Soft-bubble: A highly compliant dense geometry tactile sensor for robot manipulation

Alex Alspach[1], Kunimatsu Hashimoto[1], Naveen Kuppuswamy[1], and Russ Tedrake[1,2]

***Abstract*— Robots in your home and other environments are required to operate safely and manipulate robustly in unknown, uncertain and frequently cluttered environments. Tactile sensing is necessary when manipulating hard-to-see or occluded objects where vision falls short, and critical to maintaining awareness of grasp quality. We discuss the motivation and hardware design for a highly compliant, dense geometry tactile sensor; an inflated latex membrane with a depth sensor inside. This low cost, light weight and easy to build tactile sensor provides capabilities applicable to robust manipulation, which we demonstrate in multiple initial experiments, including successful tactile-based object classification, pose estimation and non-prehensile object manipulation. Through these experiments, we show the importance of high resolution geometry sensing for tactile tasks.**

I. Introduction

As we work towards deploying robotic systems to assist with everyday tasks, we must ensure that these robots are able to safely make contact with people, other robots, and our environment. Unlike industrial robots that can operate with certainty about their tasks and surroundings, robots designed for our homes and other unstructured environments must be able to cope with large imprecision in their knowledge of the surrounding environment. In particular, for manipulation tasks, the ability to compensate for large uncertainty through touching and feeling is increasingly perceived as the solution to coping with the perceptual challenges posed by domestic environments. These challenges include clutter, occlusions, variable lighting conditions, and never before seen objects, to name a few [1]. Due to their ability to directly capture interactions at the contacting surface, tactile sensors have the potential to be predominant when vision and other exteroceptive modalities are occluded [2] or incapable of sensing due to lack of sufficiently salient visual features [3]. Tactile sensing, coupled with the compliance needed to safely and robustly bump into and grasp objects in the environment could prove hugely beneficial in hastening the deployment of robots at home. Beyond these design requirements, for commercially viable home robots, it is also necessary to consider component cost, weight, and manufacturability.

It is well recognized that mechanical compliance is a critical element in enabling robots to cope with unforeseen contacts [4] and offers greater intrinsic robustness to uncertainty [5]. One successful strategy for endowing robots with the requisite compliance is the use of elastomers and fluids as building materials [6]. For slow speeds and low masses, the passive compliance offered by these materials reduces the impulse of a collision while a deformable contact patch spreads force out over a larger area [7]. Elastomers have made their way into and onto traditionally rigid robots, including fully compliant air-powered sensorless grippers [8], [9] where precision is not required and compliance alone provides a stable grasp. However, constrained home manipulation tasks generally require more precision and care.

In addressing the need to sense and manage physical interactions, there are several themes that have been tackled

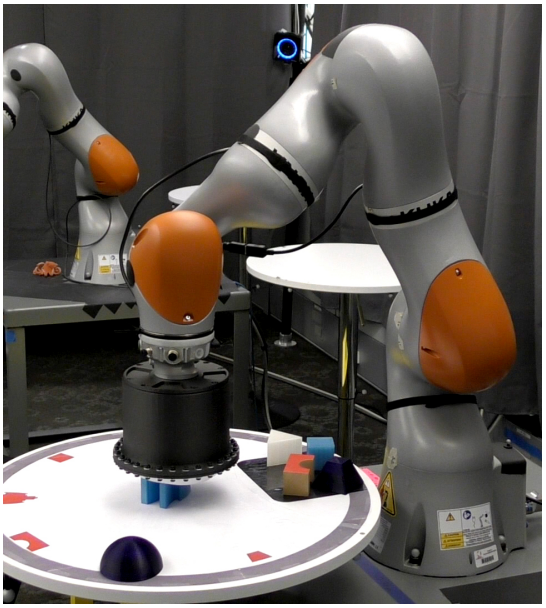

(a)

(b)

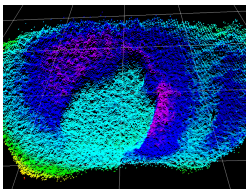

(c)

Fig. 1: (a) The soft-bubble, a high resolution tactile sensor, is attached to a Kuka IIWA and pressed against a robot-shaped block (a). The resulting depth image (b) and point cloud (c) are shown as well.

[1]Toyota Research Institute, One Kendall Square, Building 600, Cambridge, MA 02139, USA, alex.alspach@tri.global, kunimatsu.hashimoto@tri.global, naveen.kuppuswamy@tri.global, russ.tedrake@tri.global

[2]The Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, 32 Vassar St., Cambridge, MA 02139, USA, russt@mit.edu

so far: the need to sense accidental contacts and react suitably [10], tracking and controlling contact forces during intentional contact [11], tactile exploration for objects that are hard-to-see or occluded by clutter [3], classification of object type and shape inference [12] and sensing the quality of induced grasps [13]. For tactile sensing of object class (shape, material, etc.) and object state (pose, velocity, etc.), contact geometry sensing enables an understanding of surface and other physical properties [14] as well as pose refinement in order to manipulate it accurately [15]. Despite these advancements, the fundamental difficulties and open questions in modelling compliant contact mechanics [16] have limited the adoption and deployment of soft tactile sensors. Although data-driven methods have been employed as attempts to overcome the modeling difficulties [13], [17], there remains a lack of highly compliant mechanisms which also incorporate high resolution contact sensing.

