

Haptic Contour Following with the Smart Suction Cup

Sebastian D. Lee, Jungpyo Lee, and Hannah S. Stuart

Abstract—The Smart Suction Cup is a tactile sensing and gripping system designed to enhance pick-and-place operations in industrial settings. While previous research has primarily focused on utilizing this technology for haptic search in cases of initial grasp failure, this study introduces a novel application: following contours. This function is already established as an important function for object recognition and grasp planning – substantiated by numerous works using other tactile sensors. Here, we explore contour following for a flow-based tactile sensor because it is not susceptible to visual occlusions nor tactile sensor wear. Experimental validation demonstrates the Smart Suction Cup’s ability to track edges at different speeds and navigate various planar contours, showcasing rapid and robust tracking of edges. Notably, the Smart Suction Cup can reliably operate at a speed of 3 cm/s. This is one step towards the adoption of the Smart Suction Cup for real-world applications.

I. INTRODUCTION

Suction cups play a crucial role in industry for pick-and-place tasks. Their ability to astrictively grip and manipulate objects of various shapes and sizes makes them valuable tools for manufacturing and logistics operations today. However, traditional suction cups typically lack tactile sensing, limiting their adaptability and effectiveness in semi-structured and cluttered environments, such as e-commerce warehouses. Thus, prior efforts equip suction cups with sensing abilities.

Doi et al. (2020) proposed a 3-electrode capacitive sensor, which could localize contact [1]. Shahabi et al. (2023) embedded four microfluidic, carbon grease channels into a single suction cup, and demonstrated surface slope estimation [2]. Aoyagi et al. (2020) dip-coated a 2x2 array of suction cups with a thin-film polymer [3] and demonstrated real-time edge detection to trigger a regrasp. However, contour following with a sensorized suction cup has not yet been demonstrated. We seek to address this gap with the Smart Suction Cup.

The Smart Suction Cup was first introduced by Huh et al. (2021) [4]. We then introduced a haptic search method for sliding and rotating control strategies for grasping adversarial objects [5]. Implementing this method, which identifies viable grasping points when the initial grasp point fails, resulted in a 2.5x improvement in grasping during cluttered bin-picking trials. Our latest hardware is summarized in [5] and tested on PCB’s for industrial recycling applications. Tactile sensors often employ localized force or geometry transducers near the contact point. One relative advantage

S. D. Lee, J. Lee and H. S. Stuart are with the Embodied Dexterity Group, Dept. of Mechanical Engineering, University of California Berkeley, Berkeley, CA 94720 USA. (e-mail: sebastiandavidlee@berkeley.edu; jungpyolee@berkeley.edu; hstuart@berkeley.edu)

Email coorespondence to hstuart@berkeley.edu.

There is a video supplement associated with this work.




Fig. 1. The Smart Suction Cup robotic system showing: a) suction cup mounted on a robot arm, b) close-up of suction cup hardware, and c) bottom view of the suction cup showing the tool frame.

of the Smart Suction Cup hardware is that its sensitive electronics are remote from the point of contact, which reduces physical damage to the transducers and degradation of the tactile signals. Though we previously demonstrate adaptive grasping capabilities with this technology, purposeful contour following has yet to be realized.

Contour following is employed by people when interacting with the physical world as an integral part of haptic exploration. Edges are particularly salient features for humans [6], and guide the recognition of object shapes [7]. Lepora et al. (2017) demonstrate contour following with the TacTip using model-based [8] and model-free [9] approaches, the latter being more robust to complex, soft, and irregular objects. They also demonstrate this ability with DIGIT [10] and DigiTac [11]. With their 12-taxel, capacitive sensor, Martinez et al. (2013) demonstrate object shape classification with contour following [12]. These works show that robots can leverage contour following using tactile information alone in ways that are useful for automated dexterous manipulation.

