

Acoustic Sensing for Soft Pneumatic Actuators

Gabriel Zöller Vincent Wall Oliver Brock

Abstract—We propose a novel sensing method for soft pneumatic actuators. The method uses a single microphone, embedded into the actuator’s air chamber. Contact with the environment induces sound (vibration) in the actuator. The materials and the shape of the actuator reflect, refract, and attenuate the sound as it propagates inside the actuator. This produces a unique sound signature for different types of events, enabling the sensing of contact locations, contact force, and the type of contacted material. Sensing is insensitive to the inflation state of the actuator and to background noise. We demonstrate the robustness and versatility of the microphone-based sensor solution in experiments with a PneuFlex actuator. The proposed sensorization avoids the fundamental challenges of sensorizing soft pneumatic actuators, because the placement of a microphone does not negatively affect the compliance of the actuator and because a single microphone suffices for sensorization of the entire actuator, eliminating the need for an application-specific sensor layout.

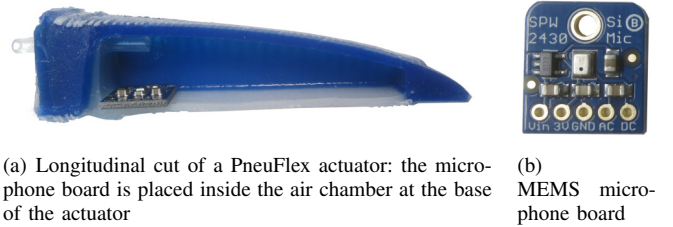
I. INTRODUCTION

Soft pneumatic actuators offer significant benefits for robotic grasping and manipulation, due to their compliance [1], [2], [3]. But their sensorization poses substantial challenges. First, most traditional, “hard” sensors, electronics, and wiring cannot not be built into soft materials without negatively affecting their compliance. Second, the complete recovery of the actuator’s high-dimensional state would require many sensors. Third, the most appropriate sensor placement depends on the application. Once the sensor layout has been built, the actuator is specialized to this application.

We demonstrate sensing of contact location, contact force, and material of the contacted object for a soft pneumatic actuator. The method uses a single microphone, embedded into the actuator’s air chamber, as shown in Figure 1. Our approach exploits acoustics: Airborne and structure-borne sound [4] is affected by the transport medium through reflection, refraction, and attenuation. These interactions between sound and matter are sufficiently complicated to produce a unique sound signature for different types of events at different locations on the actuator. We therefore propose to use a single microphone to infer from sound *what* happened and *where* it happened on the actuator.

Our microphone-based sensing approach eliminates all three of the aforementioned problems of sensorizing soft pneumatic actuators. Even though the microphone itself could be considered “hard”, we place it in the air chamber at the base of the actuator where it does not negatively affect the compliance of the actuator and cabling is trivial. And as

All authors are with the Robotics and Biology Laboratory at Technische Universität Berlin, Germany. We gratefully acknowledge financial support by the European Commission (SOMA, H2020-ICT-645599). We also thank Hugo Wöhrle for his contributions to the project.



(a) Longitudinal cut of a PneuFlex actuator: the microphone board is placed inside the air chamber at the base of the actuator (b) MEMS microphone board

Fig. 1. A microphone embedded into the air chamber of a soft pneumatic actuator enables the reliable detection of localized contact events.

a single microphone is sufficient to recover contact events all over the actuator, we do not need multiple sensors or an application-specific sensor layout.

In our experiments, we demonstrate the surprising accuracy and versatility of the proposed sensing method. We are able to reliably recover contact events and their locations, as well as contact force and material properties on a PneuFlex actuator [1]. Because such contact events play a critical role in human grasping and manipulation [5], we believe that microphone-based sensing is well-suited to advance robotic grasping and manipulation with soft pneumatic actuators.

II. RELATED WORK

For a detailed review of tactile sensing in robot hands, we refer the reader to reference [6]. Here, we discuss related work specifically suited for soft actuators and prior uses of vibration-sensing in robotics.

A. Sensor Technologies for Soft Actuators

The compliance of soft actuators requires that sensors are flexible and stretchable, including electronics and cabling. This renders many existing sensor technologies unusable [6].

