

Accurate Contact Localization and Indentation Depth Prediction With an Optics-based Tactile Sensor

Pedro Piacenza¹, Weipeng Dang², Emily Hannigan¹, Jeremy Espinal³, Ikram Hussain³,
Ioannis Kymissis² and Matei Ciocarlie¹

Abstract—Traditional methods to achieve high localization accuracy with tactile sensors usually use a matrix of miniaturized individual sensors distributed on the area of interest. This approach usually comes at a price of increased complexity in fabrication and circuitry, and can be hard to adapt for non planar geometries. We propose to use low cost optic components mounted on the edges of the sensing area to measure how light traveling through an elastomer is affected by touch. Multiple light emitters and receivers provide us with a rich signal set that contains the necessary information to pinpoint both the location and depth of an indentation with high accuracy. We demonstrate sub-millimeter accuracy on location and depth on a 20mm by 20mm active sensing area. Our sensor provides high depth sensitivity as a result of two different modalities in how light is guided through our elastomer. This method results in a low cost, easy to manufacture sensor. We believe this approach can be adapted to cover non-planar surfaces, simplifying future integration in robot skin applications.

I. INTRODUCTION

Tactile sensors for robot manipulators can be analyzed and quantified based on multiple performance criteria. From an operational perspective, these include high accuracy in establishing both the *location* of a contact and the *magnitude* of the applied force. In particular, good signal-to-noise ratio is desirable for both the contact forces that characterize incipient contact and the larger forces encountered during manipulation. From a manufacturing perspective, achieving *coverage* of potentially irregular, non-flat surfaces is also important for application to robotic fingers and palms.

One way to achieve accuracy and good coverage is by using individual taxels distributed over the surface that must be sensorized; however, this imposes miniaturization constraints on each taxel. Matrix addressing for taxel arrays reduces the complexity of the circuit and allows part of it to be implemented on a printed circuit board, but imposes 2D structure on the sensor. As recent reviews point out, achieving multiple such performance metrics traditionally leads to manufacturing difficulties and system constraints that prohibit large-scale deployment [1], [2].

Our approach is to build a tactile sensor as a continuous volume of a transparent polymer, with light emitters and receivers embedded along the perimeter (Fig. 1). Indentation




Fig. 1: Sensor consists of a square mold where LEDs and photodiodes are edge-mounted and the cavity is filled with an elastomer. We measure light transport through the elastomer to learn the location and the depth of an indentation.

of the sensor area affects how light from the emitters is transported through this medium, producing a signal that is measured by the receivers. Throughout this study, we use indentation depth as a proxy for contact force, based on a known stiffness curve for the constituent material. This method natively lends itself to covering large areas with simple-to-manufacture, low-cost sensors.

In order to achieve the performance goals stated above, namely high accuracy in both localization and indentation depth prediction, we rely on two key ideas, which are also the main contributions of this paper:

- *Leverage multiple modes of interaction between an indenting probe and a light-transporting medium in order to increase depth prediction accuracy throughout the operating range.* As we will detail in the paper, indentation of the medium affects light transport in at least two ways. First, during initial contact, the probe alters the geometry of the surface and changes the refraction of the light rays. Second, as indentation becomes deeper, the probe blocks direct paths between emitters and receivers. We design our sensor to use both modes in continuous fashion, resulting in a good sensitivity throughout a wide range of indentation depths.
- *Use data-driven methods to directly learn the mapping between a rich signal set extracted from the sensor and our variables of interest.* For our sensor, such data set can be obtained by measuring the signal between every emitter and every receiver. In the past, we have used an all-pairs approach on a piezoresistance-based sensor and showed it can lead to high localization accuracy [3].

¹Department of Mechanical Engineering, Columbia University, New York, NY 10027, USA.

{pp2511, ejh2192, matei.ciocarlie}@columbia.edu

²Department of Electrical Engineering, Columbia University, New York, NY 10027, USA.

wd2265@columbia.edu, johnkym@ee.columbia.edu

³Columbia Engineering ENG Summer Research Program.

The result of using these methods is a tactile pad that exhibits desirable performance characteristics in accuracy and sensitivity, while using a simple manufacturing method and low-cost components. While not explicitly tested here, we believe that both the fabrication technique and the data-driven signal processing approach also lend themselves to constructing pads of irregular three-dimensional geometry. Furthermore, mapping the signal between every emitter and receiver produces a rich signal set with relatively few wires. Both of these characteristics could enable easier integration into robotic fingers and palms, which is our directional goal.

