Localization and Manipulation of Small Parts Using GelSight Tactile Sensing

Rui Li, Robert Platt Jr., Wenzhen Yuan*, Andreas ten Pas*, Nathan Roscup*, Mandayam A. Srinivasan, Edward Adelson

Abstract— **Robust manipulation and insertion of small parts can be challenging because of the small tolerances typically involved. The key to robust control of these kinds of manipulation interactions is accurate tracking and control of the parts involved. Typically, this is accomplished using visual servoing or force-based control. However, these approaches have drawbacks. Instead, we propose a new approach that uses tactile sensing to accurately localize the pose of a part grasped in the robot hand. Using a feature-based matching technique in conjunction with a newly developed tactile sensing technology known as GelSight that has much higher resolution than competing methods, we synthesize high-resolution height maps of object surfaces. As a result of these high-resolution tactile maps, we are able to localize small parts held in a robot hand very accurately. We quantify localization accuracy in benchtop experiments and experimentally demonstrate the practicality of the approach in the context of a small parts insertion problem.**
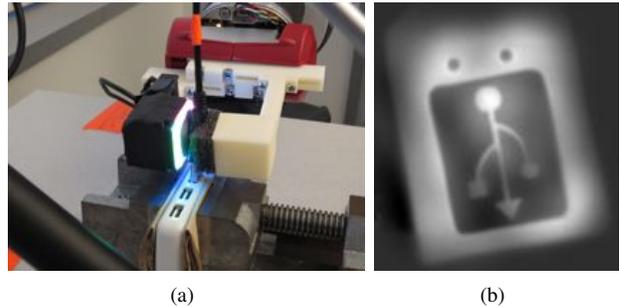
Fig. 1. (a) The robot, Baxter, inserts a USB connector into a mating hole. The robot relied on tactile sensing in order to localize the connector in the gripper and therefore to calculate how to move the hand. (b) Illustration of the tactile map that describes what the connector "feels like" to the robot. The brightness corresponds to the protrusion of the object in contact. The tactile map is used to localize the connector relative to the gripper.

I. INTRODUCTION

Small parts manipulation and insertion is an important robotics problem that has applications in manufacturing, space and hazardous environments, and medicine. A good example is the problem of grasping and inserting a USB cable, as shown in Figure 1. This insertion is challenging because the the tolerances are very low – less than plus or minus one millimeter. Nevertheless, this type of fine manipulation problem is important. Today, human factory workers are often employed to perform fine insertions and manipulation of exactly this kind.

There are two main types of approaches to performing an insertion such as the USB insertion in Figure 1. The first is force-based perception and/or control. A good example of this is the remote center of compliance (RCC) [1]. In RCC, an active or passive compliance mechanism is developed that causes the peg to slide into the hole when a downward force is applied. Alternatively, it is possible to perceive hole location based on sensed forces [2]. However, this technique is hard to use in many robot manipulation problems because of secondary load paths. Another important approach to performing fine insertions is visual servoing [3], [4]. Here, the robot vision system localizes features or fiducials both on the part to be inserted and on the mating surface. The transform between the two features is estimated using projective geometry in the reference frame of the camera and a robot motion that reduces the error relative to a desired transform is calculated. The advantage of this approach is accuracy: visual servoing can guide part insertion to within ten microns of error [4], [5]. However, it is necessary for the insertion operation to remain within view of the camera during the entire operation (a challenging requirement because the robot hand tends to occlude the mating surface). Moreover, it is necessary to be able to localize features on both the part and the mating surface consistently through time. Since this can often be challenging to do using "natural" features, fiducials are often affixed to the parts in order to facilitate consistent localization [5] (an undesirable modification of the environment).

