

A Multi-Chamber Smart Suction Cup for Adaptive Gripping and Haptic Exploration

Tae Myung Huh¹, Kate Sanders², Michael Danielczuk², Monica Li¹, Ken Goldberg², Hannah S. Stuart¹

Abstract—We present a novel robot end-effector for gripping and haptic exploration. Tactile sensing through suction flow monitoring is applied to a new suction cup design that contains multiple chambers for air flow. Each chamber connects with its own remote pressure transducer, which enables both absolute and differential pressure measures between chambers. By changing the overall vacuum applied to this smart suction cup, it can perform different functions such as gentle haptic exploration (low pressure) and monitoring breaks in the seal during strong astrictive gripping (high pressure). Haptic exploration of surfaces through sliding and palpation can guide the selection of suction grasp locations and help to identify the local surface geometry. During suction gripping, this design localizes breaks in the suction seal between four quadrants with up to 97% accuracy and detects breaks in the suction seal early enough to avoid total grasp failure.

I. INTRODUCTION

Vacuum grippers are widely used to handle objects in industry. They perform astrictive grasping or, in other words, they apply attractive forces to object surfaces through suction pressure. The uncontact suction cup has the advantage of simple operation and enables the handling of a wide range of items, including those that are delicate, large or inaccessible by a jaw gripper [1].

One major challenge in suction grasping is how to plan a contact location among a variety of object geometries. Examples of planning methods include the heuristic search of a surface normal [2] and neural network training of grasp affordance using binary success labels [3]. Wan et al. use CAD model meshes to plan a grasp resisting gravitational wrench [4], and Dex-Net 3.0 learns the best suction contact pose from a point cloud considering both suction seal formation and gravitational wrench resistance [5]. These methods rely on RGB or depth sensors, which may not perceive fine details critical to suction sealing success, e.g., texture, rugosity, porosity, etc. Vision can also become occluded in cluttered environments. We investigate how sensing can be incorporated into the suction mechanism to monitor local contact geometry through haptic exploration.

Another challenge arises during forceful manual manipulation. In industry, robotic speed is desired for time efficiency,

This paper has supplementary downloadable material available at <http://ieeexplore.ieee.org>, provided by the authors. This includes movie clips of sliding haptic exploration, FTIR test setup, and image process of FTIR video data.

¹T. M. Huh, M. Li and H. S. Stuart are with the Embodied Dexterity Group, Dept. of Mechanical Engineering, University of California Berkeley, Berkeley, CA, USA. thuh@berkeley.edu

²K. Sanders, M. Danielczuk and K. Goldberg are with AUTOLAB, Dept. of Electrical Engineering and Computer Science, University of California Berkeley, Berkeley, CA, USA.

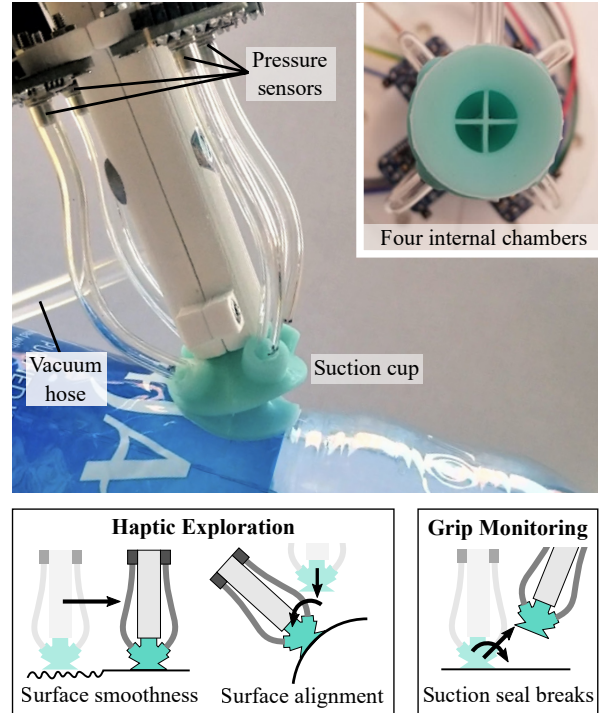


Fig. 1: The multi-chamber smart suction cup grips a water bottle. The cup has four internal chambers, each connected to a pressure transducer that provides a measure of internal flow rate. It is able to localize small breaks in the seal due to, for example, the rugosity (e.g., wrinkles, bumps, etc.) of the object surface or the application of external wrenches. Overall vacuum pressure is modulated in order to achieve different exploratory haptic procedures, such as sliding across surfaces.

however the inertial force induced by motion can cause grasping failure. Pham et al. use time-optimal path generation bounded by contact stability constraints to generate critically fast arm trajectories during pick-and-place [6]. Cheng et al. demonstrate an optimal control approach with a single suction gripper to reorient object by extrinsic dexterity, utilizing external contacts from the table [7]. Both methods utilize known inertial properties of the gripped object. These types of dynamic and forceful maneuvers could be adaptively achieved with the addition of suction cup tactile sensing, especially for objects with properties that might compromise suction seal.

