

Constructing Contact Estimation Models for Barometric Tactile Sensors

Sharmi Shah¹, Ethan Chun¹, Hongmin Kim¹, Andrew SaLoutos¹, David Nguyen¹,
 TaeWon Seo², and Sangbae Kim¹

Abstract—Barometric tactile sensors present a cheap and customizable method for adding tactile sensing to robotic platforms. These sensors consist of commercially available MEMS barometers embedded in an elastomer. However, as the sensing surface and elastomer volume increase in complexity, time-dependent material dynamics reduce sensing accuracy. We present a collection of inference and usage recommendations towards mitigating these dynamics and improving sensor force and localization resolution. Using two custom, curved, barometric tactile sensors as case studies, we demonstrate that a new data collection regime alone can improve normal force predictions by 30.4% compared to prior work. We further introduce a Binned-RNN inference architecture and demonstrate its efficacy through select ablations. Small enough to run on the sensor’s integrated microcontroller at 100Hz, we find our model achieves a minimum spatial resolution of 0.86 mm on an ellipsoid tactile sensor. Finally, we demonstrate the robustness of these sensing capabilities through freeform contact and controlled object rolling. Video demonstrations can be found at <https://youtu.be/mi7qqjssirg>.

I. INTRODUCTION

Tactile sensing plays a critical role in enhancing robotic manipulation by providing information that vision and proprioception alone cannot capture. A wide range of tactile sensing technologies have been developed to achieve this goal, including capacitive, piezoresistive, optical, and magnetic sensors [1]. Each modality has its trade-offs. For example, optical tactile sensors can offer rich spatial information and high sensitivity, but often have low bandwidth and require larger form factors due to internal camera components [2, 3]. Capacitive and magnetic sensors, while compact, can be susceptible to electrical noise and magnetic interference respectively [4, 5]. Many such sensors also require specialized fabrication equipment or cleanroom processes, which increase both cost and development time.

In contrast, barometric tactile sensors are constructed by embedding off-the-shelf MEMS barometers in commercially available elastomers. Fabrication is low-cost and straightforward: a 3D printed mold is used to cast an elastomer over a set of barometers. Once cast, the elastomer layer not only transmits contact forces but also shields the barometers.

Despite the advantages of barometric tactile sensors, the challenge of accurately interpreting their signals remains

¹ Authors are with the Biomimetic Robotics Laboratory at the Department of Mechanical Engineering, Massachusetts Institute of Technology (MIT), Cambridge, MA, 02139, USA. sharmis@mit.edu

² Author is with the Robot Design Engineering Lab at the Department of Mechanical Engineering, Hanyang University, Seoul, 04763, Korea.

This work was supported by the Advanced Robotics Lab of LG Electronics Co., Ltd. and the MOTIE in Korea, under the Global Human Resource Development for Innovative Design in Robot and Engineering (P0017306) supervised by KIAT.




Fig. 1. **Barometric tactile sensor designs.** Both tactile sensors used in this work feature a microcontroller and eight barometers in a non-planar arrangement embedded within a curved cast elastomer. The cast sphere on the spherical sensor (top row) has a radius of 10 mm. The cast ellipsoid on the ellipsoid sensor (bottom row) has semi-axis measurements of $a = 10.7$ mm, $b = 9$ mm, $c = 6.8$ mm. Refer to [6] for further design details.

significant. While MEMS barometers themselves are well-characterized and consistent, the surrounding elastomer introduces a complex, nonlinear mapping between surface contact forces and internal pressure readings. This mapping is strongly influenced by the sensor’s geometry, the properties of the elastomer, and the nature of the contact. To recover meaningful information from these signals, prior work has typically relied on data-driven methods to approximate the inverse mapping [7–9]. These models are usually compact enough to run on embedded systems for real-time inference of contact forces and locations. However, elastomers also exhibit time-dependent behaviors such as hysteresis and stress relaxation, which degrade the reliability of learned models, particularly in sensor designs with non-trivial shapes.

In our custom-designed tactile sensors, the time-dependent viscoelastic effects can lead to significant variability in prediction accuracy. To address these challenges and ensure that data-driven methods keep pace with the expanding variety of barometric tactile sensor designs, this paper presents and validates a set of practices for building lightweight and robust contact estimation models for barometric tactile sensors.

Drawing from empirical results across two sensor designs, we highlight three practices that improve inference performance in our experiments:

- 1) Building off prior work in machine learning, we find that the combined effect of collecting data at a fixed frequency, uniformly sampling across the sensor sur-

face, and including both in-contact and out-of-contact samples reduces prediction errors on our tactile sensors.

- 2) Furthermore, we find that discretizing continuous contact location parameters into bins and training with cross entropy loss reduces bias toward the center of the distribution, making prediction accuracy more uniform across the prediction range.
- 3) Finally, we find that mirroring the time-dependent behavior of the elastomer by incorporating temporal context in the network architecture reduces variation in prediction accuracy across time.

In this work, we demonstrate the success of these practices with experiments on a previously introduced sphere-shaped sensor [6] and a new ellipsoid-shaped sensor. First, we show improved force prediction accuracy relative to the previous work on the spherical sensor resulting from improvements to our data collection methods. Next, we perform ablations on our inference pipeline, demonstrating the efficacy of each of our prescribed changes. Subsequently, we test sensor robustness to material aging with a separate validation trajectory and robustness to out of distribution events with a freeform contact sequence. Finally, we demonstrate the integrated system through a controlled rolling task that requires accurate tactile feedback for successful execution.