Metal alloys, like eutectic gallium-indium (EGaIn), are liquid at room temperature and can be used to detect strain through change in resistance. Based on these alloys, various silicone-based contact and strain sensors have been proposed [7], [8]. Bilodeau et al. [9] include liquid metal strain sensors into a four-fingered gripper. Farrow and Correll [10] inject EGaIn into prefabricated silicon tubes to produce an easily customizable liquid metal strain sensor. Given a soft actuator and a desired application, it is possible to create economic and effective sensor layouts [11].

Conductive, fabric-based sensors can serve as an alternative to resistive, liquid metal-based sensors. For example, knitted fabrics with embedded silver nanowires produce a flexible capacitive sensor when sandwiched with a stretchable dielectric layer [12]. A deformable skin measures surface contact locations using the principle of electrical

impedance tomography [13]. Another sensor technology relies on stretchable optical waveguides and light attenuation for strain sensing in fluidic actuators [14].

Flexible, stretchable tactile arrays constructed from MEMS barometers measure pressure inside soft material [15]. However, the electronic components are only flexible but not stretchable, limiting their applicability.

Other sensor technologies, such as the GelSight [16] or TacTip [17], include soft-material components, but still require “hard” sensor components near the location of contact.

With all of these sensor technologies, it is necessary to place the sensor at the location where something must be sensed. Acoustic sensing holds the surprising promise of removing this constraint.

B. Acoustic Sensing

Microphones or vibration-sensing accelerometers have been used in robotics in a variety of ways: to recognize objects from the sound they make upon contact [18] or more complex interactions [19], to detect materials [20] and textures [21], [22], to identify the contents of a sealed container by shaking [23], or to detect the attainment of contact during grasping and placement of objects [24]. All of these applications measure properties of the world, rather than properties of the robot itself.

A microphone-based proximity sensor leverages the seashell effect [25], which amplifies ambient sound based on the resonance frequency of the air pocket between the microphone and the nearby object. Here, the microphone is used to measure the distance between a gripper and an object.

None of the uses of microphones we discussed so far take advantage of sound modulation that occurs in structure-borne sound as a result of the sound-carrying structure [4]. One work that makes a step into this direction uses forearm-mounted contact (piezo) microphones to detect sound (transmitted through the bones) from finger gestures, including tapping, flicking, writing, etc. [26]. Here, the structure is used to conduct sound. The only work that deliberately exploits modulation in structure-borne sound uses this effect to create a touch-based input device, buttons or touchpad, by placing an active vibration source and a vibration sensor on a hollow box [27]. The touch of a finger modulates the vibration, enabling detection of the finger’s contact location on the box.

The work presented in this paper will take advantage of sound modulation in the sound-conducting structure to “localize” sound-causing contact events. In contrast to the aforementioned work [27], our sensing is passive and therefore detects contact events, rather than contact states.

III. INSTALLING THE MICROPHONE

The pneumatically powered PneuFlex actuator, shown in Figure 1, was developed as a finger in the context of robotic grasping with soft hands [1]. During inflation, the expansion of the actuator’s silicone hull (Smooth-On Dragon Skin 10) is constrained in two ways: first, by a helically wound polyester thread to prevent radial expansion, and second, by an inextensible but flexible passive layer, embedded in

the bottom of the finger. This layer prevents longitudinal extension of the bottom of the finger, leading to a bending motion.

We embed a MEMS (micro-electro-mechanical system) condenser microphone (Adafruit SPW2430, see Figure 1) into the air chamber of the actuator, as close as possible to the base of the actuator. This part experiences the least deformation when the actuator is inflated. To install the microphone, we cut three sides of the finger, leaving the bottom

layer and the polyester thread intact. Through the opening we insert the microphone and attach it with silicone adhesive to the bottom of the finger, facing up (see Figure 2). The finger is then resealed with silicone. The cables leading to the microphone are routed through the cut and attached to the outside of the finger, also with silicone glue. This is a simple procedure that does not affect the structural integrity or expansion properties of the actuator.

IV. CLASSIFYING ACOUSTIC SIGNALS

A. Data Acquisition

Our first goal is to demonstrate the identification of contact locations based on the sound recorded by the microphone. We will employ machine learning to classify the sound. As a first step, we record sound data.

We recorded the sound that resulted from tapping the finger against an acrylic strut, as shown in Figure 3, via a USB audio interface recording at a 24bit resolution and a 48 kHz sampling rate. The microphone is rated for a range of 100 Hz to 10 kHz, in which it behaves linearly—but it records also much lower frequencies (see Figure 4).