Prior tactile sensors designed for use in suction cups provide partial information about object properties and vacuum sealing state. Aoyagi et al. coat piezoresistive polymer on a bellows suction cup to measure compression forces [8]. Doi et al. implement a capacitive proximity sensor on the base plate of a suction cup end-effector to measure the distance from the plate to the object surface [9]. These methods measure vacuum state indirectly from the deformation of the

suction cup and proximity to the object. Another straightforward approach is to monitor internal vacuum pressure of the suction cup as a discrete measure of suction sealing, as in [10]. None of these methods localize the source of a leak or measure local surface geometry. Nadeau et al. demonstrate how continuous suction-flow-rate measurement at the fingertips of a multifinger robotic hand can inform grasping and in-hand manipulation under water and note sensitivity to contact geometry [11], but do not address application to suction cups.

We hypothesize that the continuous measurement of flow rate within the suction cup itself provides new opportunities for smart suction cup dexterity. In this paper, we present a novel sensing method using inner chambers (Fig. 1), each of which connects to a pressure transducer to estimate distributed flow rates. In Section II, we describe the design and fabrication process of the suction cup, including Computational Fluid Dynamics (CFD) simulations of contact cases. Experimental results during haptic exploration in Section III suggest that this sensor can detect transitions between surface textures by gentle sliding motions, inspired by human haptic exploratory procedures [12]. This design also enables the evaluation of geometric surface normal contact through palpation. In Section IV, we monitor the suction signal during astrictive gripping as twists, resulting in external grasp wrenches, are applied. We find this sensor predicts the occurrence and location of catastrophic breaks in the suction seal around the perimeter of the cup. As discussed in Section V, the strength of this smart suction cup is in the handling unknown objects with variable surface properties that may compromise sealing.

II. DESIGN AND FABRICATION

This smart suction cup utilizes airflows inside its chambers to monitor local contacts. Internal wall structures separate the suction cup into four chambers (Fig. 1). Suction airflow is separated into each chamber and the pressure sensor connected to each chamber provides an estimate local flow rate. We implement the wall structure inside a single-bellows suction cup (Fig. 2) for its versatility on different curvatures and orientations of objects. The internal wall structure only spans the proximal portion of the suction cup, in order to maintain typical flexibility, deformation and seal formation in the distal lip.

We perform fabrication, including the chamber walls, with a single-step casting of silicone rubber. The casting mold is comprised of three parts, two outer shells and one core (Fig. 2), that are 3D printed using an SLA 3D printer (Formlabs, Form2). To ensure the clean casting of the thin wall structures (0.8mm thick), we used a syringe with a blunt needle (gauge 14) to inject uncured RTV silicone rubber (Smooth-On, MoldMax 40) and vacuum-degassed it. After curing, the outer shells are removed and the silicone suction cup is stretched and peeled off of the inner core mold. Tearing of the silicone can occur during this step, especially with harder rubbers.





Fig. 2: Casting mold and design of the suction cup. (a) Casting mold has three parts (2 Outer shells and 1 core). Molds are aligned and fixed by pins and bottom bolts. (b) The resulting suction cup is connected with vacuum connector and hoses to the pressure sensors. (c) Cross-sectional view of the suction cup shows internal and outer key dimensions.

A. System integration

The suction cup is connected with pressure transducers and a vacuum source with pressure regulation. For experimental trials, this system is then integrated with a Universal Robot arm. We used ROS to control the robot arm and collect both pressure sensor and wrist force/torque (F/T) sensor data.

Four ported pressure sensors (Adafruit, MPRLS Breakout, 24bit ADC, 0.01Pa/count with an RMS noise of 5.0 Pa) connect with the four chambers of the smart suction cup via polyurethane tubes. The suction cup and the pressure sensors attach to a 3D printed fixture (Fig. 1). This fixture is then attached to the wrist F/T sensor (ATI, Axia80, sampling rate 150Hz) on the robot arm (Universal Robots, UR-10). A microcontroller (Cypress, PSoC 4000s) is fixed to the arm proximal to the load cell and communicates with the four pressure sensors via I2C at a 166.7Hz sampling rate.

A vacuum generator (VacMotion, VM5-NA) converts compressed building air to a vacuum source with a maximum vacuum of 85kPa. A solenoid valve (SMC pneumatics, VQ110, On/off time = 3.5 / 2ms), commanded by a microcontroller, regulates the compressed air as a means of moderating vacuum intensity. During haptic exploration experiments, the valve is controlled with pulse width modulation (PWM) at a frequency of 30Hz with 30% duty cycle to approximate lower vacuum pressures. We chose this PWM setting considering the on-off time of the solenoid valve and the sampling rate of the pressure sensor. The vacuum hose that applies suction to the cup is attached at both the suction cup vacuum connector and proximal to the load cell to reduce tube movement and subsequent F/T coupling.



(a) Vertical leakage airflow (b) Horizontal leakage airflow
Fig. 3: Two cases of CFD simulation. Dark blocks are engaged objects.