A. Overview

To the best of the author’s knowledge, we are presenting the first case of contour following using a suction cup gripper. Section II provides a description of the contour-following controller design. In Section III, we summarize the hardware (Fig. 1) and describe controller bench-marking methods. In practice, speed and versatility of a tactile sensing system are both desirable, though these performance metrics can be at odds. The current study evaluates the Smart Suction

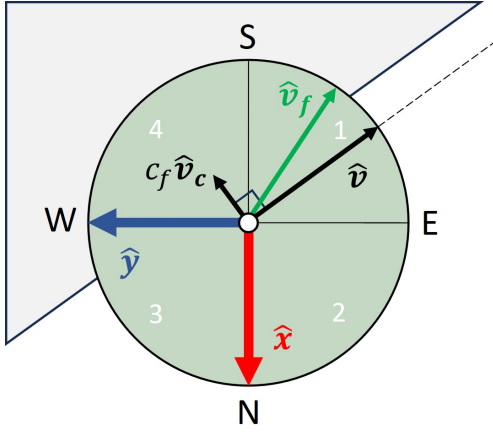


Fig. 2. Top view diagram of the Smart Suction Cup on a straight edge.

Cup's ability to track an edge at different speeds as well as the ability to navigate different planar contours. Section IV presents the experimental results. We find that at low speeds many controller designs achieve contour following on a disk, but at higher speeds success rate is more sensitive to controller parameters. We also observe successful contour following on compound geometries and filleted edges. In Section V, we discuss the effects of parameters on success envelopes, reliability and speed, and observed failure modes. We conclude in Section VI that rapid and robust contour following with the Smart Suction Cup is possible.

II. CONTROLLER DESIGN

To achieve contour following, the controller first determines the direction in which to move. For each loop iteration, the pressure differential across the chambers, ΔP , is computed for two orthogonal directions of the cup's gripping surface, \hat{x} and \hat{y} , as:

$$\Delta P_y = P_W - P_E \quad (1a)$$

$$\Delta P_x = P_N - P_S \quad (1b)$$

where P is the vacuum pressure relative to ambient pressure, and the subscripts represent cardinal direction or chamber number, e.g., P_1 is the vacuum pressure from chamber 1, and $P_S = \frac{P_1 + P_4}{2}$, while $P_N = \frac{P_2 + P_3}{2}$. With these pressure differentials, we can compute the direction vector as:

$$\mathbf{v} = -(\Delta P_x \hat{x} + \Delta P_y \hat{y}) \quad (2)$$

where \hat{x} and \hat{y} are fixed in the tool frame. In this work, we use only the direction or unit vector \hat{v} . We also define a correction unit vector, \hat{v}_c , which is orthogonal to \hat{v} . The final vector, \mathbf{v}_f , is a linear combination of \hat{v} and \hat{v}_c :

$$\mathbf{v}_f = \hat{v} + c_f \hat{v}_c \quad (3)$$

where c_f is the correction factor. Similarly, we normalize to compute unit vector, \hat{v}_f . Fig. 2 graphically shows the relationship between these vectors for a particular instance, assuming that \hat{v} has no error. The mean pressure across the four chambers is computed as,

$$\bar{P} = (P_N + P_S + P_W + P_E)/4. \quad (4)$$




Fig. 3. Correction factor as a function of pressure. When $P < P_{ref}$, the correction factor is positive. When $P > P_{ref}$, the correction factor is negative. Variables in blue are varied in our experiments, while variables in black are fixed.

We allow the correction factor magnitude, $|c_f|$, to vary linearly with the log of the mean pressure, $\log_{10}(\bar{P})$. The idea is that \bar{P} is related to how far away or into an edge the cup: \bar{P} varies more gradually when the cup is fully open (\bar{P} is low) and more dramatically when the cup is mostly covered (\bar{P} is high) [5].

Fig. 3 illustrates the piece-wise function that determines c_f . To reduce the number of experimental parameters, values for P_n and P_g were set to 10 Pa and 10 kPa, which are approximately the noise level of \bar{P} and the pressure required to grasp an object, respectively. In Fig. 3, $P_{ref} = 100$ Pa and $\beta = 1$, though these factors will be varied in experiments.