Instead, this paper explores a tactile-sensing based approach to the problem. In contrast to force sensing methods, our approach is to use tactile sensing to localize a part accurately relative to the gripper holding it. We assume that the main challenge is localizing the part in the robot hand and not localizing the mating surface. This assumption is reasonable in many scenarios where the mating surface is fixed and can be localized prior to insertion or manipulation. In this case, it is the pose of the part in the hand that is hard to estimate. The key feature of our approach is the use of an tactile map [6]. The tactile map is a model of what the object surface is expected to feel like as a function of contact configuration and is created prior to manipulation (see Figure 1 (b)). During manipulation, the robot perceives tactile information regarding the pose of the object in the hand. By registering this tactile information back to the

tactile map, the robot can localize the pose of object relative to the gripper. In this paper, we use a recently developed tactile sensor, known as GelSight [7]. The GelSight sensor reconstructs the 3D geometry of the surface of a contacting object using photometric stereo algorithms. The resolution of the resulting height map is on the order of the number of camera pixels – 320×240 in our case. We use feature-based RANSAC operating on this height map both to create tactile maps and to localize a given set of tactile information within a map.

A. Related Work

The problem of localizing an object using tactile sensing has been studied for a long time. Early work included approaches based on fitting a parameterizable object model to contact points [8], using observability theory to estimate object pose [9], and using active tactile interaction to explore objects [10]. More recent work uses Bayesian estimation. For example, Chhatpar and Branicky use particle filtering to localize a peg with respect to a hole [11]. Gadeyne and Bruyninckx use Markov localization to localize the 3-dof pose of an object [12]. Petrovskaya *et al.* localize the 6-dof pose of an arbitrary polyhedral object by making a series of surface probes [13]. Corcoran and Platt localize the 6-dof pose of an object held in a whole-hand grasp based on a set of binary contact information [14]. A couple of prior works incorporate the idea of a tactile map. Platt *et al.* use tactile maps to localize distinctive haptic features in soft materials [6]. Pezzementi *et al.* use tactile models in order to classify and localize objects (for example, raised letters from a children's play set) [15]. Another important area of related work has to do with other tactile sensors that measure deformation in a deformable membrane. For example, Hristu, Ferrier, and Brockett proposed a deformable membrane tactile sensor that operates by tracking dots printed on a deformable membrane and reconstructing the contact geometry using a finite elements approach [16]. Wettels, Smith, Santos, and Loeb, developed a sensor that measured pressure in a weakly conductive fluid fingertip at a small set of locations [17]. Torres-Jara *et al.* developed a tactile sensor that used hall effect sensors to measure membrane deformations [18]. An interesting review of human and robot approach to tactile sensing can be found in [19].

II. SENSOR DESIGN AND INTEGRATION

A. GelSight Concept

GelSight is a recently developed tactile sensing technology that can measure the geometry of a contacted surface at a resolution as fine as a few microns [7], [20], [21]. This greatly exceeds the resolution of other available tactile sensors. For example, the Takktile array sensor senses independent forces over an 8×5 grid with approximately 6 mm resolution [22]. The RoboTouch sensor from Pressure Profile systems has 24 sensor elements with 6×6 mm resolution [23]. In contrast, the GelSight technology can sense contact geometry at approximately pixel resolution – 320×240 in our current sensors.
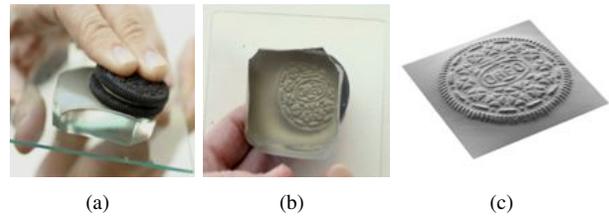

Fig. 2. Illustration of the GelSight operating concept. (a) An elastomer gel is coated with a reflective membrane; (b) the membrane deforms when the gel makes contact with an object; (c) with proper illumination, it is possible to calculate a 3D height map of the object surface using photometric methods.
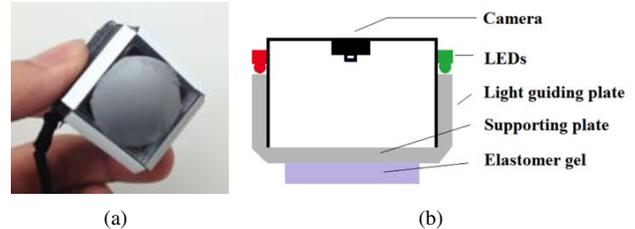

Fig. 3. GelSight sensor design suitable for mounting on a robot fingertip. (a) A fingertip GelSight sensor, and (b) its schematic diagram (side view).