II. RELATED WORK

We focus on geometry, output representation, and inference methods used by prior barometric tactile sensors. For a more comprehensive overview of tactile sensing, please refer to [1, 10].

A. Sensor Geometry

Barometric tactile sensors are appealing due to their ease of fabrication, which in principle supports a wide range of sensor geometries. However, existing barometric tactile sensors generally fall into two simple shapes: planar arrays and domes. Planar arrays can either be constructed by combining TakkTile barometer units [11] into an array [7], [12–14] or by fabricating a barometer array directly [15–18]. Importantly, the sensing surface is offset by a constant amount from all barometers. The other common sensor shape, tactile domes, typically consist of a radially symmetric array of barometers placed beneath a cast dome [8, 9, 19, 20]. The symmetries in both arrangements greatly simplify modeling.

Despite the manufacturing flexibility of barometric tactile sensors, there has been limited exploration of more complex geometries. This work focuses on developing inference methods that also generalize to sensors with non-planar and non-domed geometries. As case studies, we evaluate a previously introduced spherical sensor [6] and a newly developed ellipsoid sensor.

B. Sensor Output Representation

A perfect tactile sensor would be able to reconstruct and localize the full set of forces applied to the sensor at a given time. However, given the complexity of the aforementioned

elastomer mapping function, most works focus on returning only a subset of this data.

Previous works have focused on aggregate force prediction without contact localization [9], single point contact localization without force prediction [8], and planar contact localization and force prediction [21], among other parameterizations. Notably, Epstein et al. [19] simultaneously estimated single point contact and a three-axis force on a domed barometric fingertip tactile sensor.

We adopt this parameterization, as it captures task-relevant features for many robotic manipulation scenarios. However, the previous model using this approach often estimated a contact force even in the absence of contact due to viscoelastic effects. To address this, we extend the output space with a contact flag that explicitly predicts whether the sensor is in contact in order to help disambiguate true contact from barometer signals caused by elastomer relaxation.

C. Sensor Modeling

A variety of approaches have been used to extract contact information from barometric sensor data. Analytical methods such as finite element models [8] and mechanical approximations [7] offer physical insight, but become intractable for irregular geometries. Gaussian processes [19] and neural networks [6, 12, 22] have shown greater flexibility.

Most neural network models in prior work treat the problem as static regression. For example, Chuah and Kim [22] and Reeks et al. [12] use MLPs to predict contact forces on near-planar arrays; SaLoutos et al. [6] use an MLP to jointly regress contact location and force on a three-quarter sphere; and Hou, Zhou, and Spiers [17] combine an MLP and CNN-LSTM for object pose estimation on planar sensors.

While prior work has incorporated temporal context for downstream tasks, we instead use temporally-aware architectures to directly regress contact location and applied forces so downstream models can be designed independent of the sensor dynamics. Second, we reformulate the inference problem as classification over discretized surface location bins, improving contact localization resolution and making accuracy more uniform across the sensor.

III. SENSOR DESIGNS

Barometric tactile sensors are highly customizable, but their performance depends strongly on geometry and material volume. This section compares the two sensor designs used in this study, shown in Fig. 1, which were manufactured via the process described in [6]. The sensors share the same PCB design and elastomer material, but differ in elastomer shape.

A. Spherical Sensor Design

The first design is a previously developed spherical sensor [6] which used neural networks to estimate contact location and force from barometer readings. However, its large elastomer volume introduced pronounced viscoelastic effects, such as stress relaxation and residual stress accumulation, which degraded model accuracy. In addition, the non-planar arrangement of barometers meant that barometers were often

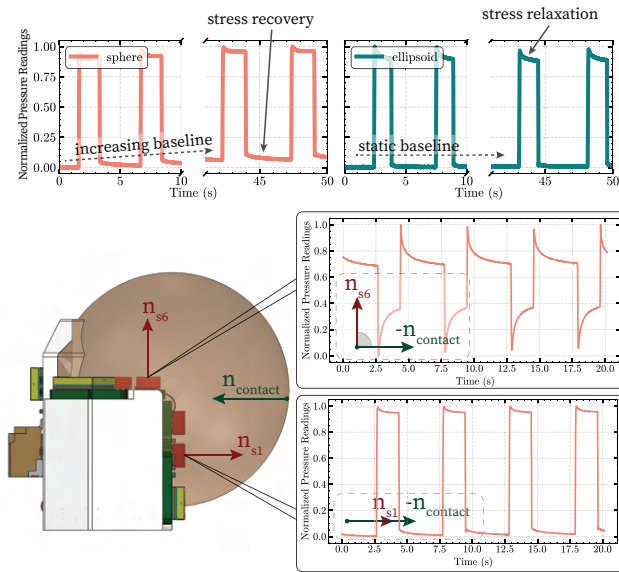


Fig. 2. **Effect of complex sensor morphologies on barometer readings.** *Top:* Normalized barometer readings over repeated contacts for both tactile sensors. With repeated contact cycles on the spherical sensor, residual stress accumulates in the elastomer. *Bottom:* Data from two different barometers in the spherical sensor is shown over repeated contacts in the direction given by n_{contact} . As the angle between a barometer’s sensing axis and the contact normal increases, the modes of mechanical interaction between the elastomer and the barometer become more variable.

at large angles to incoming contact forces, leading to greater variability in raw barometer readings, as shown in the bottom of Fig. 2.

