2014.
- [14] Aaron M Dollar, Leif P Jentoft, Jason H Gao, and Robert D Howe. Contact sensing and grasping performance of compliant hands. Autonomous Robots, 28(1):65–75, 2010.
- [15] Nicholas Farrow and Nikolaus Correll. A soft pneumatic actuator that can sense grasp and touch. In Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on, pages 2317–2323. IEEE, 2015.
- [16] Nicholas Farrow, Yang Li, and Nikolaus Correll. Morphological and embedded computation in a self-contained soft robotic hand. arxiv:1605.00354, 2016.
- [17] RONALDS Fearing. Simplified grasping and manipulation with dextrous robot hands. IEEE Journal on Robotics and Automation, 2(4):188–195, 1986.
- [18] Javier Felip and Antonio Morales. Robust sensor-based grasp primitive for a three-finger robot hand. In Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on, pages 1811–1816. IEEE, 2009.
- [19] Markus Fleger and Andreas Neyer. Pdms microfluidic chip with integrated waveguides for optical detection. Microelectronic engineering, 83(4):1291–1293, 2006.
- [20] D Goger, Hosam Alagi, and Heinz Wörn. Tactile proximity sensors for robotic applications. In Industrial Technology (ICIT), 2013 IEEE International Conference on, pages 978–983. IEEE, 2013.
- [21] Bonnie L Gray and Ronald S Fearing. A surface micromachined microtactile sensor array. In Robotics and Automation, 1996. Proceedings., 1996 IEEE International Conference on, volume 1, pages 1–6. IEEE, 1996.
- [22] Daisuke Gunji, Takuma Araki, Akio Namiki, Aiguo Ming, and Makoto Shimojo. Grasping force control of multi-fingered robot hand based on slip detection using tactile sensor. Journal of the Robotics Society of Japan, 25(6):970–978, 2007.
- [23] Greg Hellard and R Andrew Russell. A robust, sensitive and economical tactile sensor for a robotic manipulator. In Australian Conference on Robotics and Automation, pages 100–104. Citeseer, 2002.
- [24] Kaijen Hsiao, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. Robust grasping under object pose uncertainty. Autonomous Robots, 31(2-3):253–268, 2011.
- [25] Kaijen Hsiao, Paul Nangeroni, Manfred Huber, Ashutosh Saxena, and Andrew Y Ng. Reactive grasping using optical proximity sensors. In Robotics and Automation, 2009. ICRA’09. IEEE International Conference on, pages 2098–2105. IEEE, 2009.
- [26] Dana Hughes and Nikolaus Correll. Texture recognition and localization in amorphous robotic skin. Bioinspiration & biomimetics, 10(5):055002, 2015.

- [27] Leo Jiang, Kevin Low, Joey Costa, Richard J Black, and Yong-Lae Park. Fiber optically sensorized multi-fingered robotic hand. In Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on, pages 1763–1768. IEEE, 2015.
- [28] Jelizaveta Konstantinova, Agostino Stilli, and Kaspar Althoefer. Force and proximity fingertip sensor to enhance grasping perception. In Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on, pages 2118–2123. IEEE, 2015.
- [29] Torsten Kröger, Bernd Finkemeyer, Simon Winkelbach, Lars-Oliver Eble, Sven Molkenstruck, and Friedrich M Wahl. A manipulator plays Jenga. Robotics & Automation Magazine, IEEE, 15(3):79–84, 2008.
- [30] Hyung-Kew Lee, Jaehoon Chung, Sun-Il Chang, and Euisik Yoon. Normal and shear force measurement using a flexible polymer tactile sensor with embedded multiple capacitors. Microelectromechanical Systems, Journal of, 17(4):934–942, 2008.
- [31] Emanuele Luberto, Yier Wu, Gaspare Santaera, Marco Gabiccini, and Antonio Bicchi. Enhancing adaptive grasping through a simple sensor-based reflex mechanism. IEEE Robotics and Automation Letters, 2017.
- [32] Lu Ma, Mahsa Ghafarianzadeh, Dave Coleman, Nikolaus Correll, and Gabe Sibley. Simultaneous localization, mapping, and manipulation for unsupervised object discovery. In IEEE International Conference on Robotics and Automation, pages 1344–1351, 2015.
- [33] Alexis Maldonado, Humberto Alvarez, and Michael Beetz. Improving robot manipulation through fingertip perception. In Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on, pages 2947–2954. IEEE, 2012.
- [34] Alexis Maldonado, Ulrich Klank, and Michael Beetz. Robotic grasping of unmodeled objects using time-of-flight range data and finger torque information. In Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on, pages 2586–2591. IEEE, 2010.
- [35] Brian Mayton, Louis LeGrand, and Joshua R Smith. An electric field pretouch system for grasping and co-manipulation. In Robotics and Automation (ICRA), 2010 IEEE International Conference on, pages 831–838. IEEE, 2010.
- [36] MA McEvoy and N Correll. Materials that couple sensing, actuation, computation, and communication. Science, 347(6228):1261689, 2015.
- [37] Philipp Mittendorfer and Gordon Cheng. Humanoid multimodal tactile-sensing modules. Robotics, IEEE Transactions on, 27(3):401–410, 2011.
- [38] Katsuhiko Nakamura and Hiroyuki Shinoda. A tactile sensor instantaneously evaluating friction coefficients. In Transducers 01 Eurosensors XV, pages 1402–1405. Springer, 2001.
- [39] Lorenzo Natale and Eduardo Torres-Jara. A sensitive approach to grasping. In Proceedings of the sixth international workshop on epigenetic robotics, pages 87–94, 2006.
- [40] John Oberlin and Stefanie Tellex. Learning to pick up objects through active exploration. In Development and Learning and Epigenetic Robotics (ICDL-EpiRob), 2015 Joint IEEE International Conference on, pages 252–253. IEEE, 2015.