II. RELATED WORK

The use of optics for tactile sensing is not new, and has a long history of integration in robotic fingers and hands. Early work by Begej demonstrated the use of CCD sensors recording light patterns through a robotic tip affected by deformation [4]. More recently Lepora and Ward-Cherrier [5] showed how to achieve super-resolution and hyperacuity with a CCD-based touch sensors integrated into a fingertip. Johnson and Adelson used color-coded 3D geometry reconstruction to retrieve minute surface details with an optics-based sensor [6]. These studies share a common concept of a CCD array imaging a deformed fingertip from the inside, requiring that the array be positioned far enough from the surface in order to image the entire touch area. In our approach, the sensing elements are fully distributed, allowing for coverage of large areas and potentially irregular geometry. Work by Polygerinos et al. [7] use the deformation of an optic fiber to create a force transducer. This approach has the advantage that the sensing electronics do not have to be located close to the contact area.

In our work, we take advantage of multiple modes of light transport through an elastomer to increase the sensitivity of the sensor. In recent work, Patel and Correll took advantage of reflection and refraction to build an IR touch sensor that also functions as a proximity sensor [8]. Their work however does not provide means to also localize contact.

We perform contact localization by combining signals from multiple emitter-receiver pairs, a technique which we previously used in the context of a piezoresistive sensor [3]. Other sensors also use a small number of underlying transducers to recover richer information about the contact. For example, work in the ROBOSKIN project showed how to calibrate multiple piezocapacitive transducers [9], used them to recover a complete contact profile [10] using an analytic model of deformation, and finally used such information for manipulation learning tasks [11]. Our localization method is entirely data driven and makes no assumptions about the underlying properties of the medium, which could allow coverage of more complex geometric surfaces.

Our localization approach shares some of the same goals of techniques such as super-resolution and electric impedance tomography. Van den Heever et al. [12] used a similar algorithm to super-resolution imaging, combining several measurements of a 5 by 5 force sensitive resistors array into an overall higher resolution measurement. Lepora

and Ward-Cherrier [5] and Lepora et al. [13] used a Bayesian perception method to obtain a 35-fold improvement of localization acuity (0.12mm) over a sensor resolution of 4mm. Electric impedance tomography (EIT) is used to estimate the internal conductivity of an electrically conductive body by virtue of measurements taken with electrodes placed on the boundary of said body. While originally used for medical applications, EIT techniques have been applied successfully for manufacturing artificial sensitive skin for robotics [14]–[16], although with lower spatial resolution than other methods.

We rely on data-driven methods to learn the behavior of our sensors; along these lines, we note that machine learning for manipulation based on tactile data is not new. Ponce Wong et al. [17] learned to discriminate between different types of geometric features based on the signals provided by a previously developed [18] multimodal touch sensor. Current work by Wan et al. [19] relates tactile signal variability and predictability to grasp stability using recently developed MEMS-based sensors [20]. With traditional tactile arrays, Dang and Allen [21] successfully used an SVM classifier to distinguish stable from unstable grasps in the context of robotic manipulation using a Barrett Hand. Bekiroglu et al. [22] also studied how grasp stability can be assessed based on tactile sensory data using machine-learning techniques. In similar fashion, both Saal et al. [23] and Tanaka et al. [24] used probabilistic models on tactile data to estimate object dynamics and perform object recognition respectively. However, most of this work is based on arrays built on rigid substrates and thus unable to provide full coverage of complex geometry. In contrast, we apply our methods to the design of the sensor itself, and believe that developing the sensor simultaneously with the learning techniques that make use of the data can bring us closer to achieving complete tactile systems.

III. TACTILE SENSING METHOD

The fundamental sensing unit of our approach is comprised of a light emitting diode (LED) and a photodiode receiver, edge-mounted around a sensing area which is filled with polydimethylsiloxane (PDMS) (Fig. 2a). While we ultimately use sensors with multiple emitters and receivers, in this section we focus on a single emitter-receiver pair in order to discuss the underlying transduction mechanism; we will return to complete sensor design in the next section.

A. Light transport and interaction modes

The core transduction mechanism relies on the fact that as a probe indents the surface of the sensor, light transport between the emitter and the receiver is altered, changing the signal reported by the receiver. Consider the multiple ways in which light from the LED can reach the opposite photodiode: through a direct path or through a reflection. Of particular interest to us are reflections off the surface of the sensor: based on Snell’s law, due to different refractive indices of the elastomer and air, light rays hitting the surface below the critical angle are reflected back into the elastomer.




Fig. 2: The first mode of detection (b) is the result of light scattering and surface deformation. The second mode (c) is the result of the indenter tip physically blocking the direct path of light.