B. Ellipsoid Sensor Design

The second design is a new ellipsoid sensor, developed to explore how geometry affects both sensor dynamics and utility. The ellipsoid contains less elastomer, which decreases the severity of time-dependent effects and results in more stable barometer readings compared to the spherical sensor as shown at the top of Fig. 2. Beyond modeling benefits, the ellipsoid form factor supports new manipulation capabilities. Its elongated shape enables behaviors like sliding between narrow surfaces or burrowing between objects.

The limited time-dependent effects of the ellipsoid form factor simplifies the modeling problem. At the same time, more challenging designs like the spherical sensor remain desirable for their greater surface area and symmetry. Rather than constrain sensor design to favor easier modeling, this work focuses on developing small-scale models to better capture the nonlinear and dynamic behavior of the elastomer.

IV. DATA COLLECTION

Given the importance of dataset quality, our first practice focuses on improving our data collection process. We automate this process using a custom 5-axis gantry, shown in Fig. 3, that combines a 3-axis linear stage holding the tactile sensor with a 2-axis rotary stage carrying an ATI Delta SI-165-15 force sensor. The gantry allows for precise, repeatable trajectories across a wide range of sensor geometries with




Fig. 3. **Data collection hardware setup.** Ground truth sensor data for both sensor geometries is collected with a custom five-axis gantry. The tactile sensor is mounted to the 3-axis linear stage. An ATI force sensor is mounted to the lower rotary stage.

linear and rotary axis resolutions of $21.5\mu\text{m}$ and 0.0015 radians respectively.

The gantry is used to collect data on the tactile sensors to capture both the short-term time-dependent material dynamics associated with the onset and release of contact and the long-term effects of accumulated stress within the sensor. To capture short-term dynamics, we collect data at a fixed frequency and explicit contact and non-contact events are included. To provide robustness against longer term effects, we sample contact events uniformly and randomly across the surface of the sensor. With sufficient sampling, these data collection methods provide the network with a diverse collection of internal sensor states.

The datasets used for training comprise of sampled contact events on the presented tactile sensors, where each contact event is characterized by a contact location, a sample depth, and a duration. We sample 2500 contact events over three hours on both sensor designs at randomly specified locations and depths at 100 Hz. This roughly corresponds to a normal force range of 0 N to 8 N, with low forces over-sampled, and a shear force range of -3 N to +3 N. The contact duration is uniformly randomly sampled between 0.5 and 5 seconds. For each sensor and network, 70% of the data is used for training and 30% is used for testing.

To validate our updated data collection method, we trained the neural network architecture used in our previous work [6] on a new dataset collected on a spherical sensor. Compared to the average force RMSE of 1.575 N reported in the previous study, our improved data collection methods reduced the error to 1.095 N — a 30.4% decrease.

V. BINNED RNN INFERENCE METHOD

With the improved dataset, we now seek to develop lightweight architectures that satisfy the following key requirements:

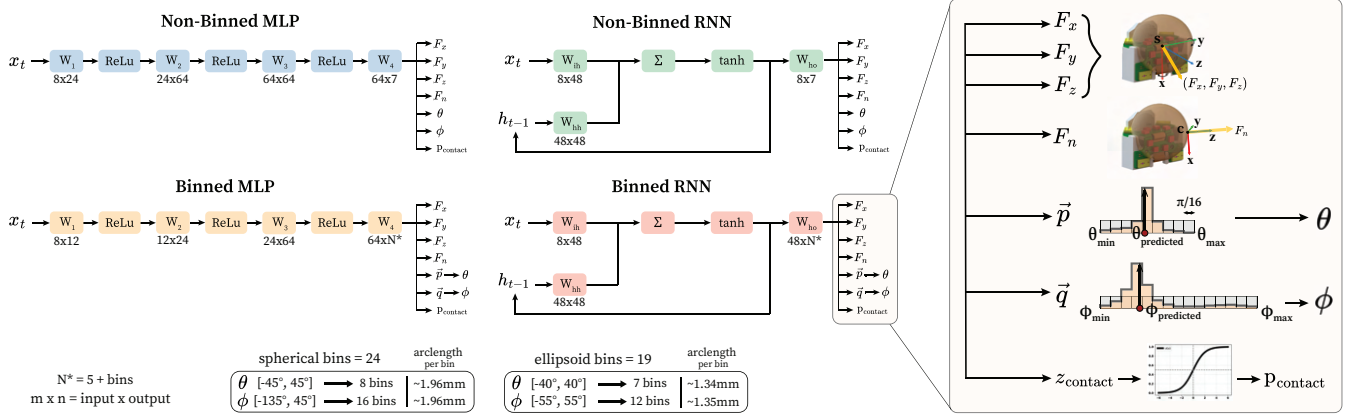


Fig. 4. **Neural network experiments.** The four neural network architectures tested are arranged in a 2x2 grid: columns group models by type (MLP vs. RNN), while rows group them by angle binning method. Binned versions include additional softmaxed output probability vectors, \vec{p} and \vec{q} , which represent discretized distributions of probabilities over θ and ϕ . These are converted to final angle estimates using a weighted average. All architectures output a contact state variable which is converted to a contact flag probability, p_{contact} , by applying a sigmoid function.

