Active Acoustic Contact Sensing for Soft Pneumatic Actuators

Gabriel Zöller   Vincent Wall   Oliver Brock

***Abstract*— We present an active acoustic sensor that turns soft pneumatic actuators into contact sensors. The whole surface of the actuator becomes a sensor, rendering the question of where best to place a contact sensor unnecessary. At the same time, the compliance of the soft actuator remains unaffected. A small, embedded speaker emits a frequency sweep which travels through the actuator before it is recorded with an embedded microphone. The specific contact state of the actuator affects how the sound is modulated while traversing the structure. We learn to recognize these changes in the sound and map them to the corresponding contact locations. We demonstrate the method on the PneuFlex actuator. The active acoustic sensor achieves a classification rate of 93% and mean regression error of 3.7 mm. It is robust against background noises and different objects. Finally, we test it on a Panda robot arm and show that it is unaffected by motor noises and other active sensors.**

I. Introduction

We use sound to sense contact anywhere on soft robotic actuators. In effect, our active acoustic contact sensor converts the whole actuator into one big contact sensor. It offers detailed measurements of the actuator's contact state while being low-cost and non-intrusive. This is important because sensors for soft actuators should not limit the actuator's compliance. The sensor technology needs to be as compliant as the actuator itself. Finding a good placement of a sensor is difficult when one needs to trade-off access to the relevant measurements with the negative effects on the compliance. Additionally, to maintain the low fabrication costs of soft actuators, soft sensors should be equally low-cost.

In previous work, we have shown that dynamic contacts create impact sounds that we can measure to locate contact points [1]. The sound transports information about the contact from the point of origin, through the actuator, to a microphone. This allows any part of the actuator to act as a contact sensor. Additionally, the actual physical sensor, i.e. the microphone, can be placed somewhere to the side, where it does not affect the actuator's compliance. However, the *dynamic* contact is only measurable until the impact sound has faded. Afterwards, when the contact is *static*, no new sound is introduced. It is difficult to acoustically sense contact in this passive case. Instead, we propose to actively emit a custom sound inside the actuator using a small embedded speaker. This way, the sensor is independent of externally created impact sounds, and can instead sense the static contact state at any time. When the active sound travels through the structure of the actuator, it is modified slightly based on the acoustic properties of the actuator [2]. These acoustic properties depend on the design and material of the actuator, but also on its current contact state. In a way, each contact state "sounds" differently. While these interactions are too complex to model analytically [3], the contact-dependent differences in sound are detectable by data-driven machine learning methods.

We demonstrate the active acoustic contact sensor on the example of the highly compliant PneuFlex actuator [4]. We embed a speaker and a microphone at different ends of the actuator's air chamber (Fig. 1), and learn to localize static contact on the actuator's hull. Using the active sound, we distinguish six contact locations with a classification rate of 93% and predict the exact location with a mean regression error of 3.7 mm. This is a significant improvement over passive and dynamic acoustic sensing. Furthermore, we show that the sensor is unaffected by different objects

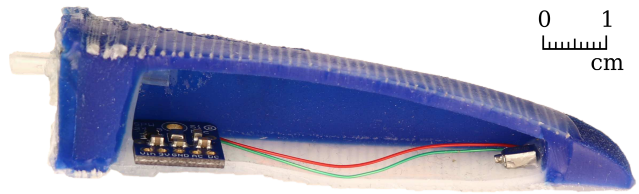

(a) Inside view of a PneuFlex actuator: Microphone (left) and speaker (right) are attached to the passive bottom layer at opposite ends of the air chamber.

(b) The speaker (top) and MEMS microphone (bottom) are small enough to fit into the actuator.

(c) Experimental setup on a Panda robot: Four sensorized actuators are mounted on an RBO Hand 2. Contact is measured with a static object in different workspace locations.

Fig. 1. For active acoustic contact sensing, we embed a microphone and a speaker into a soft pneumatic actuator. The emitted frequency sweep is affected by the specific contact state. A trained sensor model detects these subtle changes and predicts the contact location anywhere on the actuator.

We thank Marius Hebecker for his help with the experiments.
All authors are with the Robotics and Biology Laboratory, Technische Universität Berlin, Germany. We gratefully acknowledge financial support by the European Commission (SOMA, H2020-ICT-645599), the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2002/1 "Science of Intelligence" - project number 390523135 and German Priority Program DFG-SPP 2100 "Soft Material Robotic Systems".

and loud background noises. Finally, we test the sensor on a Panda robot arm and demonstrate its robustness against motor noises and neighboring active acoustic sensors.