For each of the eight contact locations (see Figure 3), we recorded five repetitions at four inflation levels of the actuator: 0 kPa, 15 kPa, 30 kPa, and 50 kPa, the latter corresponding to full inflation. Six such sets were recorded by two different experimenters. This amounts to a total of 960 sounds arising from contact events. The beginning and end of the recording for each contact event were selected manually. Each of the resulting sound recordings is labeled with the contact location to obtain our first data set.

We also want to demonstrate that sound-based sensing is able to identify contact force and the material the tap was performed on. For this, we recorded a second data set. Now the tip of the finger is the only contact location we consider. We tap the fingertip against objects made of three different materials: wood, silicone, and aluminum. We used three different tapping durations: lifting the finger off the object after 0.1 s, 0.5 s, and 1 s. For each of these, we recorded three different ranges of contact forces, measured by a force sensor

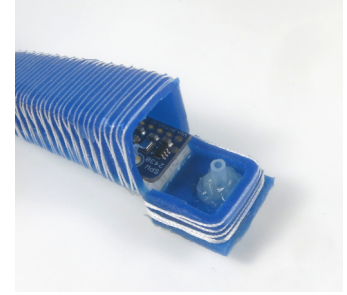


Fig. 2. To install the microphone, the actuator is cut at the base, leaving structural reinforcements intact. Then the microphone is inserted and the finger is glued shut again.

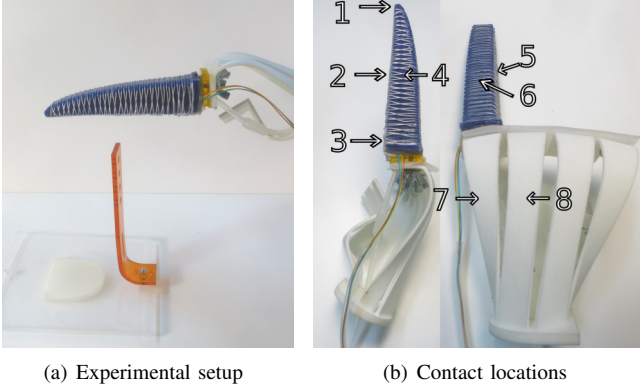


Fig. 3. Contact events were created by tapping different parts of the finger and hand scaffold against an acrylic strip (a). We used eight different locations (b), three on the bottom of the finger (1: tip, 2: middle, 3: base), three on the other sides of the finger (4: left, 5: right, 6: top), and two on the hand’s scaffold (7: mounting strut of the finger, 8: other parts of the structure). The scaffold is from a four-fingered RBO Hand 2 [1]

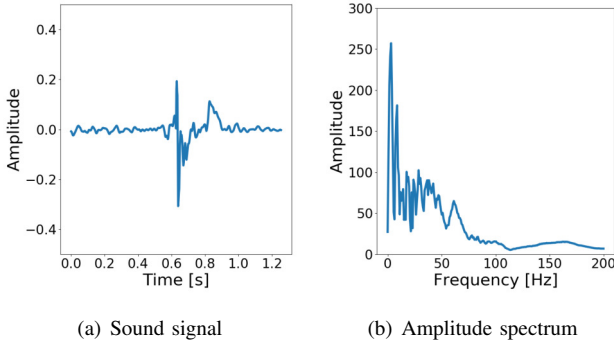


Fig. 4. Sound recorded by the microphone during a contact event (left) and the corresponding amplitude spectrum (right): the actuator is inflated to 15 kPa and contact is made at the tip

mounted underneath the object: 0 N to 1 N, 1 N to 2 N, and 2 N to 3 N. All of this was done for the same four inflation states of the actuator as above. We recorded two sets of these measurements, amounting to a total of 648 recorded sounds, making up our second data set. Each of the resulting sound recordings is labeled with the material, the tapping duration, and the contact force.

B. Input Features

To ready our data set for machine learning, we extract a consistent feature vector from each sound file. We apply a discrete Fourier transformation to obtain an amplitude spectrum, which is the real-valued absolute of the complex-valued frequency spectrum. This conversion turns variable-length sound files into feature vectors of identical length, each value corresponding to the amplitude of a specific frequency range. Figure 4 shows a sound file and the corresponding amplitude spectrum. We include the *cutoff frequency* and *sampling resolution* of the amplitude spectra as hyperparameters to the classification problem (see next section). Finally, we normalize the feature vectors to zero mean and unit variance on the training set.