B. CFD Simulation

Using Computational Fluid Dynamics (CFD) simulation (COMSOL Multiphysics, $k - \epsilon$ turbulence model), we evaluate the gripper in two suction flow cases: vertical and horizontal flow (Fig. 3). The vertical flow case emulates when the suction cup only partially contacts a surface, or when the surface's shape inhibits sealing. However, when the suction cup engages with a smooth flat surface, flow can only move inward from the outer edges of the cup, as in the horizontal flow case. This horizontal leak is common as the suction cup is wrenched from the surface after a suction seal is formed. Although the suction cup will deform under vacuum pressure, we use modeled rigid geometry in the CFD simulation. For each case, we approximate the leak flow direction with a small pipe ($D = 1\text{mm}$, $L = 7\text{mm}$) located close to one of the internal chambers as shown in Fig. 4 and Fig. 5. The boundary conditions of the vacuum pump pressures and flow rates match the experimental vacuum generator we used.

The simulation results suggest that the gripper can locate leakage flow using differences between the four pressure transducers. We defined vacuum pressure (P_{vac}) as


$$P_{vac} = P_{atm} - P_{chamber}. \quad (1)$$

In the vertical leakage flow case, P_{vac} close to the leaking orifice shows the least vacuum pressure than the others (Fig. 4a). On the other hand, the horizontal leakage causes the diagonally opposite channel to have the lowest P_{vac} (Fig. 4b). These trends are supported by the flow results in Fig. 5, where the vertical and horizontal orifices produce the highest flow rate is opposite chambers. The simulation result also shows an estimate of the pressure difference between chambers ($\sim 0.4\text{kPa}$) which can be differentiated by the selected pressure sensors.


III. HAPTIC EXPLORATION

In suction cup grasping, contact surface smoothness enables the formation of a suction seal. Surface textures and rugosity alter the quality of the seal that determines the resistible grasp wrench. Local geometry of the contact, which is influenced by surface curvature and suction cup orientation, can also compromise seal formation. However, this information is difficult to obtain from vision alone because of limited resolution and potential for occlusion.

We propose a method of haptic exploration using lower suction forces by regulating vacuum pressure. Using full



(a) Vertical Orifice (b) Horizontal Orifice
Fig. 4: CFD result of the P_{vac} measured at the sensor locations of each chamber. Bar graphs are scaled to represent the 0.4kPa from the maximum of the four P_{vac} s



(a) Vertical Orifice (b) Horizontal Orifice
Fig. 5: Cross sectional view of the pressure distribution. Arrows inside represents relative logarithmic scale of air flow velocity. The colormap unit is kPa.

vacuum power can apply large suction forces, even to non-optimal contacts, causing the object to move and making it difficult to move the cup to a new location. Using lowered vacuum pressure, the suction cup can gently slide along the surface, allowing efficient continuous exploration.

A. Experiments

First, we estimate surface texture using the smart suction cup. On a substrate with sandpapers of different roughness (120 to 600 grit), we apply both full vacuum and lowered vacuum to the suction gripper and measure the pressure sensor response. For both full vacuum and lower vacuum cases, the robot arm approaches the contact surface without a vacuum until the normal force reading from wrist F/T sensor reaches 0.5N , and then the vacuum is turned on.

Second, we test sliding haptic explorations on acrylic plates with textured-smooth surface transitions (Fig. 8). Using the robot arm, we apply an initial normal force of 0.5N and then apply low suction. We exclude the full vacuum setting because it inhibits sliding. The suction cup starts with only half contact on the textured surface. The robot then moves the suction cup laterally at 6.5 mm/sec speed while measuring the vacuum pressure in each chamber. We test on two sets of surface transitions: wavy-smooth and ribbed-smooth. We laser-cut acrylic plates in $35\text{mm} \times 51\text{mm}$ and bonded textured-smooth acrylic pairs along the longer edge using acrylic cement.

Third, we use our tactile sensing suction gripper to seek surface normal and evaluate surface curvature through palpation. We use two 3D printed test objects with different contact radii at the tip (8mm and 11mm), shown in Fig. 9. The robot changes the orientation of the end-effector from




Fig. 6: Texture evaluation result with full vacuum (top) and lowered vacuum (bottom) mode. For lowered vacuum, the valve was controlled with 30Hz, 30% duty cycle PWM. Shades represent ± 1 standard deviation. (Right) P_{vac} on 400 grit sand paper with lowered vacuum.




Fig. 7: Magnitude of discrete Fourier transform of the lower vacuum mode at 30 Hz

normal by pivoting about the tip of the object with a range of $\pm 30^\circ$ in 15° increments. At each orientation, the robot then translates the suction cup until the gripper normal force reaches a set threshold, either 0.5N or 1N. Low vacuum is applied and the pressure transducer readings are recorded. Then, we evaluate suction grasp at the surface normal with full vacuum setting.

B. Results

1) **Surface Texture evaluation:** While both full vacuum mode and lowered vacuum mode differentiate surface textures, the lowered vacuum mode interacts with the surface more gently. When engaged with the full vacuum mode, the measured P_{vac} in the suction cup increases with smoothness (Fig. 6(top)), reaching 74 kPa for the smoothest 600 grit sandpaper. In contrast, the lowered vacuum mode produces 25 times lower P_{vac} on average while still showing a similar trend between mean pressure and surface smoothness (Fig. 6(bottom)).