The GelSight principle of operation is as follows. A piece of clear elastomer gel is coated with a reflective membrane (Figure 2(a)). When an object is pressed against the membrane, it deforms to take the shape of the object's surface (Figure 2(b)). The membrane is illuminated by LEDs that project from different directions so that multiple images of the same membrane deformation are captured. A 3D height map of the surface can then be reconstructed using photometric stereo algorithms (Figure 2(c)).

B. A Fingertip GelSight Sensor

In order to use GelSight in robot manipulation tasks, we designed a version of the sensor that can be mounted in the fingertip of a robot hand (Figure 3(a)). The elastomer gel is shown at the bottom of Figure 3(b). Above it is a hollow box. At the top of the box, a camera points downward through the box and onto the gel, and captures deformations in the gel caused by contact. A key aspect of the design is the way the gel is illuminated. We illuminate the gel from four sides simultaneously in four different colors: red (R), green (G), blue (B), and white (W). Each point on the gel has a color with three values corresponding to the R, G and B channels. We do the color calibration by pressing the GelSight sensor onto a hemisphere with a diameter of 5mm (known surface normals), and record the color values at each position of the hemisphere. We then create a lookup table that corresponds a color value to a surface normal value. With the lookup table, it is possible to retrieve surface normals from the color values on any contacted surface in real time. Then, the height map is calculated using Poisson integration from the surface normals. Overall we can calculate the height map at 10 frames per second in Matlab 2013b on a 2.8 GHz Intel Core i7 running 32-bit Windows 7. It is worth noting that
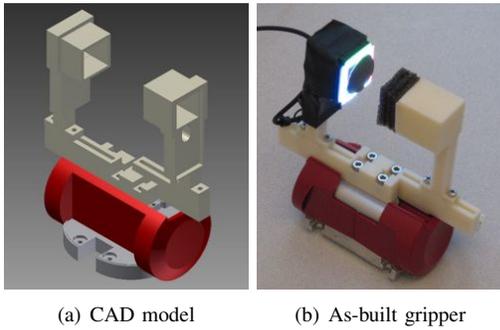
(a) CAD model   (b) As-built gripper

Fig. 4. Integration of the sensor into the Rethink Robotics Baxter hand. As (b) shows, one finger was equipped with a sensor in our experiments.

light from each of the four LEDs is directed by light guiding plates into the supporting plate and the gel. As a result, the path length of the light is maximized so as to simulate a more parallel illumination system as assumed in photometric stereo. This improves the accuracy of the resulting height map of the contacted object surface. All together, this sensor is a cube approximately 2.5 cm on a side. With a Logitech C310 camera, there is a tail due to camera dimension, which is not shown in Figure 3(a).

Figure 4 shows how the sensor is integrated and mounted into the Baxter gripper. The box at the end of each finger accommodates the illumination apparatus and the camera. The back of each finger has a slot specifically designed for mounting of the camera board from the Logitech C310 camera. As Figure 4(b) shows, only one of the fingers was equipped with a sensor. The other finger opposed the sensor with a compliant mesh pad. We have found that the shape of the elastomer gel on the sensor is important. We have explored two alternatives: gel with a flat membrane and gel with a domed membrane. While the flat membrane can make contact over a larger surface area, it can fail when the sensor is not aligned parallel to the object surface. Instead, we have used the domed membrane in our experiments (Figure 3(a)). For both flat and domed membranes, the maximum protrusion or recession that can be measured by the GelSight sensor is less than approximately ±1 mm.

III. LOCALIZATION AND MAPPING VIA IMAGE REGISTRATION

The key challenge in using tactile sensing to localize an object held in the robot hand is the creation and use of the *tactile map* of the object surface. The tactile map is a model of what the robot expects to feel as a function of the position and orientation of the object relative to the sensor. The map enables the robot to localize a grasped object in its grip. For example, Figure 1(b) illustrates a tactile map of one side of a USB connector, where the brightness corresponds to the amount of deformation in the GelSight sensor, and hence the protrusion of the contacting parts. When the robot grasps the connector (Figure 1(a)), the GelSight sensor mounted on its fingertip measures a height map of the portion of the connector surface where it is gripped. By matching this

height map with the corresponding portion of the tactile map, it is possible to localize the grasped object with respect to the gripper.