III. EXPERIMENTAL METHODS

A. Robotic System

1) *Tactile Sensor*: We use a pressure-based tactile sensing suction cup [4]. Fabrication details can be found in our prior work [5]. The latest version of the hardware uses a 3D-printed enclosure (Fig. 1b), which houses four pressure sensors (MPRLS, Adafruit), an I2C multiplexer (PCA9546, Adafruit), and a microcontroller (ESP32-S3 feather, Adafruit). The Smart Suction Cup is split into four chambers (Fig. 1c). In total, four pressure signals characterize flow differentials across the chambers in the cup and are enough to localize leakage and desired motion direction.

2) *Robot Arm*: The Smart Suction Cup serves as the end-effector. Its base is bolted to a 6-DoF F/T transducer (Axia80, ATI), which is mounted on a 6-DoF robot arm (UR-10, Universal Robots) as shown in Fig. 1a. Transducer normal force is measured and utilized to initialize experiments when the cup first approaches the object to achieve 1.5 N of compression. The arm then controls horizontal motion while maintaining a constant vertical height assuming a flat planar object; the load cell is not used during real-time contour-following control. In this study, the experiments are conducted using a desktop computer running on Ubuntu 20.04 and ROS (Noetic).




Fig. 4. Experimental setup for contour following on a disk. The cup's position is initialized to the center of the disk to ensure consistency across trials. The trials starts after the cup moves to the edge.

B. Controller Benchmarking on a Disk

In order to benchmark the performance of our controllers, contour-following is tested on a 35 mm disk, which enables testing in all orientations of the suction cup (Fig. 4). A trial is initialized at the center of the disk, which ensures consistency across trials. The cup is then moved to the edge of the disk, 32 mm away from the disk's center to ensure sufficient coverage of the cup's chambers. Contour following is arbitrarily chosen to be in the counter-clockwise direction.

We test the effect of three parameters: P_{ref} , β , and step size ΔL . Cup speed is controlled by the step size, ΔL . We expect P_{ref} should correspond to the offset of the edge relative to the suction cup, while β should correspond to how aggressively the cup corrects. By definition, $|c_f| \leq \beta$, which implies that, when $\beta = 1$, the maximum magnitude of $c_f \hat{v}_c$ is equal to that of \hat{v} . This also means the angle between them could be a maximum of 45° . For each triplet of $(P_{ref}, \beta, \Delta L)$ three trials were run and success rate recorded, where a successful trial is defined as completion of one full circumference. We extract the path lengths, mean radii, and standard deviations of the resulting trajectories and compare them to their respective controller parameters.

C. Testing on Other Objects

After selecting one promising controller found in our disk benchmarking experiments, we run contour-following trials on different types of geometries to assess the robustness of our system. First, we select primitive planar geometries designed to test the limits of the suction cup's contour-following capability. We initially test disks and holes of differing diameters, from 2 to 14 mm in steps of 2 mm (Fig. 5a). Since the cup is able to follow the edge of a 35 mm disk, we are most interested in characterizing the minimum diameter disk and hole.

We then test discontinuous corners between straight edges of varying internal (concave) and external (convex) angles, α , from 0 to 360° (Fig. 5b). Angles close to 180° resemble flat edges. Such discontinuities with no connecting curvature can generate signals that yield failures. We then test three compound shapes with varying continuous and/or discontinuous features to test how the prior results translate to more complex shapes. Finally, we perform initial trials on a simple three dimensional filleted edge, where the fillet gradually increases from a zero radius of curvature to 16 mm radius of curvature; we can observe how sharp the edge must be for our method to work.




Fig. 5. Disks, holes, and corners. Primitive planar geometries can be characterized by diameter or angle.