C. Classification of Amplitude Spectra

We want to learn a classifier that enables the identification of contact locations from a feature vector (processed amplitude spectrum). We tested three different classification algorithms: k-nearest neighbor, logistic regression, and support vector machines. We did not attempt neural networks, as we have relatively few training data compared to the high dimensionality of our feature vector.

Our classifier is based on the scikit-learn library [28] which implements all three algorithms. We performed an exhaustive parameter search (grid search) over hyperparameters. For k-nearest neighbors, we tested different distance metrics (Manhattan and Euclidean distance). For support vector machines we tested different kernels (linear, radial basis function, and the parameters γ and C). For logistic regression, we tested different regularizations (L_1 and L_2 norm). For all of the classifiers, we tested different frequency cutoffs (100 Hz, 200 Hz, 1 kHz, 3 kHz, 10 kHz) and different resolutions for sampling the amplitude spectrum (100 samples, 1000 samples).

To identify optimal hyperparameters for each classifier, we used 5-fold cross-validation. The classifiers were trained on two sets of recordings and evaluated on another set of recordings from a different date. The best classifier, by a small margin (ca. 2%) over k-nearest neighbors and a large margin (ca. 20%) over the other two algorithms, was a support vector machine (SVM) with a radial basis function kernel ($\gamma = 0.001$ and $C = 100$) on a spectrum with a cut-off frequency at 200 Hz at 1000 data points.

V. EXPERIMENTAL VALIDATION

A. Sensing Contact Locations

We trained the SVM classifier described above on the first data set and tested with a held-out set of sounds. Classification rates are summarized in the confusion matrix shown in Figure 5. The overall classification rate is 74.7 % for all contact points, including the scaffold. In the following, we provide some interpretation and explanation of the data shown in the confusion matrix.

The contact locations *left*, *right*, *top*, and *middle* cover the four sides of the actuator at the identical distance from the base. As a result, we might expect significant misclassification. And we do see this to some degree between opposing contact points: Left is misclassified as right in 25 % of the cases and middle as top in 10 % of the cases. Interestingly, however, the inverse confusion is much less likely: Right is misclassified as left only in 13 %, and top as middle only in 5 % of the cases. We interpret this to imply that the misclassifications of left as right and middle as top can also be learned better, possibly with more data or with slight design modifications.

The *structure* contact point, on the scaffold without a finger, exhibits a decent classification rate of 70 %. This suggests that the classification results remain acceptable, even when the distance the sound travels through material (and consequently the signal attenuation) is large.

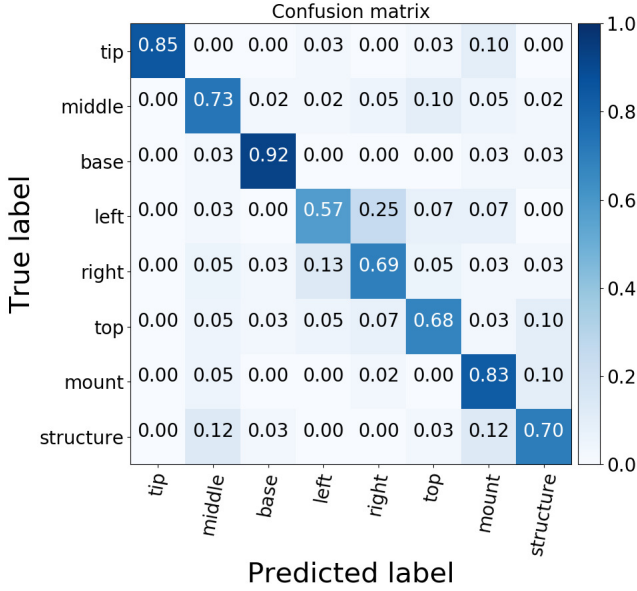


Fig. 5. Confusion matrix for the SVM-based prediction of contact location: the high values along the diagonal indicate that sound recorded from a single microphone embedded into the actuator can be used to identify contact locations.