A. Registration of A Tactile Image

In order to create a new tactile map or to localize a tactile measurement within a map, it is necessary to register one height map with respect to another. This is very similar to the well-known image mosaicing problem. However, in our case, we are mosaicing height maps rather than RGB images. Nevertheless, we have found that standard feature-based matching techniques can work well. In our scenario, it can be assumed that the two height maps will have nearly the same scale and that there will be no out-of-plane rotation. Therefore, the problem reduces to that of estimating the isometry between the two height maps. Our approach is as follows. First, we localize keypoints and feature descriptors in both images using a recently developed detection algorithm that locates robust keypoints with binary descriptors, known as BRISK [24]. Then, we calculate the best fit pose using RANSAC [25]. Hypothesis poses are sampled uniformly at random by sampling two pairs of matching keypoints. The two keypoint pairs give us a candidate translation, $\mathbf{t} \in \mathbb{R}^2$, and rotation, $R \in SO(2)$. These are combined to give us a candidate isometry. After performing several rounds of sampling, we choose the isometry with the largest number of inliers and evaluate quality of fit. To do that, we calculate the least-squares homography between inliers in one height map and inliers in the other. Because we are matching tactile sensor information, the best fit homography should be an isometry (it should have only translation and rotation components). We evaluate the "distance" of the homography to an isometry by evaluating the determinant of the rotation component of the homography. We treat the determinant as a measure of our confidence that the match is correct. Let R be the rotation component of the homography. Then our confidence measure is:

$$c = \max(1 - |1 - \det(R)|, 0). \tag{1}$$

Confidence is highest when $c = 1$. A typical result of this approach to tactile image registration is shown in Figure 5. Figure 5(a) and (b) show two tactile images of overlapping areas of a penny. Figure 5(c) shows the composite registered tactile image.

B. Mapping

In this paper, we focus on mapping only a single face or side of an object. Some objects, such as a key or a USB connector, are well modeled this way because they are nearly always grasped with one finger on each of the two large flat sides. The tactile map is created on-line in a sequential way. We start by capturing a single height map of some part of the object surface as the "root" of the map. Then, we obtain additional tactile images by touching the object surface in different configurations. Each time a new image is acquired, we attempt to match it to the tactile map. If the match confidence exceeds a threshold (0.98 in our experiments),
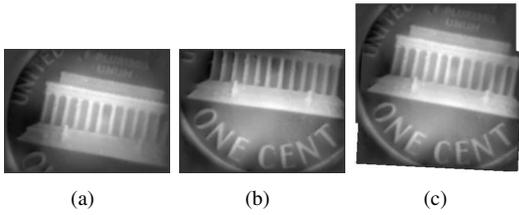
(a)   (b)   (c)

Fig. 5. (a) and (b) Two height maps created by touching different parts of a penny with the tactile sensor. The brightness corresponds to the amount of deformation in the GelSight sensor, and hence the protrusion or "height" of the contacting parts. (c) A composite height map created by registering the two components.

Fig. 6. Tactile map of the surface of a penny created by registering several partial tactile images into a single height map.

then we add it to the map. Height values in areas where the new image overlaps with the current map are averaged with height values from other images. In new areas, the height value is placed on the map directly. A complete height map of a penny is illustrated in Figure 6. In order to use the tactile map during manipulation, it is necessary to localize the pose of the map with respect to the gripper. This transform will be denoted ^gT_m, and can be estimated by measuring the pose of the "root" heightmap.