IV. RESULTS

A. Effect of Controller Parameters on Trajectory

Here, we show and analyze closed-loop trajectories to illustrate the effects of P_{ref} and β . Varying P_{ref} while keeping a constant $\beta = 1.5$ (Fig. 6a) shows that a higher $P_{ref} = 200$ Pa value biases the trajectory towards the disk's center. On the other hand, $P_{ref} = 40$ Pa roughly centers the cup on the edge. In order to better understand the effect of P_{ref} on postural bias, we compute each trajectory's mean radius \bar{r} and compare it to the disk's radius in Fig. 6a. Trajectory radii are graphically represented by the dashed circles. Mean radius offset, Δr , is the difference between \bar{r} and the radius of the disk, and it is plotted across three trials under each condition in Fig. 6c. Individual data points and error bars representing one standard deviation are plotted for each set of trials. At $P_{ref} = 40$ Pa, \bar{r} is 34.3 mm, and $\Delta r = -0.7$ mm. When $P_{ref} = 200$ Pa, Δr grows to -3.2 mm. As P_{ref} increases, \bar{r} decreases, resulting in the cup covering more of the disk.

Varying β while keeping a constant $P_{ref} = 100$ Pa (Fig. 6b), shows that this parameter more so affects the shape of the trajectory than the edge bias; more obvious oscillation appears as β increases from 1.5 to 3.5. We characterise these oscillations using radial standard deviation σ_r , which is the standard deviation of the trajectory relative to \bar{r} , and travel distance s , which is the trajectory's path length. At $\beta = 1.5$, $\sigma_r = 1.36$ mm and $s = 244$ mm, which is 111% of the disk's circumference, C , i.e. relative travel $s_c = 11\%$. At $\beta = 3.5$, s_c increases to 31% and σ_r increases to 1.89 mm. s_c and σ_r are plotted against β in Fig. 6d, which shows the average across three trials at each β . Relative travel appears to increase non-linearly with more variability at higher values of β .

B. Disk Contour Following Speed and Reliability

Through pilot studies we find that, when $P_{ref} = 70$ Pa and $\beta = 1.0$, the controller yields successful contour following at the step size of $\Delta L = 2.5$ mm. We consider $\Delta L = 2.5$ mm to be our initial "low" speed and seek to improve speed by tuning the other parameters. Figure 7 summarizes how the P_{ref} and β parameters are swept across step sizes ΔL of 2.5 mm ("low"), 5.0 mm ("intermediate"), and 7.5 mm ("high") speeds. P_{ref} is varied to six unique values between