As the probe makes initial contact with the sensor surface, the elastomer-air interface is removed from the area of contact; furthermore, surface normals are immediately disturbed. This changes the amount of light that can reach the diode via surface reflection (Fig. 2b). *This is the first mode of interaction that our transduction method captures.* It is highly sensitive to initial contact, and requires very little penetration depth to produce a strong output signal.

As the depth of indentation increases, the object penetration into the PDMS also starts to block the light rays that were reaching the photodiode through a direct, line-of-sight path (Fig. 2c). *This is our second mode of interaction.* To produce a strong signal, the probe must reach deep enough under the surface where it blocks a significant part of the diode's surface from the LED's vantage point.

We note that other light paths are also possible between the emitter and receiver. The interface between the clear elastomer and the holding structure (the bottom and side walls of the cavity) can also give rise to reflections. In practice, we have found that the elastomer and the holding plastic exhibit bonding/unbonding effects at a time scale of 5-10s when indented, creating unwanted hysteresis.

To eliminate such effects, we coat the bottom of the sensor with a 1mm layer of elastomer saturated with carbon black particles (shown in Fig. 2 by a thick black line). This eliminates bottom surface reflections, and exhibits no adverse effects, as the clear elastomer permanently bonds with the carbon black-filled layer. There are still possible reflections off the walls of the sensor; while we do not explicitly consider their effects, they can still produce meaningful

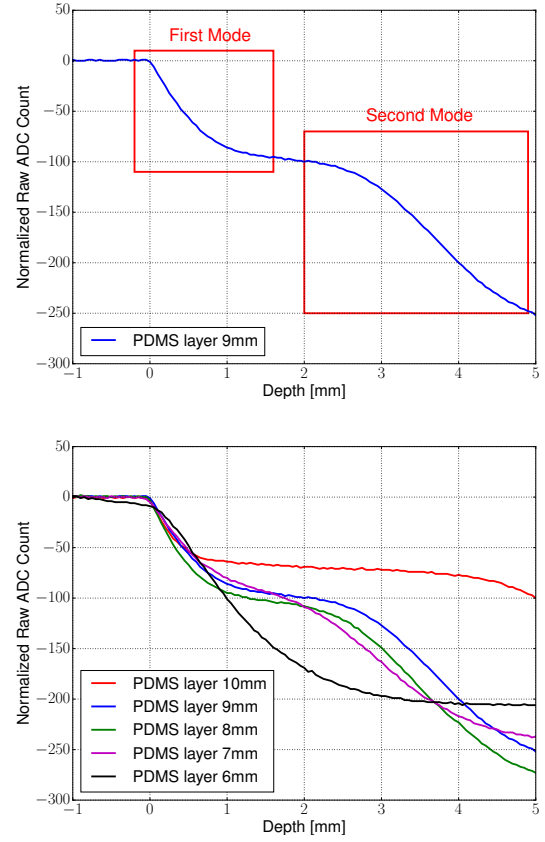


Fig. 3: The first mode of detection happens upon light contact and manifests as a sudden drop in the signal. The second mode is activated with a heavier contact, in this particular case after a depth of approximately 2mm

signals that are used by data-driven mapping algorithm.

B. Effective operating range and prototype construction

We would like our sensor to take advantage of both operating modes described above, noting that one is highly sensitive to small indentations while the other provides a strong response to deeper probes. *We thus aim to design our sensor such that these two modes are contiguous as the indentation depth increases.* The goal is to obtain high sensitivity throughout the operating range of the sensor. The key geometric factor affecting this behavior is the height of the elastomer layer, which we determine experimentally.

We constructed multiple 3D printed (black ABS material) molds with LEDs (SunLED model XSCWD23MB) placed 20mm away from the photodiode (Osram SFH 206K) on opposing sides of a square mold. These sensors have 3 LEDs on one side of the mold, and 3 corresponding photodiodes directly in front of the LEDs. The mold was filled with a transparent elastomer (PDMS, Sylgard 184, Dow Corning). PDMS and air have approximate refractive indexes of 1.4 and 1.0 [25] respectively, leading to a critical angle of 45 degrees. This means that if light hits the boundary between PDMS and air at an angle greater than 45 degrees with respect to the surface normal, the ray is reflected back into the PDMS

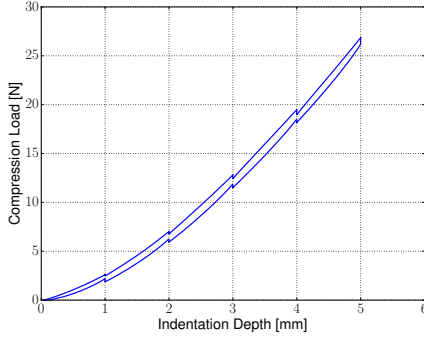


Fig. 4: Load vs indentation depth for a 1:20 ratio of curing agent to PDMS. Measurements were taken by advancing or retracting the probe in 1mm steps separated by 10s pauses

where it can be detected by the photodiode.