In this paper, we present our *Soft-Bubble*, pictured in Fig. 1, a new kind of tactile sensor that combines the advantages of highly compliant air-filled elastomeric structures with the ability to sense detailed geometric features of contacting objects. The proposed sensor captures deformation of a thin, flexible air-filled membrane using an off the shelf depth sensor. This sensor is built using accessible fabrication methods and is composed primarily of off the shelf components and materials. The resulting sensor is highly compliant, lightweight, robust to continued contact, and outputs a high resolution depth image that is ideal for manipulation applications. We demonstrate the efficacy of these features through three case studies: (a) object shape and texture classification using a deep neural network, (b) an object sorting task using the soft-bubble end-effector for non-prehensile manipulation, and (c) object pose estimation using the geometry sensed at contact. For the latter two demonstrations, a version of the proposed sensor is mounted on a Kuka IIWA arm and used to classify, explore and manipulate blocks of various shapes and sizes.

This paper is organized as follows: Section II presents a background of dense geometry and compliant tactile sensing, section III illustrates the system design and IV presents the results of experimental evaluation. Finally we discuss our findings and outlook for future work in section V.

II. Related Work

In this section, we will briefly review some of the technologies and algorithms most relevant to the soft-bubble. The field of tactile sensing has seen dramatic growth in the last 25 years [18]. A wide variety of technologies have been proposed to sense shape, texture, hardness, temperature, vibration or contact forces [19], but several open questions still remain as to how tactile sensing can be employed with the explicit goal of improving a robot's ability to manipulate the world around it. How high a spatial resolution is necessary for tactile sensing? Is either geometry or force sensing more important than the other?

High resolution tactile sensors, such as GelSight [14], GelSlim [20] and FingerVision [21], use cameras to gather large amounts of data over relatively small contact areas. GelSight in particular uses precise internal lighting and photometric stereo algorithms to generate height maps of contacting geometry. This 3D information can be used in a model-based framework [15] wherein it can be fused with external sensors, or run through particle filters [22] that capture the complicated contact mechanics, or used to directly sample the contact manifold as dictated by the manipuland and gripper geometries [23]. Alternative data-driven approaches also exist [13], although these methods have been largely employed in the domains of slip detection, grasp stability identification [17] and material property discrimination.

The soft-bubble draws influence from these camera-based tactile sensors, particularly on its use of an off-the-shelf depth camera and an opaque membrane which drapes sensed object surfaces in consistent color and reflectance properties. Mechanically, the soft-bubble is able to deform around a contacting object more freely and drastically than the gel-based sensors above, potentially making a larger slice of the object's geometric form available to the sensor. As a result of the internal depth sensor, precisely placed illumination and 3D reconstruction algorithms are not needed to capture deformation, allowing a large range of free-form membrane shapes. This lack of internal hardware requirements leads to a cheap, simple and repeatable fabrication process. The use of air over gel also simplifies fabrication, reduces cost and keeps the sensor lightweight regardless of size, making it a suitable sensor for low cost, low payload robots. Employing the resilience of latex, the sensor membrane can withstand tough treatment and worn components are easily replaced. Finally, the high friction of the membrane surface, large contact patch and ability to provide form closure via deformation around an object make this sensor well-suited for contact-heavy manipulation tasks.

III. Hardware Design

The soft-bubble dense geometry sensor design consists of three main functional components: an elastic membrane sensing surface, an airtight hull that allows pressurization of the membrane, and an internal depth sensor, as shown in the dimensioned cross-sectional view and exploded view shown in Fig. 2. The depth sensor is located inside of the airtight hull pointed at the interior surface of the membrane to measure contact induced deformation. Before inflation, the circular membrane's diameter is 150mm. Inflated, the sensing membrane forms a compliant spherical cap. The membrane is inflated to a height of 20-75mm, depending on the application. The membrane material is hand- or laser-cut 0.4mm thick latex sheet. Thinner or thicker sheet can be used with trade-offs between sensitivity and durability. For the experiments presented in Section IV, the sensing membrane is inflated to a height of 50mm with an internal air pressure of roughly 0.25psi.