Contact at the *tip* of the actuator is misclassified as contact on the *mounting* strut 10 % of the time. These two contact locations lie on at opposite ends of a line through the microphone. It is possible that the actuator might vibrate in similar ways when tapped at the extreme points of its length. This, too, may provide some guidelines for how to optimize actuator shape for microphone-based sensing.

Contact at the *base* is classified correctly 92 % of the time. Tapping at the base of the finger, right where the microphone is mounted, should be louder than other locations. This might provide a more significant feature to learn on than the vibration of the whole actuator.

Overall, these results demonstrate the feasibility of acoustic sensing of contact locations in soft-material actuators. We found it surprising that a single microphone and simple machine learning methods suffice to detect contact locations with 74.7 % accuracy. While these results do not (yet) reach the reliability of contact sensors, we are optimistic that this very first proof of concept can be improved upon in many ways. And in contrast to tactile sensors, the proposed microphone-based approach is cost-effective (the microphone costs \$5), simple to install, and unintrusive to the desired material behavior of the soft actuator.

B. Measuring Robustness to Background Noise

One might expect environmental noise to affect the sound-based identification of contact location. To test this, we repeat the previous experiments while playing back recordings of common office noises at different volumes (50 dB, 70 dB and 90 dB) from a distance of 10 cm. In Figure 6 we compare the overall classification rates for all eight contact locations at different noise levels. The first column corresponds to the original data, recorded in a quiet room at 30 dB.

The plot shows that the mean classification rate does not correlate with the volume of the background noise. The standard error of the classification rate is slightly higher in the presence of background noise, but this also does not correlate with the volume. This suggests that the proposed sound-based method for detecting contact locations is robust against background noise, even when that noise is very loud. This result can be explained by the sound insulation provided by the hull of the actuator. We performed the same test with white noise and achieved comparable results.

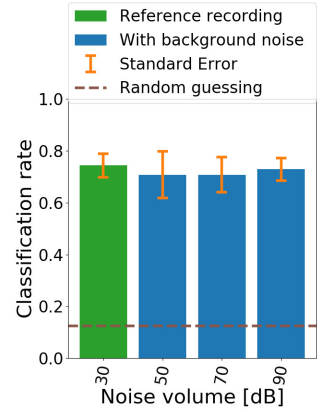


Fig. 6. The classification rate of the proposed sensor system is practically unaffected by background noise, irrespective of the noise level. The dashed line indicates the performance of guessing.

C. Sensing Contact Force

We now examine if it is also possible to predict contact forces using acoustic sensing. For this, we use the second data set, in which sounds are labeled by the three ranges of contact force. In addition, this data set includes sounds from three different object materials and three tapping durations. We again perform a hyperparameter search over the three classifiers. This search produces the same result as before: the best classifier is an SVM with a radial basis function kernel and the same hyperparameters as above. Training uses two-thirds of the recorded data, with one-third serving as test set. Please note that the recorded data is not ideal for the classification of contact forces, as the three force ranges are not clearly separated from each other but instead represent three adjacent force ranges.

Over a random-guessing baseline of 33 %, the best SVM classifier achieves a cross-validation score of 80 % and a test score of 78 %. These results demonstrate generalization of the SVM over tap duration, materials, and actuator inflation level, as these parameters vary in the training and test data.

The confusion matrix for the classification results is shown in Figure 7. High contact forces (2 N to 3 N) can be classified with 96 % accuracy. For lower contact forces (<2 N), the classification rate is lower. This might be explained by low-amplitude sound being less distinctive, even if it represents different events.

Our results demonstrate that it is possible to measure contact location *and* contact force using a single microphone.

D. Detecting Materials

To challenge the versatility of acoustic sensing even further, we attempt to predict a third type of sensor signal from the recorded sound: the material that was contacted during tapping. We again use the second data set, now attempting to learn a classifier to predict the material. We perform a

third grid search over classifiers and hyperparameters, again as described above. The best classifier is for the third time an SVM with the same hyperparameters. We train all classifiers on two-thirds of the second data set and test the trained classifiers on the remaining third.

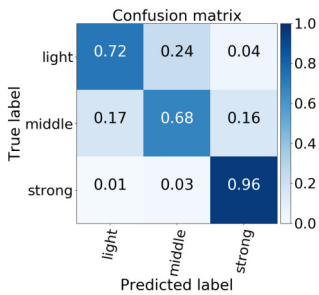


Fig. 7. Confusion matrix for the SVM-based prediction of contact forces: the high values on the diagonal indicate that contact forces can be inferred from sound recorded by a microphone embedded into the actuator. The labels light, middle, and strong correspond to the three force ranges described in the text.