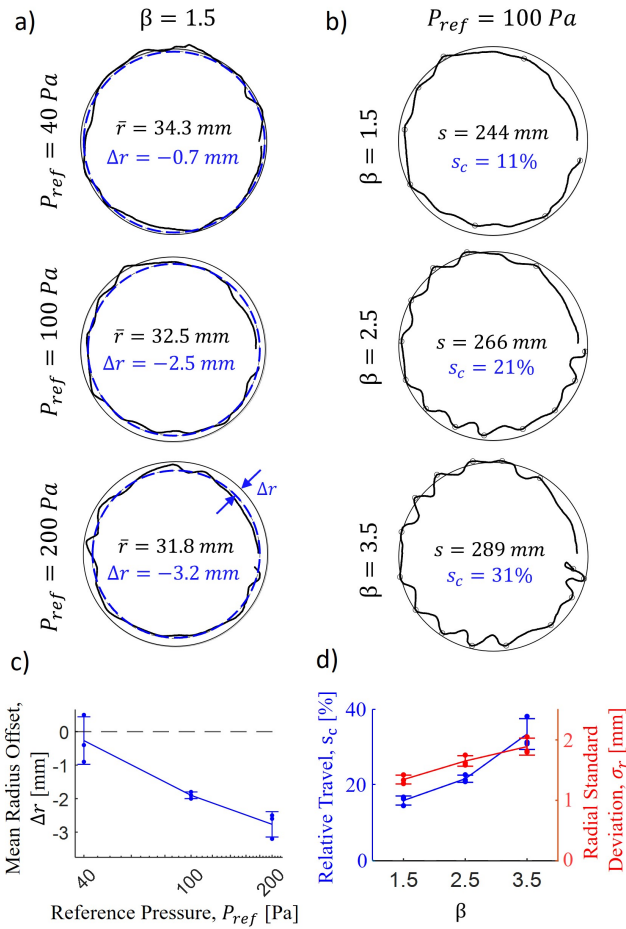


Fig. 6. Effect of controller parameters on closed-loop control trajectories. Three trajectories are plotted for two different conditions. a) P_{ref} is varied while keeping β fixed at 1.5. b) Similarly, β is varied while keeping P_{ref} fixed at 100 Pa. c) Data points and error bars representing one standard deviation are plotted for trajectory metrics. As P_{ref} increases, the mean radius offset Δr (difference between disk radius and \bar{r}) decreases, resulting in trajectories where the cup is closer to the center of the disk. d) As β increases, both relative travel s_c and radial standard deviation σ_r increase.

the range of [40, 300] Pa, while β could be any multiple of 0.5 within the range [0, 5.0]. Success rate is reported for each triplet tested (P_{ref} , β , ΔL). To reduce the number of experiments, a sparse search is conducted; pairs of (P_{ref} , β) are strategically chosen to identify the successful envelope of controller parameters for each speed, which is contoured by low success rates (i.e., the boundary of X's in Fig. 7a-c).

The parameter space at the low step size of $\Delta L = 2.5$ mm (~ 1 cm/s) yields high success rates for most combinations of P_{ref} and β . When $\beta > 3.0$, success rate is 100% for all values of P_{ref} . Due to the high reliability of contour following at this speed, the success rate of many parameter combinations are inferred from the sparse search. At the tested step sizes, a $\beta = 0.5$ or less always results in failure. Success rate is lower towards the bottom left corner at this speed; here, the trials terminate due to the cup moving out and away from the edge (Fig. 8, Failure Mode I).

The parameter space at the intermediate step size of $\Delta L = 5.0$ mm (~ 2 cm/s) shows a reduction in the size of the success envelope compared to the low step size condition.

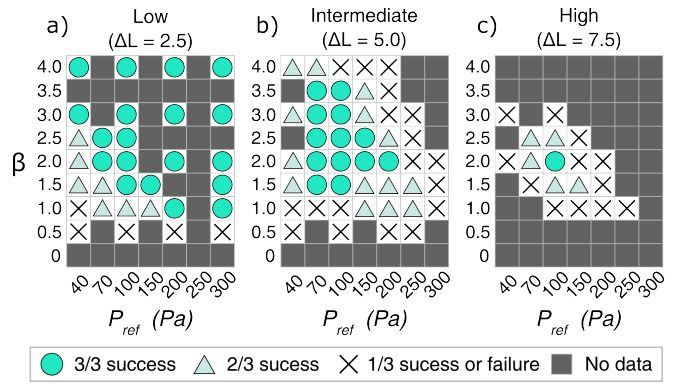


Fig. 7. Success rate maps at three different step sizes: a) 2.5, b) 5.0, and c) 7.5. At each (P_{ref} , β , ΔL) triplet, three attempts were made to follow the edge of a 35 mm disk. Circles represent 100% success rate, triangles represent 66% success, and X's represent low or no success rate. Gray boxes represent points in the parameter space that were not tested.




Fig. 8. Failure Modes. I) Out from the disk center and away from the edge. II) Into the disk center and away from the edge.

Success rate notably decreases towards the top right corner at pair ($P_{ref} = 200$ Pa, $\beta = 3.0$); trials in this regime terminate due to the cup moving too far into the surface, resulting in the cup grasping the surface (Fig. 8, Failure Mode II).

At the highest step size condition, $\Delta L = 7.5$ mm (~ 3 cm/s), we find that the controller can reliably edge-follow only when $P_{ref} = 100$ Pa and $\beta = 2.0$, i.e. it succeeds in all three trials. These P_{ref} and β settings yield successful contour following at the three tested step sizes, shown in the Supplementary Video. With different P_{ref} and β settings, the controllers at the high step size yield at least one failure. Failures occur both inside and outside of the object (Fig. 8a-b). These failure modes occur when \bar{P} falls below the noise level $P_n = 10$ Pa or exceeds $P_g = 10$ kPa, rendering an accurate measurement of directionality impossible or leading to full engagement of the surface, respectively.