A. Time of flight camera

The time of flight camera in the embodiment presented here is the PMD CamBoard pico flexx [24]. The pico flexx is

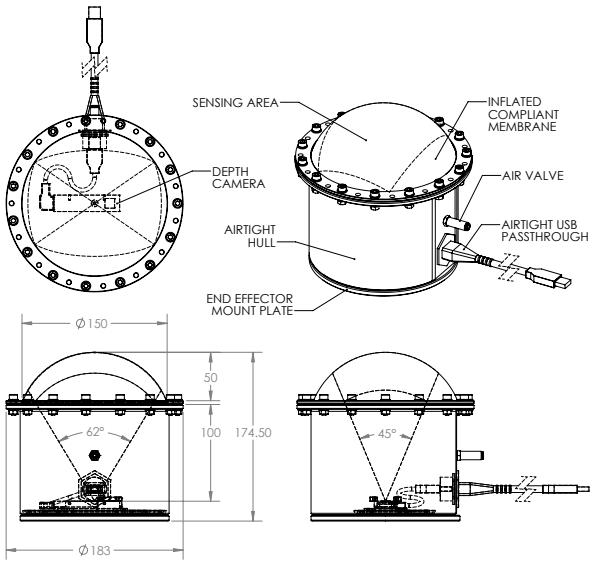

Fig. 2: Dimensioned orthographic and isometric views of the sensor assembly. All dimensions in mm.

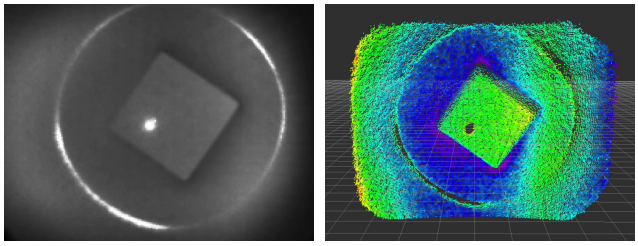

(a) IR image (b) Point cloud

Fig. 3: Glare due to excess emitter intensity

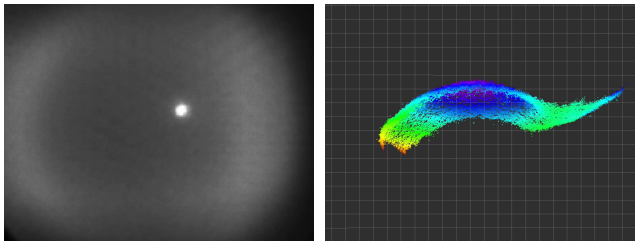

(a) IR image (b) Point cloud

Fig. 4: The dimly illuminated area to the right of the IR image forms during close range imaging as a result of the uncompensated emitter and imager FOV offset.

a development board with a depth imager intended for mobile devices, but has found its way into robotics. This depth camera was chosen for its small size (68×17×7.35mm), USB connectivity and short specified minimum sensing range (100mm). The sensor has a resolution of 224×171 pixels and can provide depth images at up to 45 fps. At close range, the depth resolution of this sensor is $\leq 2\%$ of the actual distance. The sensor weighs 8 grams and has a field of view (FOV) of 62°×45°. When inflated to 50mm, the spacial sensing resolution on the membrane is roughly 2 pixels/mm^2

The pico flexx depth camera is designed to sense within the range of 0.1-4 m. The optics are focused and the sensor is calibrated to work best in the middle of that range. When operating near the minimum range in a confined space, two issues become evident: The IR emitter is too bright and the effect of the IR emitter-imager offset, negligible at further distances, becomes non-negligible. The emitter intensity, which can not be altered in software, produces bright spots on regions of the inner membrane surface and causes sensor saturation seen in Fig. 3. It is possible to adjust exposure, but not enough to avoid saturation. Also, at this range, the offset of the emitter FOV and the imager FOV leaves a poorly illuminated region where the two FOVs fail to overlap. This asymmetric dark region in the infrared image reveals itself in depth as further away than it actually is, seen in Fig. 4. This internal nonuniform lighting can be partially mitigated by placing a diffuser (a piece of gift wrapping tape) on the IR emitter. This light diffusion enables a slightly wider view and a more symmetric depth image, but causes an increase in depth measurement noise and a significant reduction in depth accuracy.

B. Airtight hull

The hull structure provides mounting points for the depth sensor, elastic membrane, a sealed USB 3.0 passthrough and an air valve. This structure also provides mounting points for attaching the tactile sensor assembly to a robot. The dimensions of the hull and membrane are dependent on the chosen depth sensor's range and FOV specifications. At the depth sensor's minimum sensing distance, 100mm, the area covered by the FOV is 99.5cm^2. The diameter of the membrane is chosen so that the entire FOV falls onto the membrane. At 50mm inflation, this provides a sensing surface area of 175.4cm^2. The dome's total surface area is 261.4cm^2, which leaves about 86cm^2 out of the depth sensor's FOV. The height of the airtight hull is chosen so that the uninflated membrane plane is placed 100mm from the depth sensor. The diameter of the hull is chosen as to be both out of the FOV of the depth sensor, and to place the entire depth sensor FOV onto the membrane.