Our results demonstrate that a microphone embedded in a soft-material actuator can be used to sense a variety of information types: contact location, contact forces, and material type. To achieve the same with “traditional” sensing technologies, we would likely require several different sensor technologies and multiple physical sensors.

E. Applying Acoustic Sensing

To demonstrate the proposed acoustic sensing in an application, we place the RBO Hand 2 on a door handle in preparation for grasping it. The index finger of this hand is equipped with our acoustic sensor, and the contact location is predicted online from the sound signal. In future work, a controller could use this information to correct misalignments of the hand in increase grasp success.

A sliding-window predictor interprets the sound of the previous second at an average rate of 13 Hz. To also detect when *no* contact occurs, we introduce a ninth class called “none”, in addition to our previous eight contact locations (see Figure 3). To capture the different acoustic properties of the RBO Hand 2, we recorded a new data set consisting of 900 samples (5 sets, 4 inflation levels, 9 classes, 5 repetitions).

Figure 8(a) shows the experimental setting. Three still frames (b-d) show the moments the hand touches the handle with the index finger’s tip, middle, and base. The plots (e) visualize the recorded sound and the predicted contact class probabilities.

As expected, the “none”-class is most probable when the hand is not making contact. The “tip”-class is assigned high probabilities twice: When the sliding prediction window enters and when it leaves the event signal. In between, no clear prediction is made. The “middle”-class is correctly predicted with high probability, very quickly after contact

occurs. After the “base”-contact, the first few predictions are uncertain, before the correct class is predicted. This shows that different phases of the contact event, e.g. the initial contact or the following propagation, have different significance for the classes. When the relevant sound of the event is not within the sliding window, the prediction may be incorrect. Thus, the choice of the prediction window size influences the results.

This proof-of-concept experiment shows that our approach can interpret acoustic signals during hand/environment interactions in an application setting with a full RBO Hand 2 setup.

VI. CONCLUSION

We proposed a microphone-based approach to the sensorization of soft pneumatic actuators. We placed a microphone inside the air chamber of the actuator. Sound (vibration) caused by contact events anywhere on the actuator propagates through the actuator’s components and materials, thereby producing sound signatures that contain information about the event to be sensed. Using standard machine learning methods, we were able to sense contact location, contact force, and the material of the contacted object—using a single microphone. In some way, this microphone can replace multiple sensor modalities and multiple sensors of a single modality. We successfully validated the proposed acoustic sensing in experiments with a specific type of soft actuator, the PneuFlex actuator, but are convinced that acoustic sensing is applicable to any soft-material, fluidically powered actuator type.

The proposed method overcomes fundamental challenges of sensorization in soft-material robotics. Sensors do not have to be embedded into soft materials, which may adversely effect desirable material properties and complicate the design. Instead, we can place a microphone at any convenient location. In our approach, the number and placement of sensors do not have to be chosen in a task-specific manner. Instead, a single sensor detects contact events and locations anywhere on the actuator, as long as the contact event produces sound.

This is a first proof of concept and the proposed method still has limitations. Designing actuators for helpful sound propagation and actuator abilities at the same time is more difficult than designing for one of the criteria alone. Also, with the current sensorization method, it is only possible to detect sound-causing events—but not steady-state contact, contact forces, or contact normals. However, we regard the present paper as a first step. We believe that a better understanding of how to design discriminative sound propagation into soft-material robots, how to exploit multiple microphones, and how to employ active sound sensing will turn this proof-of-concept demonstration into a powerful sensorization method for soft robotics.