- 1) Prediction accuracy is independent of contact history.
- 2) Prediction error is uniform across sensing surface.
- 3) Contact is robustly detected at low applied forces.
- 4) Inference can be run on the sensor's microcontroller.

To meet these requirements, we propose a two layer recurrent neural network with binned angle outputs, shown in Fig. 4. The key components of this network are as follows.

1) *Temporal Context:* Material-induced coupling between subsequent barometer readings can cause past contact events to influence current contact predictions. To mitigate these effects, we use recurrent neural networks (RNNs). These networks maintain a hidden state which is updated at each time step based on both the current input and the previous hidden state, incorporating past interactions into current predictions.

2) *Discretized Outputs:* Uniform prediction accuracy across the sensing surface is a key requirement. However, we observe that regressing contact angles using mean squared error (MSE) loss introduces a bias toward the center of the angle distribution. This arises from implementing MSE loss on noisy or limited data, such as the data from our barometric tactile sensors.

To address this issue, we reframe the angle prediction problem as a classification task by binning the output and training with cross-entropy loss. We use one-hot encoding to encode each angle bin. Instead of directly regressing a scalar value, the model is optimized to learn a probability distribution over the angle bins, conditioned on the barometer readings. With equal sampling of the sensor surface, this method encourages uniform accuracy across the surface.

3) *Contact Detection:* Robust contact detection is particularly useful for downstream robotics tasks [6]. To achieve this, a contact detection flag is introduced as an additional output, p_{contact} .

4) *Lightweight Architecture:* A small model enables the sensor's integrated microcontroller to handle predictions directly and with minimal latency, enabling responsive sensors which are easily replaceable. For this purpose, all models

are constrained to fewer than 6000 parameters. We decided against larger non-recurrent models such as transformers or an MLP with an input history due to this constraint.

A. Binned RNN Architecture Details

Pictured in Fig. 4, the Binned RNN takes in x_t , a vector of current values from all eight barometers, and h_{t-1} , the previous hidden state. The outputs are the 3-axis force in the sensor frame $[F_x, F_y, F_z]$, the normal force F_n , the contact location as Euler angles θ , ϕ , and a contact flag, p_{contact} . Ground truth contact flag data is calculated by thresholding the ATI force sensor data at 0.1 N. The contact prediction is passed through a sigmoid function to indicate the likelihood of contact. Furthermore, normal force magnitude is predicted directly as we found that this was more robust than deriving normal force from the sensor frame force prediction.

As described earlier, the network outputs a distribution indicating the likelihood that the contact angle falls within each bin. To create the bins, the angular prediction range is split into intervals of width $\pi/16$ radians. We found that this width provided reasonable resolution while keeping the network small. The final estimates of the contact angles are calculated as the sum of the bin centers, weighted by the network predictions:

$$\hat{\theta} = \sum_{i=1}^n \hat{p}_i \theta_i \quad \hat{\phi} = \sum_{j=1}^m \hat{q}_j \phi_j \quad (1)$$

While the network does not have an explicit sense of interpolation, we find that bin prediction probabilities generally follow a uni-modal curve centered at the target value.

VI. ABLATIONS

To validate our proposed model, we perform ablations on both the architecture and binning technique. The ablations cover four architectures for each sensor: Non-Binned MLP, Binned MLP, Non-Binned RNN, and Binned RNN (Fig. 4).

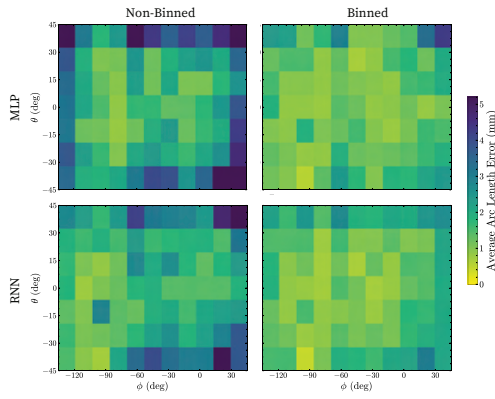


Fig. 5. **Effect of angle binning on contact localization accuracy.** Errors are defined as the arc length between the predicted and target points on the spherical sensor surface. The binned models show reduced errors towards the sensor perimeter in comparison to the non-binned models.

The RNN is trained using a history of 512 samples gathered according to the previously introduced methods. This contains about five seconds of data. We do not test RNNs on variable frequency data as the changing horizon would make the prediction task ambiguous. All networks are trained with a batch size of 50 using the Adam optimizer in PyTorch, with a 20% dropout applied to MLP architectures. The models converge after 200 to 400 epochs. The MLP trains in about one hour and the RNN trains in about two hours on an Nvidia RTX 4090 GPU.

The sections below validate the key components of the Binned RNN, namely the recurrent structure, the binned angle outputs, contact detection, and network size.

1) *Recurrent Neural Network:* Holding both binning type and sensor type constant, Table I shows that the RNN models have a lower RMSE and BCE loss on all prediction targets except for $\phi_{\text{ellipsoid}}$, despite having similar or a fewer numbers of model parameters. This suggests that the RNN is a suitable model for capturing the material dynamics.