Towards the goals of low overall cost and weight, as well as accessible fabrication, the airtight hull is 3D printed. This particular version is printed on a Markforged X7 fused deposition modeling (FDM) printer [25]. The material used is Nylon packed with chopped carbon fiber. The material stiffness added by the carbon fiber allows the part to be printed at 50% infill with thin walls, leaving internal air pockets and using less material overall. FDM printing generally results in a porous part that can not hold air, therefore the internal surfaces of the hull, holes, and the O-ring channel where the latex membrane meets the hull, are painted with two or more coats of Rust-Oleum 265495 Leak Seal Flexible Rubber Sealant. The hull can also be printed on consumer

FDM printers and sealed in the same fashion, but the result may not be as lightweight. The part can also be printed on a printer that produces non-porous parts, like a Stratasys Objet, or it can be machined from plastic or other nonporous materials.

The hull is designed with a generic mounting pattern at the base for mounting a depth sensor. A separate sensor cradle is designed to mate with this pattern and hold the sensor, as well as diffusers or filters, in place. This generic pattern allows the sensor or sensor mount to be changed without remanufacturing the entire airtight hull. The USB passthrough used in this assembly is a sealed Molex 84733-series USB 3.0 Type A connector with its female end inside the hull and its male end extending outside of the hull. A short USB 3.0 male micro to male Type A cable is used to connect the depth sensor to the USB passthrough. The valve installed into the hull for pressurization and depressurization of the sensor is a M5 threaded push-to-connect fitting (Mcmaster part number 1201N11) for 4mm OD tubing that automatically shuts when the tube is removed.

The complete sensor is assembled by installing the USB passthrough and air valve, bolting the depth sensor to the interior floor of the hull, then connecting the sensor to the USB passthrough. An O-ring is placed into the channel on top of the hull flange. The circular latex membrane is glued to a 3D printed frame while flat so that no stretching or sagging occurs while assembling. This frame and membrane subassembly is then bolted to the hull flange so that the O-ring seals between the hull's rubberize O-ring channel below and the latex membrane above. The complete assembly can then be attached to the end effector of a robot using an adpater plate. The assembled soft-bubble tactile sensor can be seen in Fig. 1a installed on a Kuka IIWA, which we like to think of as one large, dexterous finger.

C. Parallel gripper concepts

As depth sensors and their sensing ranges become smaller, this technology will be incorporated into sensors of a scale more suitable for standard robot grippers. Figure 5 shows a sensorless bubble-based parallel gripper prototype able to execute robust grasps on arbitrary objects due to its compliant, high friction latex membrane.

IV. Experiments

For an initial demonstration of the soft-bubble tactile sensor capabilities and potential applications, we conducted multiple experiments: two tactile-based object classification experiments, an object sorting manipulation experiment, and a tactile pose estimation experiment. Overall, these experiments show the utility of high resolution contact geometry sensing and compliant manipulation capabilities.

A. Classification

We use ResNet18 [26], a state-of-the-art deep neural network, as the object classifier for both of the classification experiments described in this section. As input, the network takes a depth image obtained when the soft-bubble sensor is pressed up against an object. The network outputs the probabilities for each object category.

First, we explored whether or not an object classification task could be performed using the data output by the soft-bubble. We set up seven classes for this experiment: six objects (Fig. 6) and an extra class called "no-touch" for when the soft-bubble was free from contact. The objects chosen for this first experiment are distinctly shaped blocks.

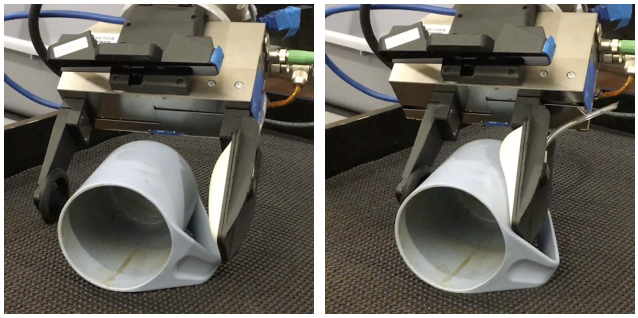

(a) Pre-grasp (b) Grasping

Fig. 5: Sensorless bubble-based parallel gripper prototype able to robustly grasp due to compliant, high friction membrane surface.