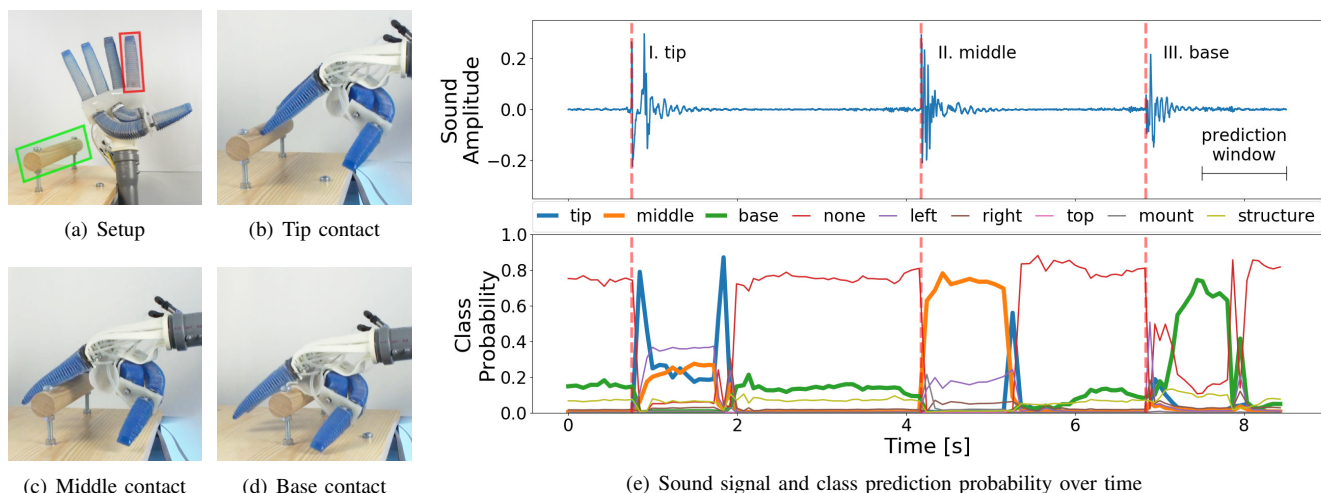


Fig. 8. Setup and results of the application experiment: (a) An RBO Hand 2 is equipped with an acoustic sensor in its index finger (red) and trained to sense contact locations with a wooden door handle (green). Over the course of the experiment, the hand makes contact with the door handle three times: At the tip (b), middle (c) and base (d) of the index finger. The contact location is predicted on the sound signal in real-time using a moving window of 1 s. The onset of the three events is marked with a vertical red line in the sound signal (e, top) and the class probabilities (e, bottom). Directly afterwards, the corresponding classes are predicted with varying success. Albeit preliminary, this indicates the contact location can be detected online on real-world hardware. A video of the experiment can be found here: <https://youtu.be/92KaQClofyc>