Prototypes with PDMS layers over 10mm showed the presence of a deadband: after a certain threshold depth the photodiode signal does not change as we indent further down until you indent deep enough to activate the second mode. To make these two modes continuous we build a set of sensors where we vary the thickness of the PDMS layer. Results can be visualized in Fig. 3. Based on these results, while the 7mm layer provides the best continuity between our two modes, the 8mm layer gives good continuity while also producing a slightly stronger signal when indented. We build our subsequent sensor to have an 8mm PDMS layer.

Another parameter to choose when building the sensor is the stiffness of the PDMS. This parameter lets us directly manipulate the mapping between indentation depth and indentation force. The stiffness of PDMS is determined by the weight ratio between the curing agent and the polymer itself. Fig. 3 corresponds to a ratio of 1:20 curing agent to PDMS, which we used in our sensors. Fig. 4 shows the mapping between indentation depth up to 5mm and force for the 6mm hemispherical tip used in our experiments.

IV. COMPLETE SENSOR DESIGN

While sensitivity to a large range of indentation depths (and forces) is important for applications in manipulation, it is not sufficient. The ability to localize touch accurately on a 2D surface embedded in 3D space is also critical. To achieve this goal, we construct our sensors with numerous light emitters and receivers mounted around the perimeter (Fig. 5). This gives rise to numerous emitter-receiver pairs, each behaving as a unit described in the previous section.

The multi-pair approach gives us a very rich signal set, with cardinality equal to the number of emitters multiplied by the number of receivers. We make the assumption that an indentation anywhere on the sensor will affect multiple such signals. We then use a data-driven approach to directly learn how these signals map to our variables of interest, such as indentation location and depth.

We validate this concept on a sensor comprised of 8 LEDs and 8 photodiodes arranged in an alternating pattern and mounted in sockets along the central cavity walls. To build




Fig. 5: The sensor design enables several receivers to be excited by a single emitter. This one-to-many relationship provides a rich set of signals when switching through all of our emitters and measuring the signals on all the receptors

the sensor we use a 3D printed square mold with exterior dimensions of 45mm x 45mm. The cavity in the mold is 32mm x 32mm. Light travels from emitters to receivers via multiple paths that cover the sensing area. This way any LED is able to excite several photodiodes (Fig. 5).

To determine the location and depth of the indentation, we will read signals from all the photodiodes as different LEDs turn on. Having 8 LEDs gives rise to 8 signals for each photodiode; plus an additional signal with all LEDs turned off. This last signal allows us to measure the ambient light captured by each diode, and leads to a total of 9 signals per diode. One important consideration with a sensor where the sensing units are exposed to ambient light is to incorporate this information such that the sensor can perform consistently in different lighting situations. To achieve this, we use the state where all LEDs are OFF as a baseline that gets removed from every other state at each sampling time.

An Arduino Mega 2560 handles switching between our 9 states, and taking analogue readings of each photodiode. The photodiode signal is amplified through a standard transimpedance amplifier circuit, and each LED is driven at full power using an NPN bipolar junction transistor. The resulting sampling frequency with this setup is 60Hz.

V. DATA COLLECTION PROTOCOL

Data collection is performed using a planar stage (Marzhauser LStep) and a linear probe located above the stage to indent vertically on the sensor with a 6mm hemispherical tip. The linear probe is position-controlled and the reference level is set manually such that the indenter tip barely makes contact with the sensor surface. The linear probe does not have force sensing capability, hence we use indentation depth as a proxy for indentation force.

Two patterns are used to indent our sensor. The *grid indentation pattern* indents the sensor on a 2mm regular grid. Taking into account the diameter of our tip, plus a 3mm margin such that we don't indent directly next to an edge, this results in 121 indent locations distributed over a 20mm x 20mm area. The order of indentation within this grid is randomized. The *random indentation pattern* indents the sensor in randomly generated locations within its workspace.

At each location we follow the same protocol. Consider the sensor surface to be the reference level, and positive depth values correspond to the indenter tip going deeper into the sensor. To discriminate touch vs. non-touch conditions, we collect data at both negative and positive depths. For depths between -10mm and -1mm , we collect one data point every 1mm . The indenter then goes down to a depth of 5.0mm taking measurements every 0.1mm . The same procedure is mirrored with the indenter tip retracting.