The high variance in P_{vac} with the lower pressure mode is induced by the frequent on-off cycles in PWM control, resulting in poor signal-to-noise ratio. However, we find that the amplitude of this oscillating P_{vac} signal provides a clean measure of surface texture. On smoother surfaces, the formation of a good suction seal allows higher variation in P_{vac} between on and off phase, causing a higher amplitude of oscillation at PWM frequency (Fig. 6 (right)). To effectively capture the amplitude of the signal, we use the magnitude of Discrete Fourier Transform (DFT) at the PWM driving frequency (30Hz)(Fig. 7), which shows similar trends as raw P_{vac} measures in Fig. 6(top). In the following subsections, we use DFT magnitude as a metric of contact evaluation.

2) **Sliding exploration:** To monitor gradual changes of the DFT during sliding motions, we use Short Time Fourier Transform (STFT) that uses Hamming windowed data of 0.5 seconds (83 samples). Each time stamp in Fig. 8 is evaluated as a center of this 0.5-second window. We define the magnitude of STFT at PWM driving frequency (30Hz) as $|STFT_{30}|$. Using this $|STFT_{30}|$ of the pressure sensors, the transition of the contact surface properties can be detected by sliding motions. In this test, contact transitions are compared with both the average $|STFT_{30}|$ over all chambers ($\overline{Ch_{all}}$) and the directional difference between the leading-edge chamber ($\overline{Ch_{right}}$) and following-edge chamber ($\overline{Ch_{left}}$).

The first transition is the initial half contact on the textured surfaces (at 1.1 seconds in Fig. 8); while the event is unnoticeable in the $\overline{Ch_{all}}$ signal, the $|STFT_{30}|$ of the contacting side $\overline{Ch_{right}}$ is greater than $\overline{Ch_{left}}$. This outcome matches expectations from the vertical orifice simulation in Fig. 4a. The following contact transitions, achieving full contact on the textured surface and then full contact on the smooth surface, can be detected by the increasing $\overline{Ch_{all}}$ over time. The ribbed surface provides less suction engagement than the wavy surface, but both are detectable. During large transitions, defined here as when $\overline{Ch_{all}}$ changes magnitude by over half of the full range signal within a short time-window (i.e., wavy half-to-full contact and ribbed-to-smooth contact), the difference in $|STFT_{30}|$ between $\overline{Ch_{right}}$ and $\overline{Ch_{left}}$ shows characteristic positive-negative shifts; at the beginning of these transitions, the leakage air on the left side flows more vertically (Fig. 4a) while its direction transitions to horizontal as more of the suction cup seals with the surfaces (Fig. 4b).

3) **Exploring surface normal of curved objects:** The surface normals of curved objects are found by the symmetric patterns of $|DFT_{30}|$ over the tested palpation angles. Fig. 10 shows how this signal varies on the different test objects and with different initial palpation forces. For all test cases, the $|DFT_{30}|$ of the left channels (Ch1, Ch2) are approximately symmetric to the right channels (Ch3, Ch4) about the surface normal ($\theta = 0^\circ$). This trend follows the simulation result in vertical orifice case (Fig. 4a), suggesting that the side with less contact produces lower $|DFT_{30}|$. When $\theta = 0^\circ$, $|DFT_{30}|$ in all channel pressures converge to a non-zero value, meaning that leaking is evenly distributed. $|DFT_{30}|$ of the right side channels are consistently lower than the left channels likely due to both slight experimental misalignment and smart cup manufacturing imperfections.

We posit that the quality of suction contact when $\theta = 0^\circ$ can be evaluated by the convergent magnitude of $|DFT_{30}|$. On the object with $R=8\text{mm}$ (Fig. 10a), $|DFT_{30}|$ at the surface normal ($\theta = 0^\circ$) is lower than maximum $|DFT_{30}|$ across all angles testes, meaning that the high curvature of the object prevents complete sealing at any orientation. In contrast, on the lower curvature $R=11\text{mm}$ object (Figs. 10b and 10c) $|DFT_{30}|$ at the surface normal is maximum over all angles tested. When the initial palpation force is not sufficient to form a full vacuum seal (Fig. 10c), the $|DFT_{30}|(\theta = 0^\circ)$ is significantly lower than in the suction grasp success case