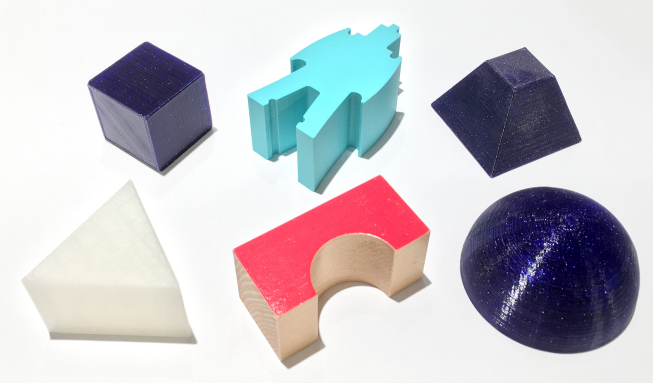

Fig. 6: Six blocks used for soft-bubble-based object classification experiment

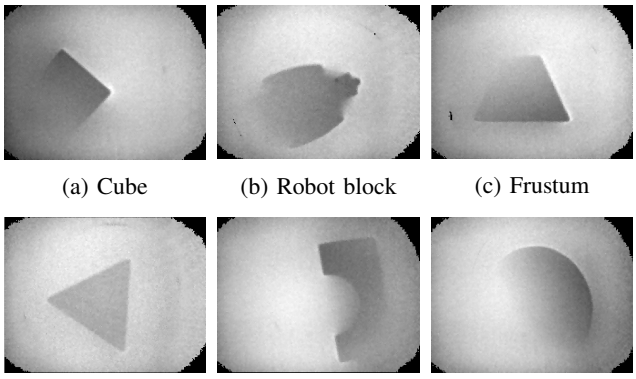

(a) Cube (b) Robot block (c) Frustum

(d) Triangular prism (e) Bridge (f) Hemisphere

Fig. 7: Depth images captured by the tactile sensor of the six objects classified.

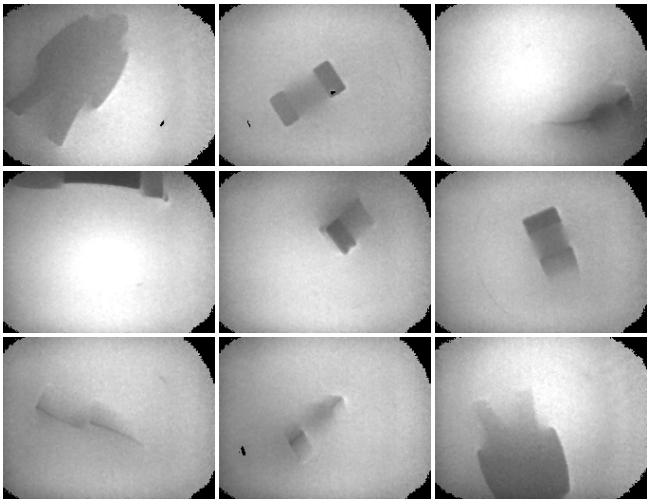

Fig. 8: Contact geometry depth images for all blocks were recorded in a wide range of angles, including partial views of local geometry, as illustrated by the selection of robot block contact geometry depth images above.

To gather data, we recorded depth image streams from the soft-bubble while each of the six objects were pressed into the bubble's membrane varying contact location, object orientations, and contact force (Fig. 7). We aimed to record contact geometry data for as wide a range of poses as possible, including partial views of local features, as illustrated by the selection of robot block depth images in Fig. 8. Depth image streams were also recorded for the "no-touch" case. The various data was collected using two separate soft-bubble assemblies with slight differences between them. The internal pressures of the bubbles were varied throughout to further diversify the data collected. These depth data streams for contacting objects were then passed to a filtering algorithm which discarded "no-touch" frames based-on a depth deviation threshold. For the "no-touch" stream, we simply sampled from all captures frames. Consequently, we collected 1,000 training images and 200 validation images for each of the seven classes. To train the network, cross-entropy loss and stochastic gradient descent were used as the loss function and optimizer, respectively. We used 0.1 for the learning rate with a 0.1 decay every 30 epochs, 0.9 for momentum, 10^{-4} for weight decay, and 32 as the mini-batch size. Note that the network was trained from scratch since publicly available pre-trained weights are usually trained using RGB images, differing from our choice of network input. The network was trained up to 80 epochs and achieved 98.14% average top-1 accuracy as the best performance throughout the training on our validation data. This result indicates that the soft-bubble can be successfully used for practical object classification tasks.

Next, we investigated the importance of input image resolution. In this experiment, we used two datasets; the previously detailed *six objects* dataset, and another one called

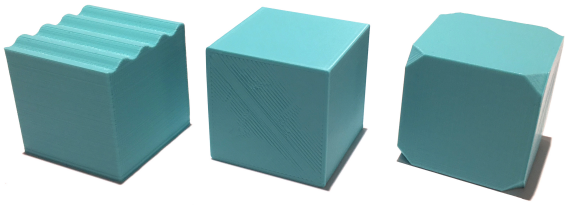

Fig. 9: Three cubes with differing surface features including a standard cube (middle), wavy face (left) and cut-off corners (right)

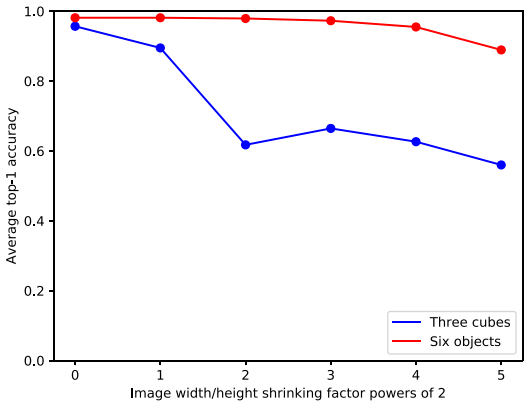

Fig. 10: Average Top-1 classification accuracy with varying image resolution

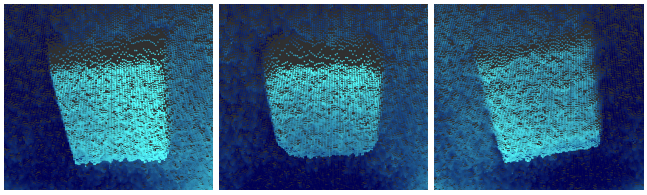

(a) Cube (b) Cut corners (c) Wave

Fig. 11: Full resolution point clouds of the cubes in Fig. 9.