C. Localization Experiments

We performed experiments to characterize the localization accuracy of our approach. The goal of localization is to locate the grasped object relative to the gripper. When the gripper grasps an object, the GelSight sensor captures a height map of a segment of the object surface. This height map is registered with the tactile map and used to localize the object. Figure 7(a) shows the experimental setup. An object was fixtured to an adjustable $x - y - \theta$ platform and the tactile sensor was fixtured in a jig above the platform. This allowed us to adjust the position of the object relative to the sensor in a controlled way while capturing tactile data. During the experiment, we moved the object to a series of measured positions and orientations relative to the sensor and captured a height map from the sensor. For each height map captured this way, we registered it with respect to the tactile map (created prior to the experiment) and thereby calculated the pose of the object relative to the sensor. By comparing the measured pose of the object relative to the sensor and the estimated pose based on the registered height map, we were able to calculate localization error statistics.

We evaluated the accuracy of orientation and translation estimates separately. In order to evaluate orientation error, we collected data using a USB connector as the object. The connector was placed at orientations between -90 and $+90$ degrees in steps of 10 degrees. The comparison between true orientation and estimated orientation is shown in Figure 7(b). The mean absolute error was found to be 1.15 degrees with average standard deviation 1.02 degrees. We performed a similar experiment to evaluate translational accuracy (see Figure 7(c)). Here, we performed experiments using a quarter placed at displacements between -6 mm and $+6$ mm with steps of 1 mm. The mean absolute error for translation localization was found to be 0.14 mm, and the average standard deviation 0.055 mm. It is likely that a portion of the translation and orientation error that we report is a result of experimental error related to manual adjustment of the jig to produce the object translation and rotation.

IV. ROBOT INSERTION EXPERIMENTS

These experiments evaluate the effectiveness of using our approach to small part localization and manipulation in the context of an insertion task.

A. Setup

The basic setup is as follows. A USB cable hangs from a jig positioned in the robot workspace. The location of the jig is known, but the pose of the USB connector itself varies because the way in which the connector hangs is unknown (see Figure 8 (a)). The objective is for the robot to grasp the connector and insert it into a USB mating hole located in a known pose. In order to accomplish this, it is necessary to localize the connector relative to the mating hole with an error of less than approximately ± 1 mm. The robot calculates the expected location of the connector based on the jig location. It reaches to that position and closes the fingers. If the robot cannot localize the grasped connector in its grip, then it releases the hand and moves a small distance and tries again (according to a simple "blind" search procedure). Once the connector has been grasped in such a way that it is localized, then, the robot proceeds with the insertion (see Figures 8(b) and 8(c)).

B. Connector Alignment

After grasping the connector, the robot moves the gripper to a pose a few centimeters above the mating hole. Then, after localizing the connector in the grasp, it calculates a target pose that will align the connector just above the hole. This occurs as follows. Since we have assumed that the mating hole is fixtured to the environment (*i.e.* the base frame), we can calculate the target transform, $^bT_o{}^*$, for the connector with respect to the base frame. This target transform denotes the pose of the connector hovering just over the hole. The pose of the gripper in the base frame, ^bT_g, is available using the forward kinematics of the robot. The transform, ^gT_m, denotes the pose of the map with respect to the gripper. This must be measured during map creation. The map is created by registering tactile images relative
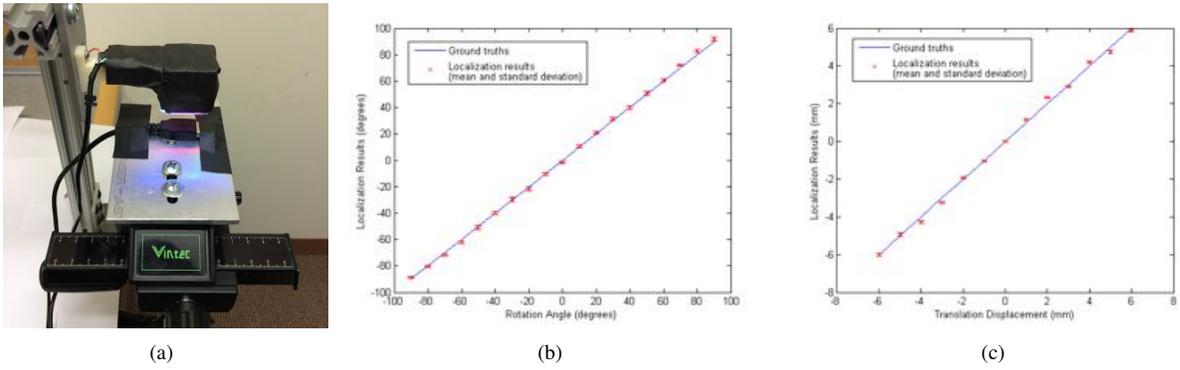
Fig. 7. Experiments characterizing the localization accuracy of our method. (a) Experimental setup. The GelSight sensor is fixed on a beam facing downwards, and the tripod and bench for hosting the USB with rotation and translation flexibilities. (b) Estimated orientation (from tactile data) as a function of true orientation. (c) Estimated translation as a function of true translation.