2) *Binning Angles:* The binning procedure is validated for the spherical sensor by computing the arc length error (ϵ) as the shortest distance between the target (θ_1, ϕ_1) and predicted (θ_2, ϕ_2) points on a sphere of radius R :

$$\epsilon = R \cos^{-1} (\sin \phi_1 \sin \phi_2 + \cos \phi_1 \cos \phi_2 \cos(\theta_1 - \theta_2)) \quad (2)$$

For each architecture, we plot the average arc length error across the surface of the spherical sensor, shown in Fig. 5. The results show that non-binned networks have higher error toward the sensor edges, while binned networks have the desired, more uniform accuracy across the sensor surface.

The impact of binning on localization precision is illustrated in Fig. 6, which shows that binning consistently reduces the variance of angular error across both sensor types and network architectures. In addition to lowering variance, binning also improves the spatial resolution of the sensor. As shown in Table I, all binned architectures yield lower angular RMSEs than their non-binned counterparts. For example, the Binned RNN achieves RMSEs in (θ, ϕ) of $(0.1516, 0.1842)$

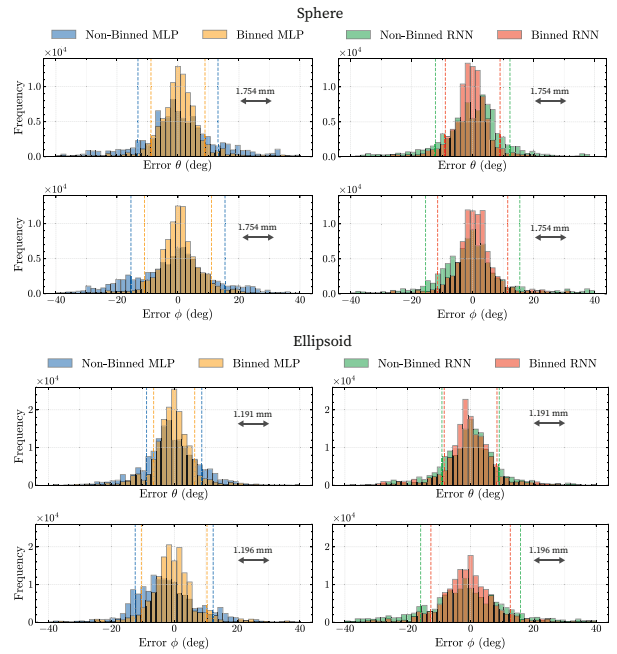


Fig. 6. **Angular prediction error histograms.** Sensor angular prediction error distributions across models for both contact angles. One standard deviation lines are shown on each plot. The binned architectures have consistently lower standard deviations of errors than non-binned.




Fig. 7. **Robustness of contact flag.** Ground truth force data over repeated contacts for the spherical sensor. Shaded regions mark where the predicted contact flag or thresholded predicted normal force detects contact. Explicit prediction of contact is more robust than thresholding predicted forces.

radians on the spherical sensor and $(0.1263, 0.1968)$ radians on the ellipsoid. These angular errors correspond to arc lengths of approximately $(1.52 \text{ mm}, 1.84 \text{ mm})$ and $(0.86 \text{ mm}, 1.35 \text{ mm})$ on the respective sensor surfaces.

3) *Contact Prediction:* Fig. 7 shows that for the spherical sensor, the contact flag, calculated by thresholding the contact prediction at 0.5, is more robust than thresholding force predictions. Force thresholding reliably detects the onset of contact, but the material dynamics introduce ambiguity when determining when contact has ended. In contrast, the contact flag detects the onset and release of contact.

4) *Lightweight Architecture:* Once deployed to the sensor's microcontroller, the RNN evaluates in 7.87 ms while the MLP evaluates in 5.60 ms. These inference times allow the full firmware loop, including sampling of eight barometers and model evaluation, to run at a fixed frequency of 100 Hz.

TABLE I

RMSE AND BINARY CROSS ENTROPY (BCE) SCORES FOR RNN AND MLP ARCHITECTURES ACROSS SENSORS AND PARAMETERIZATIONS

Sensor	Architecture	Binning Type	# Params	RMSE						BCE
				$f_x(N)$	$f_y(N)$	$f_z(N)$	$f_n(N)$	$\theta(\text{rad})$	$\phi(\text{rad})$	p_{contact}
Sphere (Prior Work)*	MLP	Non-Binned	5425	–	–	–	1.0412	0.1785	0.2111	–
Sphere	RNN	Binned	4205	0.4997	0.4807	0.4756	0.6213	0.1516	0.1842	0.0872
		Non-Binned	3127	0.5385	0.5551	0.5565	0.6494	0.2135	0.2698	0.0965
	MLP	Binned	3905	0.7926	0.6371	0.7108	1.0390	0.1644	0.2301	0.1541
		Non-Binned	5555	0.6900	0.6095	0.7071	0.8995	0.2290	0.2709	0.1386
Ellipsoid	RNN	Binned	3960	0.3638	0.3433	0.5527	0.6477	0.1263	0.1968	0.1875
		Non-Binned	3127	0.3372	0.3252	0.5388	0.6126	0.1460	0.2184	0.1754
	MLP	Binned	3580	0.4081	0.3952	0.7605	0.8567	0.1264	0.1875	0.2335
		Non-Binned	5555	0.3958	0.3559	0.6847	0.7856	0.1665	0.2156	0.2834

* In prior work on the sphere [6], $[f_x, f_y, f_z]$ were expressed in the contact frame, whereas in this work they are given in the sensor base frame. Therefore, the RMSE of $f_z(N)$ in the prior work is equivalent to RMSE of $f_n(N)$ in the current work while RMSE f_x and RMSE f_y are incomparable.