Each measurement i results in a tuple of the form $\Phi_i = (x_i, y_i, d_i, p_{j=1}^1, \dots, p_{j=1}^8, \dots, p_{j=9}^1, \dots, p_{j=9}^8)$ where (x_i, y_i) is the indentation location in sensor coordinates, d_i is the depth at which the measurement was taken and (p_j^1, \dots, p_j^8) corresponds to the readings of our 8 photodiodes at state $j \in [1, 9]$. For each diode, states 1 through 8 correspond to each one of the 8 LEDs being ON, and state 9 corresponds to the case where all the LEDs are OFF. We thus have a total of 75 numbers comprised in each tuple Φ_i . These tuples are analyzed as described in the following section.

VI. ANALYSIS AND RESULTS

Our main objective is to learn the mapping between our photodiode readings (p_j^1, \dots, p_j^8) to the indentation location and depth (x_i, y_i, d_i) .

We have found that we obtain higher performance by splitting this problem into two components. First, we use a classifier to determine if touch is occurring; this classifier is trained on both data points with $d_i < 0$ and $d_i \geq 0$. We use a linear SVM as our classifier of choice for this problem.

If the SVM classifier predicts touch is occurring, we use a second stage regressor that outputs predicted values for (x_i, y_i, d_i) . This regressor is trained only on training data with $d_i \geq 0$. We use a kernelized ridge regressor with a Laplacian kernel and use half of the training data to calibrate the ridge regression tuning factor λ and the kernel bandwidth γ through grid search. Results presented in this section were obtained with $\lambda = 2.15e^{-4}$ and $\gamma = 5.45e^{-4}$.

To train our predictors we collected four grid pattern datasets, each consisting of 121 indentations, and each indentation containing 161 datapoints at different depths. Aiming for robustness to changes in lighting conditions, two of these datasets were collected with the sensor exposed to ambient light and the other two datasets were collected in darkness. The feature space used for training has a dimensionality of 64, since the all LEDs OFF state is first subtracted from all other signals and not used as a stand-alone feature.

TABLE I: Touch vs. no touch classification success rate

Ground Truth	Depth Value	Ambient Light test set	Dark test set
No touch	-0.4 mm	1.0	1.0
No touch	-0.2 mm	0.99	1.0
Touch	0 mm	0.04	0.10
Touch	0.2 mm	0.32	0.33
Touch	0.4 mm	0.51	0.65
Touch	0.6 mm	0.76	0.90
Touch	0.8 mm	0.96	0.96
Touch	1.0 mm	0.98	0.98

The metric used to quantify the success of our regressor is the magnitude of the error for both the localization and depth accuracy. In case of the classifier, the metric is the percentage of successful predictions. Each test dataset is collected on a *random indentation pattern* which contains 100 indentation events. The results presented here are those obtained by testing our models against two different test datasets: one collected in ambient light and another collected in the dark.

Classification results in the region of interest for the touch vs no-touch case are summarized in Table I. Note that these results are sliced at a certain depth, and they aggregate the classification performance across all locations. With both test datasets, the classifier has difficulty detecting touch at 0mm, where the tip of the indenter is barely making contact with the sensor. However, at 0.6mm depth the dark test dataset already provides a 90% success rate in the classification, while the ambient light test dataset provides 76% success rate. According to the mapping presented on Fig. 4, a depth of 0.6mm corresponds to an indentation force of approximately 2 Newtons. This represents the minimum indenter force our sensor is capable of detecting.

Regression results for localization on the light dataset are presented on Fig. 7 for a few representative depths. At depth 0.1mm the signals are still not good enough to provide accurate localization, but as we indent further down localization improves well beyond sub-millimeter accuracy.

The regression for depth can only be visualized for a specific location. Fig. 6 shows the performance at three different locations in our ambient light test dataset. Since this regressor is trained only on contact data, it is not able to predict negative depths which causes error to be greater at depths close to 0mm. At 0.5mm and deeper, the depth prediction shows an accuracy of below half a millimeter.

Accuracy in localization and depth for both of our test datasets is presented in more detail in Tables II and III.

VII. CONCLUSIONS

In this paper, we explore the use of light transport through a clear medium as the core transduction method for a tactile pad. We develop touch sensors with multiple light emitter and receivers around the perimeter of a clear elastomer pad. By measuring light transport between each emitter-receiver pair we collect a very rich signal set characterizing indentation performed anywhere on the sensor surface. We then mine this data to learn the mapping between our signals and the variables characterizing indentation. The result is a low-cost, easy to manufacture tactile pad that displays many desirable characteristics: sensitivity to initial contact, and sub-millimeter accuracy in predicting both indentation location and depth throughout most of the operating range.