REFERENCES

- [1] R. Deimel and O. Brock, "A novel type of compliant and underactuated robotic hand for dexterous grasping," *The International Journal of Robotics Research*, vol. 35, no. 1-3, pp. 161–185, 2016.
- [2] C. Eppner, R. Deimel, J. Álvarez-Ruiz, M. Maertens, and O. Brock, "Exploitation of environmental constraints in human and robotic grasping," *International Journal of Robotics Research*, vol. 34, no. 7, pp. 1021–1038, 2015.
- [3] C. Eppner and O. Brock, "Planning grasp strategies that exploit environmental constraints," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 4947–4952.
- [4] L. Cremer, M. Heckl, and B. A. T. Petersson, *Structure-Borne Sound: Structural Vibrations and Sound Radiation at Audio Frequencies*, 3rd ed. Springer-Verlag, 2005.
- [5] R. S. Johansson and J. R. Flanagan, "Coding and use of tactile signals from the fingertips in object manipulation tasks," *Nature Reviews Neuroscience*, vol. 10, no. 5, pp. 345–359, 2009.
- [6] Z. Kappassov, J.-A. Corrales, and V. Perdureau, "Tactile sensing in dexterous robot hands — Review," *Robotics and Autonomous Systems*, vol. 74, no. A, pp. 195–220, 2015.
- [7] D. M. Vogt, Y. L. Park, and R. J. Wood, "Design and characterization of a soft multi-axis force sensor using embedded microfluidic channels," *IEEE Sensors Journal*, vol. 13, no. 10, pp. 4056–4064, 2013.
- [8] J. W. Boley, E. L. White, G. T.-C. Chiu, and R. K. Kramer, "Direct writing of Gallium-Indium alloy for stretchable electronics," *Advanced Functional Materials*, vol. 24, no. 23, pp. 3501–3507, 2014.
- [9] R. A. Bilodeau, E. L. White, and R. K. Kramer, "Monolithic fabrication of sensors and actuators in a soft robotic gripper," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2015, pp. 2324–2329.
- [10] N. Farrow and N. Correll, "A soft pneumatic actuator that can sense grasp and touch," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2015, pp. 2317–2323.
- [11] V. Wall, G. Zöllner, and O. Brock, "A method for sensorizing soft actuators and its application to the RBO Hand 2," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 4965–4970.
- [12] A. Atalay, V. Sanchez, and O. Atalay, "Batch fabrication of customizable silicone-textile composite capacitive strain sensors for human motion tracking," *Advanced Materials Technologies*, vol. 2, no. 9, p. 1700136, 2017.
- [13] F. Visentin, P. Fiorini, and K. Suzuki, "A deformable smart skin for continuous sensing based on electrical impedance tomography," *Sensors*, vol. 16, no. 11, p. 1928, Nov. 2016.
- [14] H. Zhao, K. O'Brien, S. Li, and R. F. Shepherd, "Optoelectronically innervated soft prosthetic hand via stretchable optical waveguides," *Science Robotics*, vol. 1, no. 1, p. eaai7529, 2016.
- [15] L. P. Jentoft, Y. Tenzer, D. Vogt, J. Liu, R. J. Wood, and R. D. Howe, "Flexible, stretchable tactile arrays from MEMS barometers," in *International Conference on Advanced Robotics*, 2013.
- [16] W. Yuan, S. Dong, and E. H. Adelson, "GelSight: High-resolution robot tactile sensors for estimating geometry and force," *Sensors*, vol. 17, no. 12, p. 2762, 2017.
- [17] N. F. Lepora, K. Aquilina, and L. Cramphorn, "Exploratory tactile servoing with biomimetic active touch," *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 1156–1163, 2017.
- [18] E. Torres-Jara, L. Natale, and P. Fitzpatrick, "Tapping into touch," in *Proceedings of the 5th International Workshop on Epigenetic Robotics*, 2005, pp. 78–86.
- [19] J. Sinapov, T. Bergquist, C. Schenck, U. Ohiri, S. Griffith, and A. Stoytchev, "Interactive object recognition using proprioceptive and auditory feedback," *The International Journal of Robotics Research*, vol. 30, no. 10, pp. 1250–1262, 2011.
- [20] E. Krotkov, R. Klatzky, and N. Zumel, "Robotic perception of material: Experiments with shape-invariant acoustic measures of material type," in *Experimental Robotics IV*. Springer, 1997, pp. 204–211.
- [21] W. W. Mayol-Cuevas, J. Juarez-Guerrero, and S. Munoz-Gutierrez, "A first approach to tactile texture recognition," in *1998 IEEE International Conference on Systems, Man, and Cybernetics*, 1998, vol. 5, Oct. 1998, pp. 4246–4250.
- [22] J. M. Romano and K. J. Kuchenbecker, "Creating realistic virtual textures from contact acceleration data," *IEEE Transactions on Haptics*, vol. 5, no. 2, pp. 109–119, 2012.
- [23] C. L. Chen, J. O. Snyder, and P. J. Ramadge, "Learning to identify container contents through tactile vibration signatures," in *2016 IEEE International Conference on Simulation, Modeling, and Programming for Autonomous Robots (SIMPAP)*, Dec. 2016, pp. 43–48.
- [24] J. M. Romano, K. Hsiao, G. Niemeyer, S. Chitta, and K. J. Kuchenbecker, "Human-inspired robotic grasp control with tactile sensing," *IEEE Transactions on Robotics*, vol. 27, no. 6, pp. 1067–1079, 2011.
- [25] L.-T. Jiang and J. R. Smith, "Seashell effect pretouch sensing for robotic grasping," in *IEEE International Conference on Robotics and Automation*, 2012, pp. 2851–2858.
- [26] B. Amento, W. Hill, and L. Terveen, "The sound of one hand: A wrist-mounted bio-acoustic fingertip gesture interface," in *CHI '02 Extended Abstracts on Human Factors in Computing Systems*, 2002, pp. 724–725.
- [27] J. Liu, C. Wang, Y. Chen, and N. Saxena, "VibWrite: Towards finger-input authentication on ubiquitous surfaces via physical vibration," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, 2017, pp. 73–87.
- [28] F. Pedregosa et al., "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.