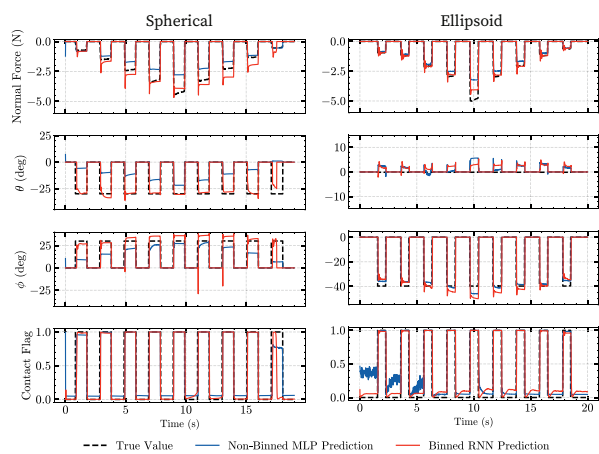


Fig. 8. **Robustness to material aging.** MLP and Binned RNN models are compared on data collected several weeks after initial data collection. The Binned RNN demonstrates increased robustness to distributional shift.

VII. EXPERIMENTAL VALIDATION

To evaluate the robustness and accuracy of our models across both sensors, we perform a series of downstream validation tasks. We first evaluate our system performance on a validation trajectory collected several weeks after the train and test data. Next, we demonstrate the efficacy of our sensor during freeform contact. Finally, we demonstrate the improved accuracy of our sensor during a precise cylinder rolling task.

A. Validation Trajectory

Changes in elastomer material properties over time cause barometer readings to shift between data collection and deployment. To ensure the robustness of our system, we test the models on a trajectory collected several weeks after the initial data collection. This trajectory holds contact angle constant while increasing and decreasing applied force.

To address material aging, a calibration procedure is implemented. Barometer readings are first tared using current unloaded baseline values, then shifted to a fixed reference using target offsets computed from no-load segments in the

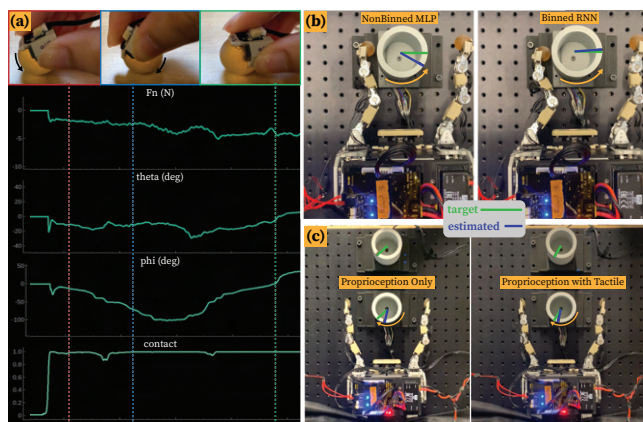


Fig. 9. **Validation experiments.** (a) Freeform sensor rolling along ϕ axis. (b) Cylinder rolling experiments for spherical sensor for target angle of 90° . (c) Performance of proprioceptive versus combined proprioceptive and tactile measurements for interactive control of cylinder orientation.

training set. This is a naive correction aimed to reduce large distributional shifts, but any remaining offset must be handled by the model itself.

To ensure valid predictions, the contact flag was used to mask force and angle measurements when no contact was detected. As shown in Fig. 8, with calibration applied to both models, the Binned RNN produces more accurate force and angle readings on the spherical sensor than the MLP. It also reduces angle drift during contact and decouples force and angle predictions. For the ellipsoid sensor, performance between models is more similar, which is to be expected given the lower elastomer volume.

B. Freeform Contact

Motivated by the unstructured nature of current manipulation tasks, we sought to qualitatively evaluate the sensor performance with the Binned RNN inference method during rapid contacts and smooth rolls across the sensor surface. An illustration of the task is shown in Fig. 9a and demonstrated in the supplementary video¹. Despite not being trained on

¹<https://youtu.be/mi7qqjssirg>

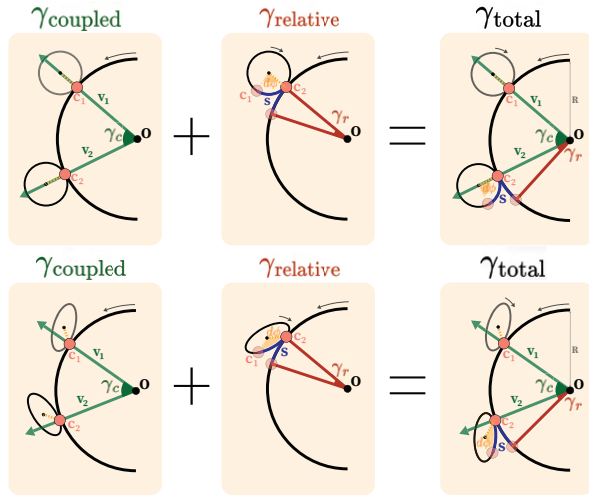


Fig. 10. **Method to estimate cylinder rotation via tactile estimates.** Cylinder rotation is calculated as the sum of coupled movements and relative rotational movements between initial contact point, c_1 , and final contact point, c_2 . A combination of sensor contact location predictions and hand proprioception are used to calculate each component.

hand-rolled data, the model detects light contacts and exhibits smooth transitions in angle predictions as the sensor is rolled.