Our sensor achieves sub-millimeter accuracy over a 400mm^2 workspace, while reporting indentation depth within 0.5mm throughout most of the operating range. This sensitivity is the result of leveraging two different modes in which the indenter affects light transport in the PDMS: the first mode being capable of detecting small indentations during initial contact, and the second mode detecting larger




Fig. 6: Regression results for depth prediction at a central location (a), an edge location close to the x-axis (b) and another edge location close to the y-axis of the sensor (c). These are some of the random locations in our test dataset

indentations for stronger contact. We use depth here as a proxy for indentation force, based on a known stiffness curve for our sensor; it is also possible to adjust the PDMS layer thickness and stiffness. A stiffer PDMS layer will require a larger force to activate the second sensing mode.

Future work will be focused on adapting this method for arbitrary geometries, such that any surface can be covered with this method. There is also the possibility of learning additional variables, such as shear forces or torsional friction. Other important factors were also left out of this preliminary study, such as an analysis of temporary properties, especially regarding signal drift that might arise over long term periods. Repeatability, hysteresis and sensitivity to environmental factors are also important metrics that should be further analyzed. Another improvement could be to incorporate multi-touch capability into our detection algorithms.

We believe that ultimately the number of variables that can be learned and how accurately we can determine those variables depends on the raw data that can be harvested from the sensor. Increasing the number of sensing units or even incorporating different sensors embedded into the elastomer can extend the sensing modalities in our sensor. Different learning methods like deep neural networks might also improve the performance or capability of the sensor. All these ideas will be explored in future work.

REFERENCES

- [1] R. Dahiya, G. Metta, M. Valle, and G. Sandini, "Tactile sensing: From humans to humanoid," *IEEE Trans on Robotics*, vol. 26, no. 1, 2010.
- [2] H. Yousef, M. Boukallel, and K. Althoefer, "Tactile sensing for dexterous in-hand manipulation in robotics: A review," *Sensors and Actuators A: physical*, vol. 167, no. 2, 2011.
- [3] P. Piacenza, Y. Xiao, S. Park, I. Kyriassis, and M. Ciocarlie, "Contact localization through spatially overlapping piezoresistive signals," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, 2016.
- [4] S. Begej, "Planar and finger-shaped optical tactile sensors for robotic applications," *IEEE J. Robot. Autom.*, vol. 4, no. 5, 1988.
- [5] N. F. Lepora and B. Ward-Cherrier, "Superresolution with an optical tactile sensor," in *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*, Sept 2015, pp. 2686–2691.
- [6] F. Schneider, J. Draheim, R. Kamberger, and U. Wallrabe, "Process and material properties of polydimethylsiloxane (pdms) for optical mems," *Sensors and Actuators A: Physical*, vol. 151, no. 2, pp. 95–99, 2009.
- [7] P. Polygerinos, D. Zbyszewski, T. Schaeffer, R. Razavi, L. D. Seneviratne, and K. Althoefer, "Mri-compatible fiber-optic force sensors for catheterization procedures," *IEEE Sensors J.*, vol. 10, 2010.