(a) Grasping the USB connector | (b) Insertion of the connector into the mating hole | (c) Insertion closeup
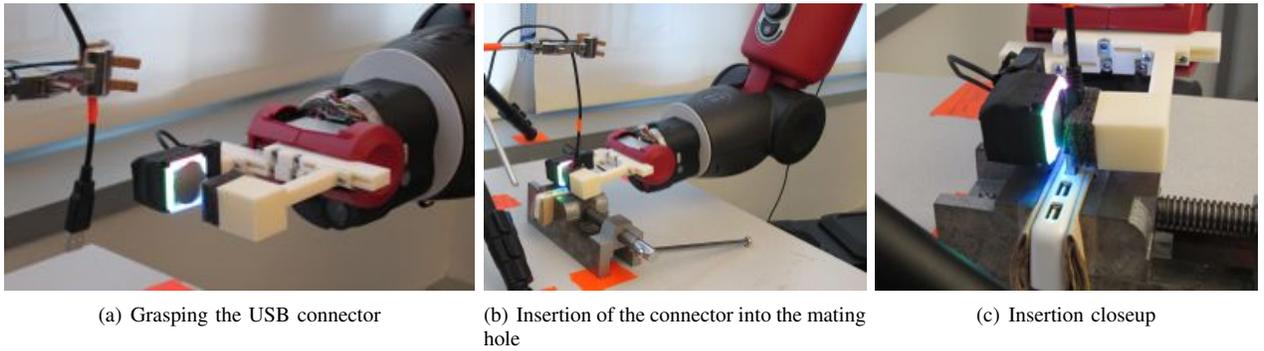
Fig. 8. USB connector insertion experimental scenario.

to a single root image. The pose of this root image with respect to the gripper must be measured and stored as ^gT_m. Finally, ^mT_o denotes the pose of the object in the map frame. This transform is calculated using feature-based height map registration. Given all of these transforms, we can calculate the target gripper pose as follows. The transforms are related by:

$$^bT_o = {^bT_g}\,{^gT_m}\,{^mT_o}\,.$$

Given the target transform, $^bT_o^*$, we calculate

$$^bT_g{}^* = {^bT_o}{}^*\,({^mT_o})^{-1}\,({^gT_m})^{-1}. \qquad (2)$$

$^bT_g{}^*$ describes the desired configuration of the gripper in the robot base frame and is used as input to an inverse kinematics solver or a Cartesian controller.

C. Connector Insertion

After localizing the USB connector in the hand, we solve Equation 2 for the target gripper pose, solve the inverse kinematics problem, and moved the gripper to the target location. Rather than using the joint position controller that ships with the Baxter SDK, we developed our own joint position controller. We found that the position error integrator implemented by the SDK did not perform well when the hand contacted the environment. Instead, our controller calculates velocity commands using a position control law: $v^* = s^*(q^* - q)/\|q^* - q\|$, where s^* denotes the desired joint

speed. This control law is accurate without using integration, and we found it to be stable.

After moving the gripper to the target pose, the USB connector was directly above the mating hole. At this point, the robot pressed the connector directly down into the hole. Because we wanted to limit the amount of force that we applied through our sensor, we did not require the USB cable to be fully inserted in order to count the insertion as a success. We only required the connector to be completely inside the mating hole so that continued downward pressure would cause the connector to become completely mated (see Figure 8(c)).