C. Performance on Other Planar Geometries

Due to the cup's structure and size, there are limits to how small or large a feature can be before the cup is unable to track its edge. To determine the geometric limits of the cup, contour following is tested on disks and holes of decreasing size as well as sharp internal and external corners. $\Delta L = 0.5$ mm while $P_{ref} = 100$ Pa and $\beta = 2.0$. We find that the cup is able to track a 12 mm diameter disk, but a 10 mm disk proves challenging. The cup is able to track a 2 mm hole, the smallest hole tested. These same controller parameters are tested on convex and concave corners. We find that the suction cup performs well on the edge of external corners with angles α of 15° or greater. However, the cup is unable to successfully follow internal corners greater than 190° .




Fig. 9. Successful contour following on compound geometries. Each trajectory is placed below each test object and scaled to the same size. Dashed circles represent the size of the suction cup relative to the shapes.

Beyond primitive features, arcs and corners can be concatenated to form compound geometries. Contour following is tested on three distinct compound geometries: a volute, decreasing semi-circles, and decreasing semi-squares. We find that the Smart Suction Cup successfully follows the edges of these three compound geometries, as seen in Fig. 9, and shown in the Supplementary Video. However, we note that the exact shape of the robot’s trajectory is slightly different from the object’s shape – this is expected given trajectories observed on earlier disk trials.

D. Filleted Edge

We also explore the Smart Suction Cup’s ability to follow a filleted edge. To do so, we run contour following on a variable fillet edge (Fig. 10). The trials start on the non-filleted portion of the edge and proceed to follow the filleted edge. Cross-sections at various distances along the fillet are illustrated in Fig. 10a. The maximum fillet radius is 16 mm.

Across five trials, top view trajectories show that the cup’s trajectories stays on top of the filleted portion of the edge, which eventually leads to the cup moving out and away from the edge (Fig. 10b). The pressure differential drops to zero between 120 mm and 160 mm, with a mean of 137.5 ± 12.7 mm. Therefore, the Smart Suction Cup appears able to follow a fillet with a radius of about 14.5 mm or less.

V. DISCUSSION

A. Understanding Performance on the Disk

Through a parameter sweep, we observe success envelopes, as in Fig. 7, their size decreasing with larger steps. Shrinking originates from the bottom left and top right corners, where both P_{ref} and β approach their tested limits. In the bottom left corner, decreasing P_{ref} increases \bar{r} (Fig. 6a), while decreasing β results in lighter oscillations (Fig. 6b); coupled together, they increase the likelihood of the cup veering out (Mode I, Fig. 8a). At low P_{ref} , a high β value improves success rate by providing corrections to

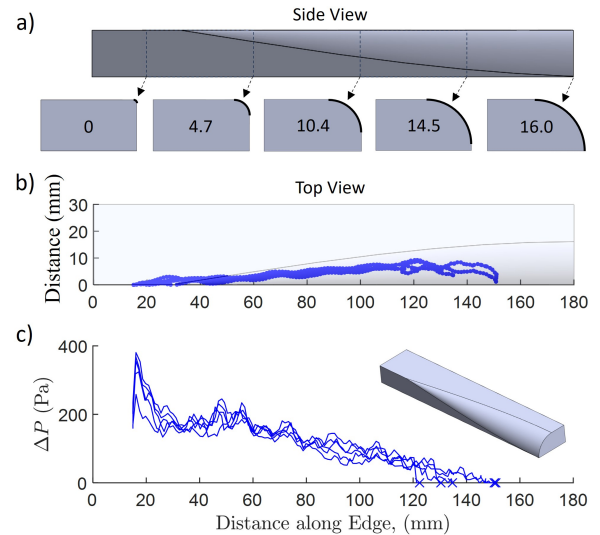


Fig. 10. Contour-following trials on a variable 3D fillet. a) Side view and fillet radius size at given cross-sections. b) Top view trajectory of five trials. c) Pressure differential data goes to zero as fillet size increases.