The model is further tested by deploying it on a newly fabricated sensor of the same design in a similar setup to Fig. 9a. Although the model was not trained on this exact sensor, it still produces smooth transitions in angle predictions as the sensor is rolled along its axes. This preliminary success of model transfer across sensors of the same geometry suggests a better understanding of the underlying pressure distribution within the elastomer. These results indicate that future models could be trained by acquiring a small dataset from a new sensor to fine-tune a preexisting general sensor model.

C. Cylinder Rolling

To evaluate the neural network improvements in predicting contact state over longer duration trials, an algorithm was implemented to roll a cylinder to a desired angle using only tactile and proprioceptive feedback.

1) *Experimental Setup:* A free spinning 3 cm radius cylinder was centered above an AS5048A rotary encoder, used to provide ground truth measurement of the cylinder’s rotation. A multi-jointed robotic finger, shown in Fig. 9b, was affixed to the bench and used to actuate the tactile sensors.

We compare the baseline Non-Binned MLP model against the new proposed Binned RNN to highlight the combined effect of temporal context and binning. For each sensor-model combination, we test target angles of 15° , 30° , 90° , and 120° , with fifteen trials conducted per angle. These angles span partial and continuous rolling motions.

2) *Rolling Algorithm:* The algorithm proceeds as follows: (1) tactile estimates are used to establish contact with the cylinder, (2) a rolling motion is performed while continuously estimating the rotation angle — if contact is lost, the rolling motion is terminated, (3) the finger is commanded to

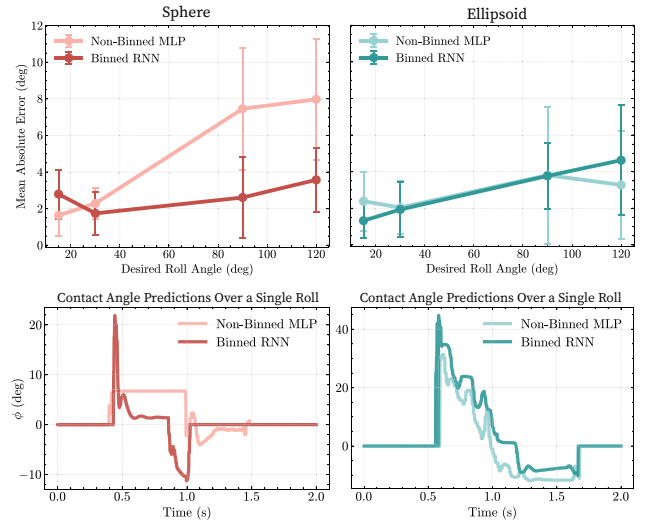


Fig. 11. **Rolling experiments.** *Top:* Mean absolute error by target angle across rolling experiment trials. The RNN consistently outperforms the MLP for the spherical sensor, while performance across angles is more comparable for the ellipsoid sensor. *Bottom:* Contact angle predictions during a roll. The MLP performs poorly on the sphere due to limited sensitivity to contact location changes.

break contact, return to a pre-contact pose, and repeat the process. This cycle continues until the estimated cumulative rotation of the cylinder matches the desired angle. The radius and location of the object were given as inputs. In a future setup, these parameters could be inferred through proprioceptive feedback with a two-fingered grasp.

To estimate the total cylinder rotation angle, the contributions from a coupled rotation mode and relative rotation mode are summed as pictured in Fig. 10 and described by Equations 3 and 4.

$$\gamma_{\text{total}} = \gamma_{\text{coupled}} + \gamma_{\text{relative}} \quad (3)$$

$$\gamma_{\text{coupled}} = \cos^{-1} \left(\frac{\mathbf{v}_1 \cdot \mathbf{v}_2}{\|\mathbf{v}_1\| \|\mathbf{v}_2\|} \right) \quad \gamma_{\text{relative}} = \frac{s}{R} \quad (4)$$

3) *Rolling Results:* Fig. 11 demonstrates that for the spherical sensor, the Binned RNN consistently outperforms the Non-Binned MLP in calculating the cylinder’s rotation during the rolling task, with lower mean absolute error and less variance. One hypothesis for this result is that the Binned RNN is able to resolve smaller changes in contact angle. Since the contact point on the spherical sensor travels a short arc length during the roll, this task demands high angular prediction resolution from the models. The bottom of Fig. 11 illustrates that the MLP often fails to capture small angular changes, especially early in the roll, resulting in periods where its angle estimate remains nearly static, exacerbating error. For the ellipsoid sensor, both models perform similarly, consistent with the results from the validation dataset.

To interactively control the orientation of the cylinder, the rolling algorithm was modified to follow a target orientation provided in real time by a user (Fig. 9c). Cylinder rotation was approximated first using proprioception only and then with added tactile estimates. As shown in the supplementary

video, proprioception provides a rough estimate of cylinder orientation, but the addition of tactile input closes the gap between the estimated and the ground truth positions.