D. Insertion Experiment and Results

Our system performed 36 USB insertions with two failures using the Rethink Robotics Baxter robot. On each trial, the USB cable was placed in the jig as shown in Figure 8(a). The robot reached forward from the same starting pose toward a fixed pose with respect to the jig and closed the gripper. If the system failed to localize the USB cable in its grasp, it opened the gripper and moved 0.75 cm forward and tried again. If that failed, then the system opened the gripper and moved 1.5 cm back. This process repeated until the USB cable was localized. This procedure often succeeded, although multiple regrasps were sometimes required. However, if the robot failed to grip and localize the connector with high confidence ($c \geq 0.98$ in Equation 1), perhaps because the
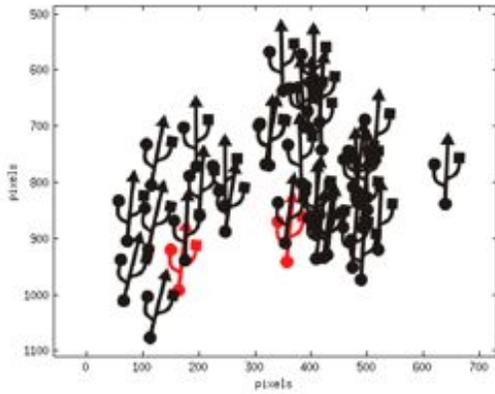
Fig. 9. The set of 36 gripper-connector poses experienced during our experiments. The poses shown in black were successful. The two in red failed. The distances on the axes are shown in pixels. Each pixel corresponds to approximately 0.005 mm of connector displacement.

connector was somehow pushed out of the 1.5 cm regrasp region, then the trial was stopped and the jig reset manually. This procedure resulted in the set of 36 relative gripper-connector configurations shown in Figure 9. Of these 36 grasps, 34 of the subsequent insertions succeeded. Grasp attempts where the connector was not localized successfully were not included in this dataset. Poses for the successful insertions are shown in black. The two failures are shown in red. While we are not certain of the causes of the two failed insertions, they were likely caused by inaccuracies in our joint position controller and the Baxter hardware. As a robot with series elastic actuators, the accuracy of Baxter's joint position control is fundamentally limited. In this work, we rely on our controller to eliminate these errors. However, there is always some joint error present because we do not use an integrator (because we contact the environment).

V. Conclusion

Fine parts manipulation and/or insertion is very challenging because of the fine tolerances that are typically involved. The key challenge in this kind of task is locating the part in the grasp. The precise pose of the part may be uncertain at the time of grasping or it may shift in the grasp as the robot moves. In either case, it is typically necessary to re-localize the part precisely just prior to insertion. In this paper, we explore an approach to localizing the part in the hand using tactile information. A key part of this work is our development and use of a robot fingertip version of the GelSight tactile sensor. This sensor delivers height maps of the surface of objects in the grasp at a much finer resolution than what is otherwise available. As a result of this key capability, we are able to use mapping and localization techniques to localize parts in the grasp very accurately.

Acknowledgments

References

[1] J. Craig, *Introduction to Robotics, Third Edition.* Pearson Prentice Hall, 2005.

[2] W. Bluethmann, R. Ambrose, M. Diftler, E. Huber, A. Fagg, M. Rosenstein, R. Platt, R. Grupen, C. Breazeal, A. Brooks, A. Lockerd, R. Peters, O. Jenkins, M. Mataric, and M. Bugajska, "Building an autonomous humanoid tool user," in *IEEE Int'l Conf. on Humanoid Robots*, 2004.

[3] R. Murray, Z. Li, and S. Sastry, *A Mathematical Introduction to Robotic Manipulation.* CRC Press, 1994.

[4] B. Yoshimi and P. Allen, "Integrating real-time vision and manipulation," in *Proc. of 13th Hawaii Int'l Conf. on System Sciences*, vol. 5, January 1997, pp. 178–187.