- [8] N. C. R. Patel, "Integrated force and distance sensing for robotic manipulation using elastomer-embedded commodity proximity sensors," in *Robotics: Science and Systems*, 2016.
- [9] G. Cannata, S. Denei, and F. Mastrogiovanni, "Towards automated self-calibration of robot skin," in *IEEE Intl. Conference on Robotics and Automation*, 2010, pp. 4849–4854.
- [10] L. Muscare, L. Seminara, F. Mastrogiovanni, M. Valle, M. Capurro, and G. Cannata, "Real-time reconstruction of contact shapes for large area robot skin," in *IEEE Int. Conf. Robot. Autom.*, 2013.
- [11] B. Argall and A. Billard, "Learning from demonstration and correction via multiple modalities for a humanoid robot," in *BIO Web of Conferences*, vol. 1. EDP Sciences, 2011, p. 00003.
- [12] D. J. van den Heever, K. Schreve, and C. Scheffer, "Tactile sensing using force sensing resistors and a super-resolution algorithm," *IEEE Sensors Journal*, vol. 9, no. 1, pp. 29–35, Jan 2009.
- [13] N. F. Lepora, U. Martinez-Hernandez, M. Evans, L. Natale, G. Metta, and T. J. Prescott, "Tactile superresolution and biomimetic hyperacuity," *IEEE Trans. on Robotics*, vol. 31, no. 3, pp. 605–618, June 2015.
- [14] A. Nagakubo, H. Alirezai, and Y. Kuniyoshi, "A deformable and deformation sensitive tactile distribution sensor," in *IEEE Intl. Conf. on Robotics and Biomimetics*. IEEE, 2007, pp. 1301–1308.
- [15] Y. Kato, T. Mukai, T. Hayakawa, and T. Shibata, "Tactile sensor without wire and sensing element in the tactile region based on eit method," in *Sensors, 2007 IEEE*. IEEE, 2007, pp. 792–795.
- [16] D. S. Tawil, D. Rye, and M. Velonaki, "Touch modality interpretation for an eit-based sensitive skin," in *Robotics and Automation (ICRA), 2011 IEEE International Conference on*. IEEE, 2011, pp. 3770–3776.
- [17] R. D. Ponce Wong, R. B. Hellman, and V. J. Santos, "Haptic exploration of fingertip-sized geometric features using a multimodal tactile sensor," in *Proc SPIE Defense, Security and Sensing Conference*, 2014.
- [18] N. Wettels, V. J. Santos, R. S. Johansson, and G. E. Loeb, "Biomimetic tactile sensor array," *Adv Robot*, vol. 22, no. 8, 2008.
- [19] Q. Wan, R. P. Adams, and R. Howe, "Variability and predictability in tactile sensing during grasping," in *IEEE International Conference on Robotics and Automation*, 2016.
- [20] Y. Tenzer, L. Jentoft, and R. Howe, "The feel of mems barometers: Inexpensive and easily customized tactile array sensors," *IEEE Robotics & Automation Magazine*, vol. 21, no. 3, 2014.
- [21] H. Dang and P. K. Allen, "Grasp adjustment on novel objects using tactile experience from similar local geometry," in *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ Intl. Conf. on*, pp. 4007–4012.
- [22] Y. Bekiroglu, J. Laaksonen, J. A. Jorgensen, V. Kyrki, and D. Kragic, "Assessing grasp stability based on learning and haptic data," *IEEE Transactions on Robotics*, vol. 27, no. 3, pp. 616–629, 2011.
- [23] H. P. Saal, J. A. Ting, and S. Vijayakumar, "Active estimation of object dynamics parameters with tactile sensors," in *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ Intl. Conf. on*, pp. 916–921.
- [24] D. Tanaka, T. Matsubara, K. Ichien, and K. Sugimoto, "Object manifold learning with action features for active tactile object recognition," in *IEEE/RSJ Intl. Conf. on Intelligent Rob. and Sys.*, 2014.
- [25] M. K. Johnson and E. H. Adelson, "Retrographic sensing for the measurement of surface texture and shape," in *Computer Vision and Pattern Recognition (CVPR)*, 2009, pp. 1070–1077.