- [41] Radhen Patel, Jorge Canardo Alastuey, and Nikolaus Correll. Improving grasp performance using inhand proximity and dynamic tactile sensing. In Int. Symp. on Experimental Robotics (ISER), Tokyo, Japan, 2016.
- [42] Radhen Patel, Jorge Canardo Alastuey, and Nikolaus Correll. Improving grasp performance using in-hand proximity and force sensing. In Int. Symposium on Experimental Robotics (ISER), Tokyo, Japan, 2016.
- [43] Radhen Patel and Nikolaus Correll. Integrated force and distance sensing using elastomer-embedded commodity proximity sensors. In Proceedings of Robotics: Science and Systems, 2016.
- [44] Radhen Patel, Rebecca Cox, Branden Romero, and Nikolaus Correll. Improving grasp performance using in-hand proximity and contact sensing. arXiv preprint arXiv:1701.06071, 2017.
- [45] Joseph M Romano, Kaijen Hsiao, Günter Niemeyer, Sachin Chitta, and Katherine J Kuchenbecker. Human-inspired robotic grasp control with tactile sensing. IEEE Transactions on Robotics, 27(6):1067–1079, 2011.
- [46] Jonathan Rossiter and Toshiharu Mukai. An led-based tactile sensor for multi-sensing over large areas. In Sensors, 2006. 5th IEEE Conference on, pages 835–838. IEEE, 2006.
- [47] Marco Santello, Martha Flanders, and John F Soechting. Patterns of hand motion during grasping and the influence of sensory guidance. Journal of Neuroscience, 22(4):1426–1435, 2002.
- [48] Jean-Philippe Saut, Serena Ivaldi, Anis Sahbani, and Philippe Bidaud. Grasping objects localized from uncertain point cloud data. Robotics and Autonomous Systems, 62(12):1742–1754, 2014.
- [49] Jivko Sinapov, Darren Earl, Derek Mitchell, and Heiko Hoffmann. Interactive audio-tactile annotation of 3d point clouds for robotic manipulation.
- [50] Johan Tegin and Jan Wikander. Tactile sensing in intelligent robotic manipulation—a review. Industrial Robot: An International Journal, 32(1):64–70, 2005.
- [51] Andreas ten Pas and Robert Platt. Using geometry to detect grasp poses in 3d point clouds. In Intl Symp. on Robotics Research, 2015.
- [52] Yaroslav Tenzer, Leif P Jentoft, and Robert D Howe. The feel of mems barometers: inexpensive and easily customized tactile array sensors. Robotics & Automation Magazine, IEEE, 21(3):89–95, 2014.
- [53] Federico Tombari, Samuele Salti, and Luigi Di Stefano. Unique signatures of histograms for local surface description. In European conference on computer vision, pages 356–369. Springer, 2010.
- [54] Sha Ye, Kenji Suzuki, Yosuke Suzuki, Masatoshi Ishikawa, and Makoto Shimojo. Robust robotic grasping using ir net-structure proximity sensor to handle objects with unknown position and attitude. In Robotics and Automation (ICRA), 2013 IEEE International Conference on, pages 3271–3278. IEEE, 2013.