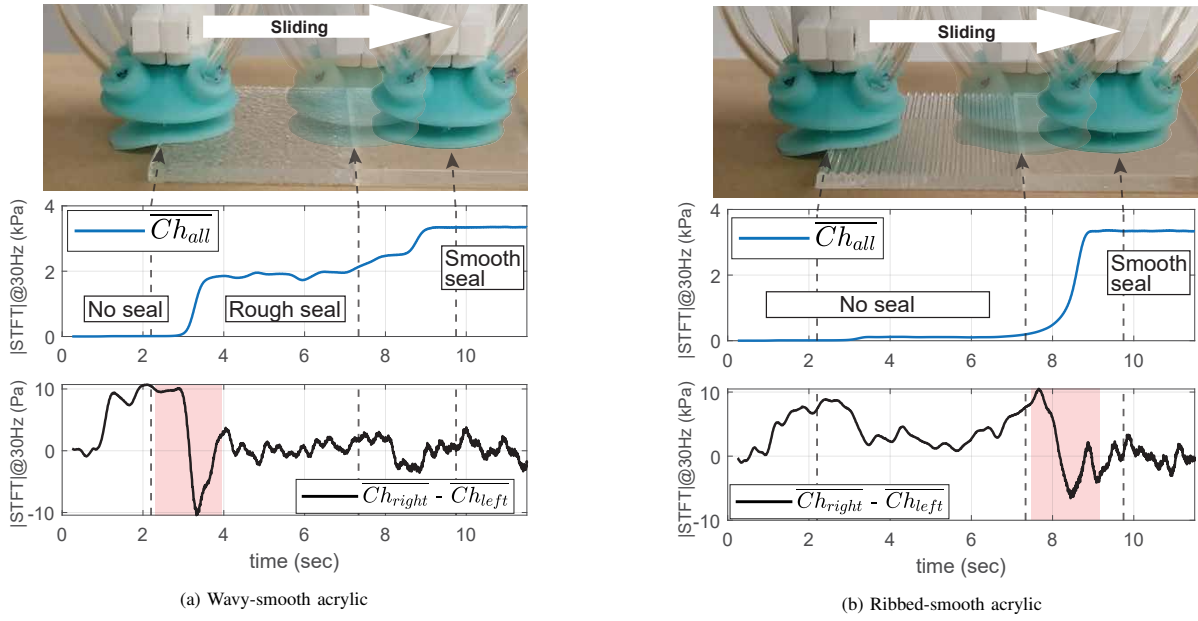


Fig. 8: Sliding test results on (a) wavy-smooth acrylic and (b) ribbed-smooth acrylic pairs. (Top) Snapshot of the sliding motion at the time marked with dashed line. (middle) Average $|STFT_{30}|$ of all four channels (bottom) Difference of $|STFT_{30}|$ between right two channels and left two channels. Red-shades are when large transition (no seal to seal) occurs.




Fig. 9: Experiment setup for seeking a surface normal on two different objects

(Fig. 10b).

IV. MONITORING DETACHMENT

Monitoring leaks in the seal of a suction cup can help in preventing grasp failures during robotic manipulation, especially of unknown objects that may have variable surface properties. We hypothesize that it may be possible to detect and localize breaks in the suction seal early enough to avoid total astrictive grasp failure. As a first step toward an adaptive controller capable of such behaviors, we examine the capability of the gripper to monitor spatial vacuum seal states and localize the leakage airflow in both current and future states in a controlled experiment.

A. Dataset

We estimate the contact seal states around the perimeter of the suction cup while collecting time-series tactile sensing and wrist F/T sensor data as the suction cup is detached from a smooth flat plate. The ground-truth contact seal states is measured using a Frustrated Total Internal Reflection (FTIR) setup, as in Li et al. [13]. A 13mm thick smooth acrylic plate is illuminated from the side with red color filtered LED strips. When the suction cup engages with the acrylic, the contact regions appear bright red to a high speed digital camera (240 FPS) that records images from the opposite side.

The suction cup is forcefully detached from the acrylic substrate under maximum vacuum for 740 separate trials

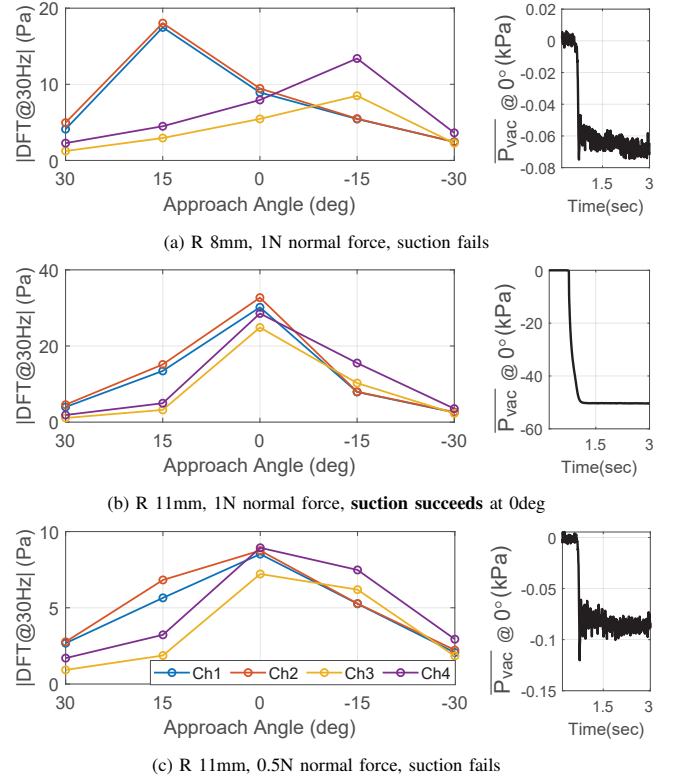


Fig. 10: Test results of seeking surface normal. (Left) Each data-point represents the mean of three trials. Three different curvature-initial normal load pairs all show symmetric trends of $|DFT_{30}|$ about the surface normal (0 deg). The level of $|DFT_{30}|$ at surface normal corresponds to suction grasp success/failure. (Right) mean $|P_{vac}|$ of 4 channels with full vacuum setting at the surface normal (0 deg).