VIII. CONCLUSION

In this work, we propose a set of recommendations to enable robust contact estimation models for barometric tactile sensors with strong viscous effects: (1) collect data at a fixed frequency for time-dependent inference, sample uniformly and randomly over the sensor surface to reduce bias, and include both contact and non-contact samples to capture transients; (2) discretize continuous contact localization parameters to reduce bias toward the mean and improve model learning; and (3) incorporate temporal context to account for the viscoelastic effects of the elastomer. We use two custom designed barometric tactile sensors as case studies and perform ablations to validate our recommendations. Finally, we deploy our models on each sensor and test robustness on a separately collected validation trajectory, demonstrate free-form, out of distribution inference, and showcase a long horizon cylinder rotation task to exhibit the stability and accuracy of the sensor measurements. With these methods we are one step closer to unlocking the promise of a cheap, customizable, yet highly capable tactile sensor.

REFERENCES

- [1] Akihiko Yamaguchi and Christopher G Atkeson. “Recent progress in tactile sensing and sensors for robotic manipulation: can we turn tactile sensing into vision?” In: *Advanced Robotics* 33.14 (2019).
- [2] Mike Lambeta et al. “Digit: A novel design for a low-cost compact high-resolution tactile sensor with application to in-hand manipulation”. In: *IEEE Robotics and Automation Letters* 5.3 (2020).
- [3] Lunwei Zhang, Yue Wang, and Yao Jiang. “Tac3D: A Novel Vision-based Tactile Sensor for Measuring Forces Distribution and Estimating Friction Coefficient Distribution”. In: (Feb. 2022).
- [4] Mahmoud Meribout et al. “Tactile sensors: A review”. In: *Measurement* 238 (2024).
- [5] Ilker S. Bayer. “MEMS-Based Tactile Sensors: Materials, Processes and Applications in Robotics”. In: *Micromachines* 13.12 (2022).
- [6] Andrew SaLoutos et al. “Design of a multimodal fingertip sensor for dynamic manipulation”. In: *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2023.
- [7] Jacob W Guggenheim et al. “Robust and inexpensive six-axis force–torque sensors using MEMS barometers”. In: *IEEE/ASME Transactions on Mechatronics* 22.2 (2017).
- [8] Pedro Piacenza, Sydney Sherman, and Matei Ciocarlie. “Data-driven super-resolution on a tactile dome”. In: *IEEE Robotics and Automation Letters* 3.3 (2018).
- [9] Gyuwon Kim and Donghyun Hwang. “BaroTac: Barometric three-axis tactile sensor with slip detection capability”. In: *Sensors* 23.1 (2022).
- [10] Zhanat Kappassov, Juan-Antonio Corrales, and Véronique Perdereau. “Tactile sensing in dexterous robot hands”. In: *Robotics and Autonomous Systems* 74 (2015).
- [11] Yaroslav Tenzer, Leif P Jentoft, and Robert D Howe. “The feel of MEMS barometers: Inexpensive and easily customized tactile array sensors”. In: *IEEE Robotics & Automation Magazine* 21.3 (2014).
- [12] Christian Reeks et al. “Angled sensor configuration capable of measuring tri-axial forces for pHRI”. In: *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2016.
- [13] Abhinav Grover et al. “Learning to detect slip with barometric tactile sensors and a temporal convolutional neural network”. In: *2022 International Conference on Robotics and Automation (ICRA)*. IEEE. 2022.
- [14] Adam J Spiers et al. “Single-grasp object classification and feature extraction with simple robot hands and tactile sensors”. In: *IEEE Transactions on Haptics* 9.2 (2016).
- [15] Thomas De Clercq, Anatolii Sianov, and Guillaume Crevecoeur. “A soft barometric tactile sensor to simultaneously localize contact and estimate normal force with validation to detect slip in a robotic gripper”. In: *IEEE Robotics and Automation Letters* 7.4 (2022).
- [16] Aiden Shaevitz, Matthew L Johnston, and Joseph R Davidson. “Design, characterization, and modeling of barometric tactile sensors for underwater applications”. In: *2023 IEEE International Conference on Soft Robotics (RoboSoft)*. IEEE. 2023, pp. 1–6.
- [17] Jian Hou, Xin Zhou, and Adam Spiers. “Location and orientation super-resolution sensing With a cost-efficient and repairable barometric tactile sensor”. In: *IEEE Transactions on Robotics* 41 (2025).
- [18] Risto Kōiva et al. “Barometer-based tactile skin for anthropomorphic robot hand”. In: *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2020.
- [19] Lindsay Epstein et al. “Bi-modal hemispherical sensors for dynamic locomotion and manipulation”. In: *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2020.
- [20] Thijs Van Hauwermeiren et al. “Integrated barometric pressure sensors on legged robots for enhanced tactile exploration of edges”. In: *IEEE Robotics and Automation Letters* (2024).
- [21] Thomas De Clercq et al. “Soft barometric tactile sensor utilizing iterative pressure reconstruction”. In: *IEEE Access* (2024).
- [22] Meng Yee Chuah and Sangbae Kim. “Improved normal and shear tactile force sensor performance via least squares artificial neural network (lsann)”. In: *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2016.