- [55] F Zhu and JW Spronck. A capacitive tactile sensor for shear and normal force measurements. Sensors and Actuators A: Physical, 31(1):115–120, 1992.

Appendix A

Algorithms

(1) The following variables are introduced in Algorithm 1 and 2:

- SAI_{0-15} , contact detection for each sensor. Sensors 0-6 are along the inside of the left finger, with 0 being closest to the palm. Sensors 7 and 15 are the tip sensors. Sensors 8-14 are along the inside of the right finger, with 8 being closet to the palm.
- S_{0-15} , sensor (x, y, z) locations using the same indexing.
- P_{0-3} , (x, y, z) coordinates
- C_o, C_f , (x, y, z) initial centroid estimation, (x, y, z) final centroid.
- $\{C_o\}, \{C_f\}$, the sets of initial and final centroids
- x_{0-15} , calculated proximity using the sensor value plugged into Equation 3.1
- ε , proximity threshold to the surface of an object
- $\Delta_{X,Y,Z}$, some offset in the X, Y , and/or Z direction
- R, t , rotation matrix, translation matrix
- T_c , camera-to-robot transformation
- θ , angle used to rotate the wrist
- D_{cup} , measured diameter of the object between the grippers
- D_{handle} , diameter in *cm* of the cup at the widest part where the handle is located

Algorithm 1: Scene Calibration

```

1 Function find_object_centroid(object_pose)
2   move_hand_to(object_pose +  $\Delta_Z$ );
3   while  $x_{7,15} > \epsilon_{table}$  do
4      $\lfloor$  move_hand_to(current_pose -  $\Delta_Z$ );
5     while !SAI6 do
6       move_hand_to(current_pose +  $\Delta_X$ );
7       if timeout then
8          $\lfloor$   $C_o = object\_pose + \Delta_{X,Y}$ ;
9          $\lfloor$  return find_object_centroid( $C_o$ );
10    while ( $x_6 < x_{14}$ ) || ( $x_{14} < x_6$ ) do
11       $\lfloor$  move_hand_to(current_pose  $\pm \Delta_X$ );
12    close_grippers();
13     $P_0 = S_6$ ;
14     $P_1 = S_{14}$ ;
15    open_grippers();
16    while  $x_{6,14} < \epsilon_{cube}$  do
17       $\lfloor$  move_hand_to(current_pose +  $\Delta_Y$ );
18    close_grippers();
19     $P_2 = S_6$ ;
20     $P_3 = S_{14}$ ;
21    return COM( $P_{0-3}$ );
22 Function calibrate_scene( $\{C_o\}$ )
23   for  $C_o$  in  $\{C_o\}$  do
24      $\lfloor$   $C_f = find\_object\_centroid(C_o - \Delta_{X,Y})$ ;
25      $\lfloor$   $\{C_f\}.add(C_f)$ ;
26   if size( $\{C_o\}$ ) < 3 then
27      $\lfloor$   $R = identity\_matrix(3 \times 3)$ ;
28      $\lfloor$   $t = ave(\{C_f\} - \{C_o\})$ ;
29   else
30      $\lfloor$   $R, t = LLS(\{C_f\}, \{C_o\})$ ;
31    $T_c = T_c * R + t$ ;
32   return  $T_c$ ;

```

Algorithm 2: Proximity Sensors Handle Search

```

1 Function find_cup_handle(cup_pose)
2   move_hand_to(object_pose +  $\Delta_Z$ );
3   while  $x_{7,15} > \epsilon_{table}$  do
4     | move_hand_to(current_pose -  $\Delta_Z$ );
5   while !handle_found do
6     | rotate_grippers( $\theta$ );
7     |  $x_{left} = \max(x_4, x_5, x_6)$ ;
8     |  $x_{right} = \max(x_{12}, x_{13}, x_{14})$ ;
9     |  $D_{cup} = \text{abs}(x_{left} - x_{right})$ ;
10    | if  $D_{cup} \geq D_{handle}$  then
11      | |  $handle\_found = True$ ;
12    | else if  $\theta > 180^\circ$  then
13      | | return (0, 0, 0);
14    | if  $x_{left} < x_{right}$  then
15      | | return  $S_{left} + x_{left}$ ;
16    | else
17      | | return  $S_{right} + x_{right}$ ;

```