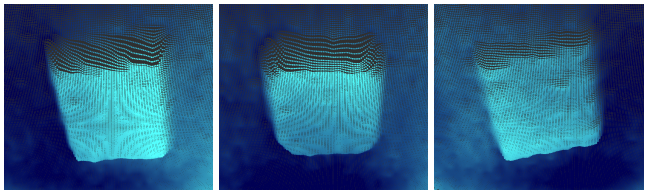

(a) Cube (b) Cut corners (c) Wave

Fig. 12: Lower resolution (length and width multiplied by 2^{-2}) point clouds of the *three cubes*. It is at this resolution that a drop-off in accuracy can be seen in Fig. 10.

the *three cubes* dataset which consists of four classes: "no-touch" and three types of cubic objects, congruent in size, but with unique surface features (Fig. 9). These cubic objects were chosen so that differences in shape between objects

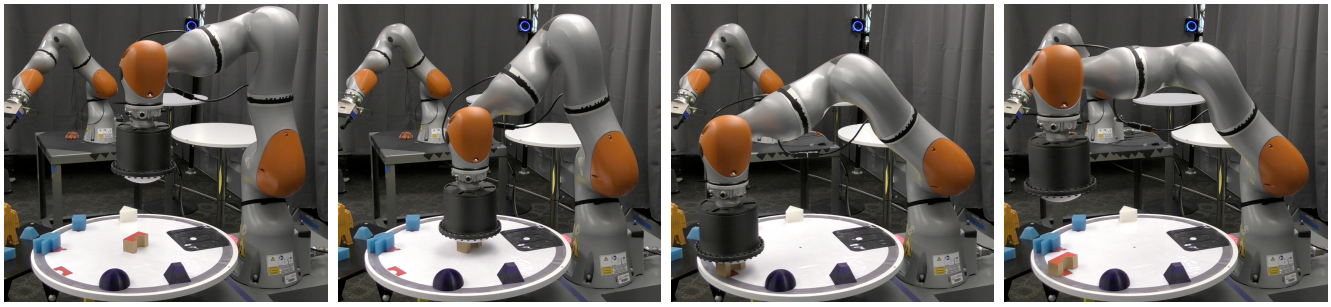

(a) (b) (c) (d)

Fig. 13: The robot positions itself above an object (a), presses against and classifies the object (b), pushes it along the table to the appropriate location (c) then retracts to repeat (d).

are subtle, especially as input image resolution is decreased. The *three cubes* dataset contains 1,000 training data and 500 validation images per category, gathered via the previously mentioned procedure. The same network architecture used before is again used for the *six objects* and *three cubes* datasets. The network is trained from scratch and seperately on *six objects* and *three cubes* datasets. The training is repeated six times for each dataset, each with varying image resolutions; the width and height of the input image is first multiplied by 2^{-N}, where N is a hyperparameter which varies from 0 to 5 over each training cycle, then is resized to 224×224 pixels to fit the input size of our network. We measured average top-1 accuracy as the metric for this experiment. The results are shown in Fig. 10. This illustrates that the accuracy is reasonably good at higher resolutions but drops as input image resolution decreases. This tendency is especially evident on the *three cubes* dataset where the differences between each object are more subtle than those of the *six objects* dataset. The accuracy for the three cubes with varying surface features drops significantly at $N = 2$. The point clouds corresponding to the original resolution and quartered resolution where $N = 2$ are shown in Figs. 11 and 12, respectively. From these results, we can argue that the higher the resolution, the better the accuracy will be for classification tasks. The results also suggest that tactile resolution must be carefully chosen based on task demands.

B. Robot Object Sorting

We integrated the soft-bubble, the six object classifier from the first experiment and a robot to demonstrate the ability to perform a simple but real world task. The soft-bubble is attached to the end effector of a Kuka IIWA robot arm, shown in Fig. 1a. We put each of the six blocks from the first experiment down on the center of a table in front of the robot. The task for the robot is to reach down, touch the blocks, and classify each object using the tactile information. Once identified by the object classifier, the robot pushes the object toward a preassigned area on the table. All of the robot's motions, i.e. pushing down, sliding objects and going back to the initial pose, are scripted. Classification using tactile information is performed between the pushing down and sliding away motions. Fig. 13 illustrates the robot's execution of the sorting task.