Fig. 11: Frustrated Total Internal Reflection (FTIR) test setup and robot motion with an example twist. Dashed arrow is a linear velocity and the purple arrow is an axis of angular velocity.




Fig. 12: Vacuum pressure measured at each chamber during detaching motion. Insets are gray-scaled image frames from the FTIR sensor.




Fig. 13: Photo of center point tracking in the image sequence. The characteristic points (blue 'X' in images) around the center of suction cup are tracked throughout the video sequence. The center point (black circle) is the mean coordinate of all tracked points.

(Fig. 11). The UR-10 robot arm applies 450g of initial vacuum suction loading away from the substrate. Then the robot applies a twist to the end-effector until it detaches. The angular speeds are set to $30^\circ \pm 5^\circ/\text{sec}$ and the axis of rotation is set by φ in $[0^\circ, 180^\circ]$ with 22.5° increments and θ in $[0^\circ, 330^\circ]$ with 30° increments. Each pair of φ and θ is tested at least six times, and a uniformly distributed perturbation in $[-0.5, 0.5] \times \text{increment}$ is added to φ and θ . The translation velocity is chosen randomly between $[-1, 1]$ cm/sec for (x,y) axis and $[0.6, 2]$ cm/sec for (z) axis. At the end of each detachment, the high-speed video is synchronized with other sensor data by blinking the FTIR LED and aligning timestamps.

An example sequence is shown in Fig. 12. As the suction cup deforms under the detaching twist, the vacuum seal (bright ring) also deforms and then leaks. The P_{vac} measure in each chamber does not vary while the vacuum seal remains, but lowers at the onset of the leak. The difference in P_{vac} among chambers shows expected trends based on the




Fig. 14: Process of obtaining contact state label from a single frame. (Left) the quadrants are defined around the tracked center point. For each quadrant, the averages of normalized intensity in a small radial band ($dr = 16$ pix) of subsectors (9 per quadrant) are evaluated along the radial direction (r). The minimum peak among nine subsectors is considered the contact state. (Right) the contact label is computed for all quadrants, resulting in a 4×1 vector.

Threshold (th)	Q1	Q2	Q3	Q4
0.5	0.50	0.65	0.6	0.40
0.6	0.34	0.46	0.59	0.44
0.7	0.19	0.14	0.87	0.08

TABLE I: Break occurrence rate of each quadrant where the contact label first decreases below thresholds (0.5-0.7) in 740 trials. Multiple quadrants can break the threshold at the same time, making the rate sum greater than 1.0.

horizontal leakage CFD simulation in Fig. 4b. The example in Fig. 12 shows the least P_{vac} in the chamber opposite to the leak point; the leak occurs near 'Ch2' and the pressure drops most rapidly for 'Ch4'. The large deformations of this suction cup are a consequence of using soft rubber to ease the fabrication process, described in Section II. To estimate local vacuum seal states in this highly deformable suction cup, we trained an artificial neural network using P_{vac} signals and wrist F/T sensor data as inputs.

For training labels, we process the video data to obtain time-series contact states. We segment the suction cup into 4 quadrants around the center, and we define the contact state label of a single video frame as a vector of 4 elements that represents the contact of each sector. Because the center point moves, as shown in Fig. 12, we track the center using the Kanade-Lucas-Tomasi (KLT) algorithm [14] featured in MATLAB. As shown in Fig. 13, the characteristic points (~ 15 points) around the center are selected from the first contact image by using minimum eigenvalue corner detection algorithm [14]. Then the KLT algorithm tracks the points in the video frame sequence. The mean of the coordinates of tracked points are assumed to be the center point. If the algorithm loses track of more than half of the points, the center point is extrapolated from previous center trajectory.

Then the contact label is computed around the tracked center point as shown in Fig. 14. The four quadrants are redefined every frame from the tracked center point and the interpolated end-effector orientation that is sampled at 10Hz. We assume that the contact state of each sector is dictated by the weakest, or darkest, contact points along the suction seal. For each sector, we use the normalized image intensity values as a metric of contact states, and compute the average intensity in a small subsector band ($dr = 16$ pixels) along

Model	MSE ↓	Break Quadrant Accuracy ↑			Median Break Time Error (ms)		
		$th = 0.5$	$th = 0.6$	$th = 0.7$	$th = 0.5$	$th = 0.6$	$th = 0.7$
FT Only ($h=30$ ms)	0.003	86%	82%	88%	6 ± 18	0 ± 24	0 ± 44
Vac Only ($h=30$ ms)	0.006	80%	97%	82%	18 ± 18	24 ± 12	-24 ± 138
FT+Vac ($h=30$ ms)	0.006	95%	82%	82%	24 ± 6	24 ± 12	-12 ± 144

TABLE II: **Learning to predict suction failure.** Mean squared error (MSE) of predicted contact state, break quadrant accuracy (BQA), and median break time error (MBTE) (in timesteps with interquartile range shown) for models trained on vacuum pressure, force/torque data, or both. BQA and MBTE are reported for different failure cutoff thresholds $th = 0.5, 0.6$, and 0.7 . Positive MBTE values indicate the model is overly conservative, while negative MBTE values indicate the model does not predict the breakage until it is imminent; in practice, an overly conservative model such as FT+Vac may be more beneficial to a system that must quickly react to an imminent breakage. All models learn to localize the first leakage point in both space and time for lower thresholds, but accuracy decreases as the threshold increases.