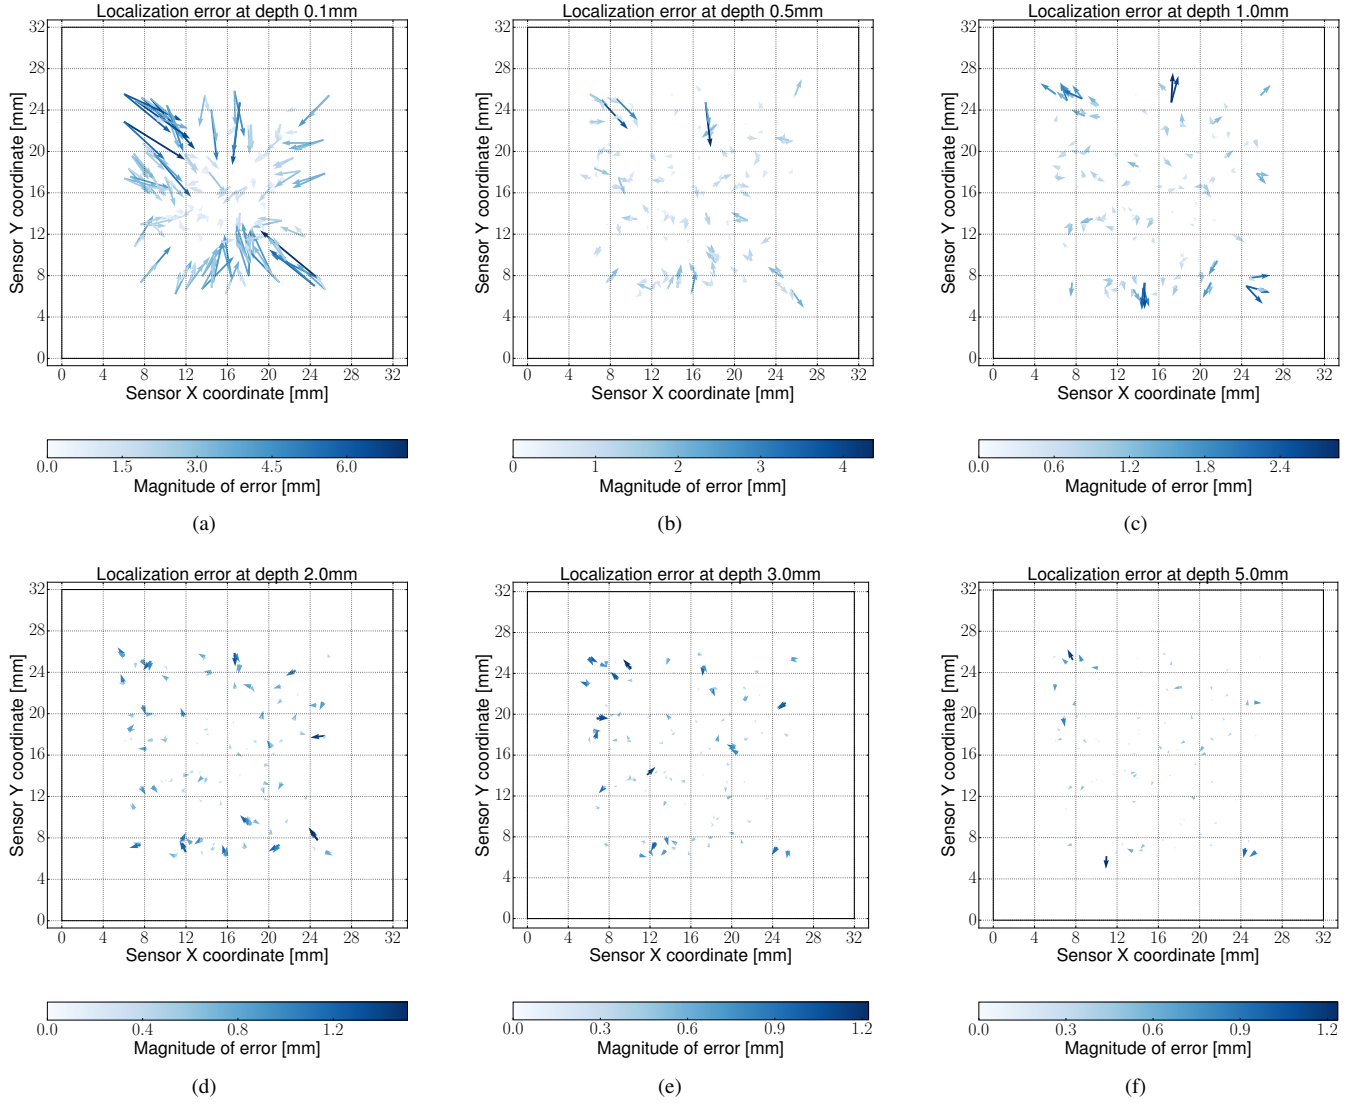


Fig. 7: Localization results for ambient light test dataset. Each arrow represents one indentation in our test set; the base is at the ground truth location while the tip of the arrow shows the predicted location.

TABLE II: Localization and depth accuracy for ambient light dataset

Depth	Localization Accuracy			Depth Accuracy		
	Median Err.	Mean Err.	Std. Dev	Median Err.	Mean Err.	Std. Dev
0.1 mm	2.189	2.369	1.508	0.167	0.167	0.079
0.5 mm	0.717	0.838	0.616	0.050	0.060	0.048
1.0 mm	0.641	0.744	0.530	0.048	0.065	0.051
2.0 mm	0.421	0.500	0.316	0.040	0.048	0.037
3.0 mm	0.362	0.413	0.256	0.037	0.047	0.041
5.0 mm	0.314	0.359	0.222	0.081	0.092	0.062

TABLE III: Localization and depth accuracy for dark dataset

Depth	Localization Accuracy			Depth Accuracy		
	Median Err.	Mean Err.	Std. Dev	Median Err.	Mean Err.	Std. Dev
0.1 mm	3.502	3.832	2.202	0.182	0.183	0.072
0.5 mm	1.102	1.333	0.985	0.050	0.058	0.044
1.0 mm	0.779	0.840	0.535	0.056	0.064	0.047
2.0 mm	0.547	0.657	0.452	0.047	0.056	0.046
3.0 mm	0.410	0.487	0.322	0.039	0.050	0.042
5.0 mm	0.316	0.392	0.288	0.089	0.109	0.072