prevent Failure Mode I. However, at a high step size, β is unable to compensate quickly enough. Biasing the cup into the edge with an intermediate value of P_{ref} therefore helps to prevent Mode I. From the top right corner, the success envelope rapidly shrinks as step size increases, often terminating in Failure Mode II (Fig. 8b). As \bar{r} decreases with higher P_{ref} , the cup biases towards the disk center (Fig. 6a), while higher β yields stronger oscillations (Fig. 6b), increasing trajectory amplitude and the likelihood of over-correcting into the object.

The effects of P_{ref} , β , and ΔL are likely coupled in ways not captured in the present study, warranting further examination. This could also support the development of improved controller methods in future work. For example, to prevent Failure Mode II, vacuum Pulse Width Modulation (PWM) can be triggered when \bar{P} exceeds a threshold, maintaining flow differential without surface grasping, as introduced in [4]. Alternatively, adaptive controller parameters show promise, like reducing step size to slow down the arm speed when \bar{P} is especially high.

B. Speed and Reliability

We find that the continuous-sliding contour-following policy can yield reliable (3/3 trial successes) and rapid (~ 3 cm/s) performance on a disk. Contour following has been demonstrated at various speeds in prior literature. For instance, the TacTip completes one full revolution in 112 s around a 105 mm diameter disk, equivalent to 3.0 mm/s (7.5% \varnothing/s relative to its 40 mm body diameter) [9]. A 12-taxel, capacitive-based tactile sensor reaches 0.7 mm/s [12]. In contrast, our Smart Suction Cup reaches a top speed of 30 mm/s, about 130% \varnothing/s relative to its size, $\varnothing_{out} = 23$ mm (Fig. 1b). This performance difference may stem from factors like reduced computation time and increased sampling rate afforded by a lower-resolution sensor. Lepora et al., (2022) benchmark DIGIT’s [10] and DigiTac’s [11]




Fig. 11. Cup lip crimping. The sequence of snapshots shows how the cup’s lips fold in, resulting in flow blockage that misdirects the cup’s motion.

contour-following ability but do not report trial completion times¹. Using common benchmarking methods, that include execution speed, for standardizing contour following assessment would assist in comparing performance in future work.

C. Robustness to Various Geometries

The control policy tested appears robust, as it generalizes to various object geometries. However, one failure mode arises when attempting to navigate internal corners with $\alpha > 190^\circ$: the lip wedges itself into the corner and then folds in on itself, covering the chambers (Fig. 11). Consequently, the cup likely cannot fully circumnavigate shapes like a “clover” or a “heart” due to their concave corners. One solution is to detect this failure mode and switch to a hopping controller, as proposed in [5]. Another failure mode occurs with a 10 mm disk, smaller than the cup’s internal diameter, $\varnothing_{in} = 11$ mm (Fig. 1c). Due to substantial flow leakage, \bar{P} is insufficient to accurately measure or grasp most objects. Flow leakage diminishes ΔP , leading to less accurate estimation of \hat{v} .

When looking at the compound geometries, our method succeeds on the “decreasing semi-circles” compound object, despite its inability to track the 10 mm disk. This may be because, as observed in Fig. 9 for the “decreasing semi-circles” object, the controller disregards and slides past the smallest arc, at the top right corner of the object. The cup behaves similarly with the smallest feature of the “decreasing semi-squares.” We also note a distortion effect for almost all object corners, where the cup creates geometric artifacts. This is likely due to unmodeled flow conditions that change as the covered area of the cup varies widely at these intersections. However, this effect does not appear to disrupt the overall contour following task completion and is left as an area for future work.

The cup is able to track a filleted edge as well, but fails as the fillet radius increases. In this study, we keep vertical height constant while moving horizontally. As a result, larger radius fillets lead to a lower ΔP , making it more challenging to estimate \hat{v} . Regulating the applied contact force may partially address this issue, since a higher contact force could increase overall flow blockage and result in a more noticeable flow differential. We plan to study Smart Suction Cup contour following on even more complex 3D geometries in future work.