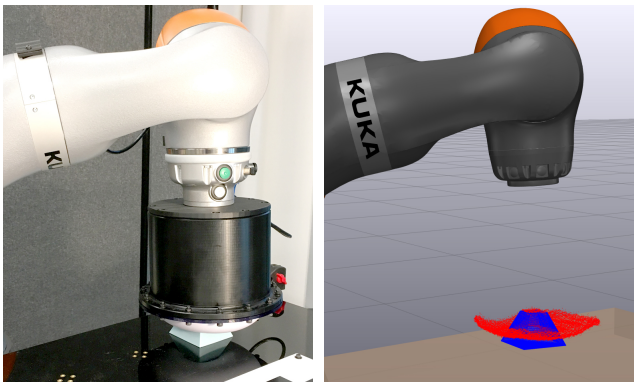

(a) Robot touches object (b) ICP result

Fig. 14: Pose estimation using the soft-bubble. Optimal pose of the pyramidal-frustum object rendered in blue along with the point-cloud from the soft-bubble in red.

C. Pose Estimation

As a third demonstration of the capabilities of the soft-bubble, we explore pose estimation of known objects using the depth image captured by the sensor. To estimate the pose, we used the well known Iterative Closest Point (ICP) algorithm [27].

The extrinsic of the pico flexx camera with respect to the end effector was computed from the CAD diagram of the soft-bubble assembly and was used for transforming the captured point-cloud into the appropriate world frame (base of the Kuka robot). For the purposes of this experiment, the Kuka was commanded to execute various Cartesian trajectories to reach out and touch the object in question. Each of these trajectories resulted in a penetration of up to 4cm (Fig. 14a). Over the course of the experiment, configurations of the robot and point clouds from the sensor were logged.

The ICP implementation of PCL [28] was utilized and the optimal pose was computed from a set of 12 initial orientations of the object in question. For each initial orientation, the lower 25% of the model was cropped out

in order to capture only the visible surfaces for the ICP computation. The image in Fig. 14b shows the converged result for the pose of the pyramidal frustum depicted as the object model superimposed onto the measured point-cloud showing near perfect correspondence. The rate of ICP computation was approximately 5Hz. It can clearly be seen that the pose of contacted objects with simple geometries can be computed using this sensor. More detailed analysis of the pose accuracy, effect of complicated geometries and sensor resolution are ongoing.

V. DISCUSSION AND OUTLOOK

As mentioned in Section III, the large form factor of the tactile sensor presented here is dependent on the specifications of the chosen depth sensor. While we used one of the smallest and most short-range-capable depth sensors on the market, a smaller depth sensor with a shorter sensing range and a wider field of view is needed in order to widely deploy the soft-bubble on grippers and space-limited parts of a robot. Not only would a smaller depth sensor lead to smaller tactile sensors, it also leads to interesting possibilities for larger scale bubbles with membranes so large that multiple synced depth sensors could work together to produce a fused point cloud. Future versions of this sensor will be located on various parts of a robot, not only end effectors.

The soft-bubble currently senses geometry only. Future work includes static and dynamic modeling of the soft mechanics so that contact pressure on the membrane may be estimated based on deformation. With the addition of dots or other trackable features on the inner surface of the membrane, shear forces and moments can be estimated as well. Modeling will also allow the sensor contact mechanics and output to be simulated. Methods will be developed for calibrating the tactile sensor's depth output, as well as for quantifying measurement error and sensor noise.

In Section IV, we show that classification and pose estimation tasks can be performed using depth images from the bubble and argue that image resolution affects the classification accuracy, especially when the differences in objects shape or surface features are subtle. We believe that this is true for other tasks involving tactile sensing like object pose estimation, although more work needs to be done to verify this. Future experiments include dual arm manipulation, tactile-based pose refinement and exploratory techniques for manipulating occluded and hard-to-see objects.

REFERENCES

[1] C. C. Kemp, A. Edsinger, and E. Torres-Jara, "Challenges for robot manipulation in human environments [grand challenges of robotics]," *IEEE Robotics & Automation Magazine*, vol. 14, no. 1, pp. 20–29, 2007.

[2] M. Bjorkman, Y. Bekiroglu, V. Hogman, and D. Kragic, "Enhancing visual perception of shape through tactile glances," in *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*. IEEE, 2013, pp. 3180–3186.

[3] A. Jain, M. D. Killpack, A. Edsinger, and C. C. Kemp, "Reaching in clutter with whole-arm tactile sensing," *The International Journal of Robotics Research*, vol. 32, no. 4, pp. 458–482, 2013.

[4] H. Abidi and M. Cianchetti, "On intrinsic safety of soft robots," *Frontiers in Robotics and AI*, vol. 4, p. 5, 2017. [Online]. Available: http://www.frontiersin.org/article/10.3389/frobt.2017.00005

[5] B. S. Homberg, R. K. Katzschmann, M. R. Dogar, and D. Rus, "Robust proprioceptive grasping with a soft robot hand," *Autonomous Robots*, pp. 1–16, 2018.