[5] W. Meeussen, M. Wise, S. Glaser, S. Chitta, C. McGann, P. Mihelich, E. Marder-Eppstein, M. Muja, V. Eruhimov, T. Foote, J. Hsu, R. Rusu, B. Marthi, G. Bradski, K. Konolige, B. Gerkey, and E. Berger, "Autonomous door opening and plugging in with a personal robot," in *IEEE Int'l Conf. on Robotics and Automation*, 2010.

[6] R. Platt, F. Permenter, and J. Pfeiffer, "Using bayesian filtering to localize flexible materials during manipulation," *IEEE Transactions on Robotics*, vol. 27, no. 3, 2011.

[7] M. Johnson and E. Adelson, "Retrographic sensing for the measurement of surface texture and shape," in *IEEE Int'l Conf. on Computer Vision and Pattern Recognition*, 2009.

[8] P. Allen and P. Michelman, "Acquisition and interpretation of 3-d sensor data from touch," *IEEE Transactions on Robotics and Automation*, vol. 6, no. 4, pp. 397–404, 1990.

[9] Y. Jia and M. Erdmann, "Pose and motion from contact," *International Journal of Robotics Research*, vol. 18, no. 5, pp. 466–490, 1999.

[10] A. Okamura and M. Cutkosky, "Feature detection for haptic exploration with robotic fingers," *International Journal of Robotics Research*, vol. 20, no. 12, pp. 925–938, 2001.

[11] S. Chhatpar and M. Branicky, "Localization in robotic assemblies with position uncertainty," in *IEEE Int'l Conf. on Intelligent Robots and Systems*, 2003, pp. 2534– 2540.

[12] K. Gadeyne and H. Bruyninckx, "Markov techniques for object localization with force-controlled robots," in *10th Int'l Conf. on Advanced Robotics*, 2001.

[13] A. Petrovskaya, O. Khatib, S. Thrun, and A. Ng, "Bayesian estimation for autonomous object manipulation based on tactile sensors," in *IEEE Int'l Conf. on Robotics and Automation*, 2006, pp. 707–714.

[14] C. Corcoran and R. Platt, "Tracking object pose and shape during robot manipulation based on tactile information," in *IEEE Int'l Conf. on Robotics and Automation*, vol. 2, 2010.

[15] Z. Pezzementi, C. Reyda, and G. Hager, "Object mapping, recognition, and localization from tactile geometry," in *IEEE Int'l Conf. on Robotics and Automation*, 2011.

[16] D. Hristu, N. Ferrier, and R. Brockett, "The performance of a deformable-membrane tactile sensor: Basic results on geometrically-defined tasks," in *IEEE Int'l Conf. on Robotics and Automation*, 2000.

[17] N. Wettels, L. Smith, V. Santos, and G. Loeb, "Deformable skin design to enhance response of a biomimetic tactile sensor," in *IEEE Int'l Conf. on Biomedical Robotics and Biomechatronics*, 2008.

[18] E. Torres-Jara, I. Vasilescu, and R. Coral, "A soft touch: Compliant tactile sensors for sensitive manipulation," *SAIL Technical Report MIT-CSAIL-TR-2006-014*, 2006.

[19] R. Dahiya, G. Metta, M. Valle, and G. Sandini, "Tactile sensingfrom humans to humanoids," *IEEE Transactions on Robotics*, June 2009.

[20] M. Johnson, F. Coley, A. Rajz, and E. Adelson, "Microgeometry capture using an elastomeric sensor," in *SIGGRAPH*, 2011.

[21] R. Li and E. Adelson, "Sensing and recognizing surface textures using a gelsight sensor," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.

[22] Y. Tenzer, L. Jentoft, and R. Howe, "Inexpensive and easily customized tactile array sensors using mems barometers chips," *IEEE*, 2012.

[23] P. P. Systems, "Products - RoboTouch," *http://www.pressureprofile.com/products-robotouch.php*.

[24] S. Leutenegger, M. Chli, and R. Siegwart, "Brisk: Binary robust invariant scalable keypoints," in *IEEE Int'l Conf. on Computer Vision*, 2011.

[25] M. Fischler and R. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, pp. 381–395, 1981.