Fig. 15: Architecture of the deep learning model used to learn contact state. Takes force, torque, and pressure features at timestep t and the previous step's hidden cell as input, and outputs a new hidden state along with features that are passed through a fully connected layer to produce the contact state scores for each section of the suction cup.

the radial direction (Fig. 14). The minimum peak of mean intensity values among the nine subsectors is assumed as that quadrant's contact label. Table I shows that catastrophic vacuum seal break, defined here as dropping below a 50% or 60% contact state threshold ($th = 0.5$ or 0.6), occurs evenly throughout the four quadrants. Moderate seal weakening ($th = 0.7$) occurs more at a particular quadrant, which may be caused by mechanical bias from the fabrication process.

The time series of the contact state labels in all 4 quadrants (size: $t \times 4$) are the output data, and the time series of 4-channel P_{vac} and 6DOF F/T sensor data (size: $t \times (4+6)$) are the input data of the deep learning training. All timestamps of F/T sensor data and contact state labels are interpolated to match the P_{vac} timestamps.

B. LSTM Training

To predict both the current seal state as well as future seal states, we trained a long short term memory (LSTM) network [15] on the given input data. The architecture is represented in Fig. 15. Specifically, we wish to predict the contact state of a suction cup at time $t + h$ given the F/T and pressure sensor data from time $[0, t]$, where h is the prediction horizon. We aim to predict h ms into the future based on current sensor data as this would allow for preventative measures to be executed before suction failure occurs on the actual machine if this model were employed in a system. We use two LSTM cells with 10 input features and 200 hidden layer units, followed by a fully-connected layer with 200 input features and 4 output features. The

memory unit of the LSTM allows the model to learn key features of time series data, and 200 hidden units allows for the model to develop a robust representation of the input data sequence while still being fast to train. At prediction step t' , sensor data for time $t = t'$ is passed into the LSTM and contact state scores are produced for time $t = t' + h$. The four contact state outputs correspond to one quadrant of the suction cup. We trained and validated on 592 detachment sequences, leaving 148 sequences as a test set. We train the network for 100 epochs using the Adam optimizer and a learning rate of 0.001. The training takes about 5 minutes on an NVIDIA Titan X GPU.

C. Results

To gauge the ability of the LSTM to predict both the seal state over time as well as its ability to predict impending seal breakage, we introduce three evaluation metrics: 1) mean squared error (MSE), which measures the ability of the LSTM to track the contact seal state across the entire sequence, 2) break quadrant accuracy (BQA), which measures the ability of the LSTM to localize the quadrant of the first seal break across each sequence, and 3) median break time error (MBTE), which measures the ability of the LSTM to localize the seal breakage in time.

Mean squared error is calculated between the normalized predicted and ground-truth contact values, averaged over the length of the sequence. Break quadrant accuracy is calculated by comparing the first quadrant to dip below a contact state threshold th in each of the ground-truth and predicted series. Median break time error (with the interquartile range also indicated in the table) is calculated by averaging the error in time steps between the first time the first quadrant drops below th in the predicted and ground-truth series, respectively. The latter two metrics are particularly useful for an adaptive controller, which must robustly identify catastrophic seal failure before it occurs.

Table II shows the results of training the network using just F/T sensor data (FT Only), just vacuum sensor data (Vac Only), and using both inputs (FT+Vac). In all cases, the network learns to predict the contact state with low MSE, and the FT+Vac model can additionally learn to robustly predict the first quadrant that will leak below a threshold of 0.5. However, as the failure cutoff threshold increases, the predictions are less accurate in both space and time.




Fig. 16: Contact state prediction results for one seal quadrant of a previously unseen trial. Model prediction delay is set to $h = 5$, so the contact state prediction at time t on the graph is produced using sensor data through time $t - h$. Figure labels $y = 0.5, 0.6$, and 0.7 thresholds used to calculate Median Break Time Error, which is also labeled at $th = 0.7$. Raw error used to calculate the MSE is also labeled at $t = 420ms$.

This result suggests that the network accurately identifies impending catastrophic failures correctly, but may not be able to completely model stretching or bending effects that can cause smaller amounts of leakage. The median break time error results suggest that the force/torque data alone can accurately identify the breakage time, but adding the vacuum pressure inputs causes the model to predict a breakage more conservatively. Since positive MBTE values indicate an overly conservative prediction (i.e., the model predicts a breakage will happen in 30 ms when it will actually occur in 54 ms), this conservative nature may actually benefit the system in practice, as it can react with more time before the breakage. In contrast, the F/T Only model can sometimes be too optimistic (negative MBTE).