[6] M. Rus, Daniela & Tolley, "Design, fabrication and control of soft robots," *Nature*, vol. 521, no. 6, pp. 467–475, 2015.

[7] J. Kim, A. Alspach, and K. Yamane, "3d printed soft skin for safe human-robot interaction," in *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*. IEEE, 2015, pp. 2419–2425.

[8] J. R. Amend, E. Brown, N. Rodenberg, H. M. Jaeger, and H. Lipson, "A positive pressure universal gripper based on the jamming of granular material," *IEEE Transactions on Robotics*, vol. 28, no. 2, pp. 341–350, April 2012.

[9] S. R. Inc. (2018) Core technology. [Online]. Available: http://www.softroboticsinc.com/core-technology

[10] S. Haddadin, A. De Luca, and A. Albu-Schäffer, "Robot collisions: A survey on detection, isolation, and identification," *IEEE Transactions on Robotics*, vol. 33, no. 6, pp. 1292–1312, 2017.

[11] B. Siciliano and L. Villani, *Robot force control*. Springer Science & Business Media, 2012, vol. 540.

[12] M. Meier, M. Schopfer, R. Haschke, and H. Ritter, "A probabilistic approach to tactile shape reconstruction," *IEEE Transactions on Robotics*, vol. 27, no. 3, pp. 630–635, 2011.

[13] R. Calandra, A. Owens, M. Upadhyaya, W. Yuan, J. Lin, E. H. Adelson, and S. Levine, "The feeling of success: Does touch sensing help predict grasp outcomes?" in *Proc. of the 1st Annual Conference on Robot Learning (CoRL)*, 2017.

[14] W. Yuan, S. Dong, and E. H. Adelson, "Gelsight: High-resolution robot tactile sensors for estimating geometry and force," *Sensors*, vol. 17, no. 12, p. 2762, 2017.

[15] G. Izatt, G. Mirano, E. Adelson, and R. Tedrake, "Tracking objects with point clouds from vision and touch," in *Robotics and Automation (ICRA), 2017 IEEE International Conference on*. IEEE, 2017, pp. 4000–4007.

[16] J. T. S. Li, S. Lyu and W. Burgard, "A comparative study of contact models for contact-aware state estimation," in *IEEE International Conference on Intelligent Robots and Systems*, September 2015.

[17] Y. Bekiroglu, J. Laaksonen, J. A. Jorgensen, V. Kyrki, and D. Kragic, "Assessing grasp stability based on learning and haptic data," *IEEE Transactions on Robotics*, vol. 27, no. 3, pp. 616–629, 2011.

[18] J. Tegin and J. Wikander, "Tactile sensing in intelligent robotic manipulation–a review," *Industrial Robot: An International Journal*, vol. 32, no. 1, pp. 64–70, 2005.

[19] R. S. Dahiya, G. Metta, M. Valle, and G. Sandini, "Tactile sensingfrom humans to humanoids," *IEEE transactions on robotics*, vol. 26, no. 1, pp. 1–20, 2010.

[20] E. Donlon, S. Dong, M. Liu, J. Li, E. Adelson, and A. Rodriguez, "Gelslim: A high-resolution, compact, robust, and calibrated tactile-sensing finger," *arXiv preprint arXiv:1803.00628*, 2018.

[21] A. Yamaguchi and C. G. Atkeson, "Implementing tactile behaviors using fingervision," in *2017 IEEE-RAS 17th International Conference on Humanoid Robotics (Humanoids)*, Nov 2017, pp. 241–248.

[22] L. Zhang and J. C. Trinkle, "The application of particle filtering to grasping acquisition with visual occlusion and tactile sensing," in *Robotics and automation (ICRA), 2012 IEEE international conference on*. IEEE, 2012, pp. 3805–3812.

[23] M. C. Koval, N. S. Pollard, and S. S. Srinivasa, "Pose estimation for planar contact manipulation with manifold particle filters," *The International Journal of Robotics Research*, vol. 34, no. 7, pp. 922–945, 2015.

[24] P. Technologies. (2018) Camboard pico flexx. [Online]. Available: http://pmdtec.com/picofamily/flexx/

[25] M. Inc. (2018) Markforged x7. [Online]. Available: http://markforged.com/x7/

[26] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.

[27] G. K. Tam, Z.-Q. Cheng, Y.-K. Lai, F. C. Langbein, Y. Liu, D. Marshall, R. R. Martin, X.-F. Sun, and P. L. Rosin, "Registration of 3d point clouds and meshes: a survey from rigid to nonrigid." *IEEE transactions on visualization and computer graphics*, vol. 19, no. 7, pp. 1199–1217, 2013.

[28] R. B. Rusu and S. Cousins, "3d is here: Point cloud library (pcl)," in *Robotics and automation (ICRA), 2011 IEEE International Conference on*. IEEE, 2011, pp. 1–4.