¹These studies do not necessarily test for top speed. The comparison here serves as a starting point for benchmarking performance in terms of task completion time.

VI. CONCLUSION

This work is the first demonstration of contour following with the Smart Suction Cup. Through parametric studies, we tuned a controller that enables rapid and reliable contour following around a disk. The controller is robust to simple planar geometries, including arcs and corners, as well as compound shapes and filleted edges. Real-world objects have complex 3D geometries and variable friction. Friction in particular could cause large deformations in shear, disrupting airflow and altering the results. Object-specific controller adjustments, and contour following for object recognition, should be considered as future work. Future work should also address challenges in vision-obscured or high-uncertainty scenarios, such as in opaque or cluttered enclosures.

ACKNOWLEDGMENT

This work was supported in part by InnoHK of the Government of the Hong Kong Special Administrative Region via the Hong Kong Centre for Logistics Robotics (award #049151) and by the University of California, Berkeley.

REFERENCES

- [1] S. Doi, H. Koga, T. Seki, and Y. Okuno, “Novel proximity sensor for realizing tactile sense in suction cups,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 638–643.
- [2] E. Shahabi, F. Visentin, A. Mondini, and B. Mazzolai, “Octopus-inspired suction cups with embedded strain sensors for object recognition,” *Advanced Intelligent Systems*, vol. 5, no. 2, p. 2200201, 2023.
- [3] S. Aoyagi, M. Suzuki, T. Morita, T. Takahashi, and H. Takise, “Bellows suction cup equipped with force sensing ability by direct coating thin-film resistor for vacuum type robotic hand,” *IEEE/ASME Transactions on Mechatronics*, vol. 25, no. 5, pp. 2501–2512, 2020.
- [4] T. M. Huh, K. Sanders, M. Danielczuk, M. Li, Y. Chen, K. Goldberg, and H. S. Stuart, “A multi-chamber smart suction cup for adaptive gripping and haptic exploration,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 1786–1793.
- [5] J. Lee, S. D. Lee, T. M. Huh, and H. S. Stuart, “Haptic search with the smart suction cup on adversarial objects,” *IEEE Transactions on Robotics*, 2023.
- [6] M. A. Plaisier, W. M. Bergmann Tiest, and A. M. Kappers, “Salient features in 3-d haptic shape perception,” *Attention, Perception, & Psychophysics*, vol. 71, no. 2, pp. 421–430, 2009.
- [7] S. J. Lederman and R. L. Klatzky, “Hand movements: A window into haptic object recognition,” *Cognitive psychology*, vol. 19, no. 3, pp. 342–368, 1987.
- [8] N. F. Lepora, K. Aquilina, and L. Cramphorn, “Exploratory tactile servoing with active touch,” *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 1156–1163, 2017.
- [9] N. F. Lepora, A. Church, C. De Kerckhove, R. Hadsell, and J. Lloyd, “From pixels to percepts: Highly robust edge perception and contour following using deep learning and an optical biomimetic tactile sensor,” *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 2101–2107, 2019.
- [10] M. Lambeta, P.-W. Chou, S. Tian, B. Yang, B. Maloon, V. R. Most, D. Stroud, R. Santos, A. Byagowi, G. Kammerer *et al.*, “Digit: A novel design for a low-cost compact high-resolution tactile sensor with application to in-hand manipulation,” *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 3838–3845, 2020.
- [11] N. F. Lepora, Y. Lin, B. Money-Coomes, and J. Lloyd, “Digitac: A digit-tactip hybrid tactile sensor for comparing low-cost high-resolution robot touch,” *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 9382–9388, 2022.
- [12] U. Martinez-Hernandez, G. Metta, T. J. Dodd, T. J. Prescott, L. Natale, and N. F. Lepora, “Active contour following to explore object shape with robot touch,” in *2013 World Haptics Conference (WHC)*. IEEE, 2013, pp. 341–346.