Ablation experiments on the prediction horizon parameter, h , suggest a linear relationship between the MSE and h , as the MSE increases by approximately 0.002 for every 60 ms h is increased by over a range of $30 \leq h \leq 330$. Data indicates that higher horizon values correspond to large, negative BTE scores, such that predictions become inaccurate and late, on average, as the horizon grows.

V. DISCUSSION

A single-bellows suction cup gripper can examine object surfaces by haptic exploration and monitor detachment details of the vacuum seal. The four-chamber cup design generates differential airflow between the chambers, measured with remote pressure transducers such that sensitive electronics are not exposed to physical damage or fatigue. We find that lowered vacuum pressure facilitates gentle interactions with objects. During haptic exploration, this smart suction cup responds to surface textures, transitions between different surfaces, surface normal and local curvature of touched surfaces. The next step is to further characterize how these gentle and exploratory haptic signals correlate with strong attractive grasp wrench limits, so that the gripper can apply more versatile planning and adaptive control on a wide variety of objects.

This tactile sensor detects spatial and temporal details of suction cup detachments when applied to an LSTM network. In these trials, the suction cup sensor (<\$80 for the raw components) predicts the quadrant and time of

initial breaks with comparable accuracy as using a wrist force/torque sensor (ATI Axia80, ~\$3,000). However, it is not our intention to create an equivalent to the common wrist load-cell. While our results highlight the smart suction cup's sensitivity to external loads through deformation, we only test on a smooth, flat, fixed acrylic plate. Future work will characterize the effect of unknown, variable surface properties on the prediction of leaks, which we expect to accentuate the importance of internal suction cup flow sensing over the measurement of external grasp forces alone.

VI. ACKNOWLEDGMENTS

T.M.Huh was supported by the University of California at Berkeley Embodied Dexterity Group. M. Danielczuk was supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. 1752814. The work of M. Li was supported by the NASA Space Technology Research Fellowship, under Grant #80NSSC19K1166.

REFERENCES

- [1] N. Correll, K. E. Bekris, D. Berenson, O. Brock, A. Causo, K. Hauser, K. Okada, A. Rodriguez, J. M. Romano, and P. R. Wurman, "Analysis and observations from the first amazon picking challenge," *IEEE Trans. Automation Science and Engineering*, vol. 15, no. 1, pp. 172–188, 2016.
- [2] D. Morrison, A. W. Tow, M. McTaggart, R. Smith, N. Kelly-Boxall, S. Wade-McCue, J. Erskine, R. Grinover, A. Gurman, T. Hunn *et al.*, "Cartman: The low-cost cartesian manipulator that won the amazon robotics challenge," in *ICRA*. IEEE, 2018, pp. 7757–7764.
- [3] A. Zeng, S. Song, K.-T. Yu, E. Donlon, F. R. Hogan, M. Bauza, D. Ma, O. Taylor, M. Liu, E. Romo *et al.*, "Robotic pick-and-place of novel objects in clutter with multi-affordance grasping and cross-domain image matching," in *ICRA*. IEEE, 2018, pp. 3750–3757.
- [4] W. Wan, K. Harada, and F. Kanehiro, "Planning grasps with suction cups and parallel grippers using superimposed segmentation of object meshes," *IEEE Transactions on Robotics*, 2020.
- [5] J. Mahler, M. Matl, X. Liu, A. Li, D. Gealy, and K. Goldberg, "Dex-net 3.0: Computing robust vacuum suction grasp targets in point clouds using a new analytic model and deep learning," in *ICRA*. IEEE, 2018, pp. 5620–5627.
- [6] H. Pham and Q. C. Pham, "Critically fast pick-and-place with suction cups," *ICRA*, vol. 2019-May, pp. 3045–3051, 2019.
- [7] X. Cheng and M. T. Mason, "Manipulation with suction cups using external contacts," in *ISRR*, 2019.
- [8] 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.
- [9] S. Doi, H. Koga, T. Seki, and Y. Okuno, "Novel proximity sensor for realizing tactile sense in suction cups," in *ICRA*. IEEE, 2020, pp. 638–643.
- [10] C. Eppner, S. Höfer, R. Jonschkowski, R. Martín-Martín, A. Sieverling, V. Wall, and O. Brock, "Lessons from the amazon picking challenge: Four aspects of building robotic systems," in *Robotics: science and systems*, 2016.
- [11] P. Nadeau, M. Abbott, D. Melville, and H. S. Stuart, "Tactile sensing based on fingertip suction flow for submerged dexterous manipulation," in *ICRA*. IEEE, 2020, pp. 3701–3707.
- [12] S. J. Lederman and R. L. Klatzky, "Extracting object properties through haptic exploration," *Acta psychologica*, vol. 84, no. 1, pp. 29–40, 1993.
- [13] M. S. Li, D. Melville, E. Chung, and H. S. Stuart, "Milliscale features increase friction of soft skin in lubricated contact," *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4781–4787, 2020.
- [14] J. Shi *et al.*, "Good features to track," in *Conference on computer vision and pattern recognition*. IEEE, 1994, pp. 593–600.
- [15] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.