The presented active acoustic contact sensor turns the whole pneumatic actuator into a contact sensor without affecting its compliance. The sensor measures static contact by transmitting an active sound through the actuator, where it is changed depending on the actuator's contact state. Because this fundamental effect is present in all soft actuators, we believe that the active acoustic contact sensor is applicable to soft actuators in general.

II. Related Work

The challenge in soft robotic sensorization is to find approaches that are applicable to the highly compliant structures, while measuring the relevant information. While well-established approaches exist to sense contact, forces, and more for traditional, "hard" robots [5], [6], soft roboticists are still evaluating different approaches. We will first give a brief overview of soft sensing in general, and then survey in more detail the research field of acoustic sensing.

A. Soft Sensing

Existing soft sensing approaches can be organized by their intrusiveness, i.e. how much do the sensors affect compliance, and their level of detail, i.e. how precisely do the sensors measure the actuator's state. The least intrusive methods use external vision systems that track the position and shape of the actuator [7]. While this does not influence the actuator at all, these methods are limited to pose measurements and fail quickly in case of occlusions, which are common during manipulation tasks. Soft tactile arrays, on the other hand, are placed directly on the actuator [8]. This allows precise measurements even during occlusions, but only at the specific sensor location and often at the cost of some compliance. Between these two extremes, liquid metal strain sensors [9], [10] and optical waveguides [11], [12] aggregate measurements over a larger area. This limits the level of detail, as the exact origin of a measurement is unclear. But needing fewer sensors also reduces the influence on compliance. To further reduce the impact on compliance, the sensor itself may be placed somewhere it does not affect the actuator, while using the actuator itself to transmit information from the point of measurement to the recording device. TacTip [13] and GelSight [14] are examples of visual sensors that use light to transmit contact information. The benefit of this method is that the actuator remains unaffected, while detailed measurements are possible from anywhere on the actuator. The same method is used for acoustic sensing.

B. Acoustic Sensing

In acoustic sensing, sound is used as an information carrier for various measurements. For example, sound is used for distance measurements in mobile robots [15] and in robotic grippers [16]. While these approaches are limited to the pre-contact phase, they show that sound carries useful information. To obtain information about the contact itself, several approaches analyze the characteristics of sounds created by the contact and use these to recognize and classify objects [17], [18], [19]. In robotic manipulation, similar techniques have been used to explore objects interactively and with multi-modal measurements. Contact sound is created by tapping objects with a compliant hand [20] and combined with vision, haptic, and/or joint torque measurements to categorize objects [21], [22], [23]. These works have studied the characteristics of contact sounds and identified effective sound features and their representations, which we will also use in this work. However, these approaches all use impact sounds created by a *dynamic contact*. To measure the *static contact state*, after the impact sound has faded, we need to *actively* introduce a new sound.

In active acoustic sensing, sound is artificially added to the system to induce a measurable acoustic signal. Ono et al. have used this technique to sensorize rigid objects [24], but stated that it would be difficult to apply their method to soft materials. Mujibiya et al. proposed a human-worn sensor which measures contact using transdermal ultrasound signals [25], but it requires the receiver to be mounted on the object that makes contact. In soft robotics, active acoustic sensing was recently used by Takaki et al. [26] who used a small speaker and microphone inside a pneumatic bellow actuator to measure its length by modeling the shift in resonance frequencies. In this paper, we use a similar hardware setup. However, we employ a model-free, data-driven approach to extract detailed contact information from the active sound signal. In effect, we turn the complete soft actuator into a contact sensor.

III. Building an Active Acoustic Contact Sensor

The active acoustic contact sensor works by emitting a sound into the actuator and recording it after it traveled through the structure. The contact location is then calculated by detecting small changes in the sound signal, which are specific to each contact state. We describe the fabrication of such a sensor by the example of the pneumatic PneuFlex actuator [4]. However, we attempt to keep the description as general as possible to be easily applicable to other actuators.

1) Hardware Setup: First, we add the audio components to the PneuFlex actuator. It is important to place the components somewhere with minimal influence on the actuator's compliance while maintaining enough distance between speaker and microphone for the sound to be modulated based on the current contact state. We embed a MEMS condenser microphone (Adafruit SPW2430) at the base-end of the air chamber. A small balanced armature speaker (Knowles RAB-32063-000) is placed at the tip-end to maximize the distance (Fig. 1(a)). We attach both devices to the passive bottom-layer of the actuator, because it experiences the least deformation during use. To minimize the influence on the compliance, we route the speaker cables in a curved pattern along the passive layer. Both components are connected to a USB audio interface (MAYA44 USB+) recording at a sample rate of 48 kHz with 32 Bit precision. We chose these audio components for their small form factor, which fits well

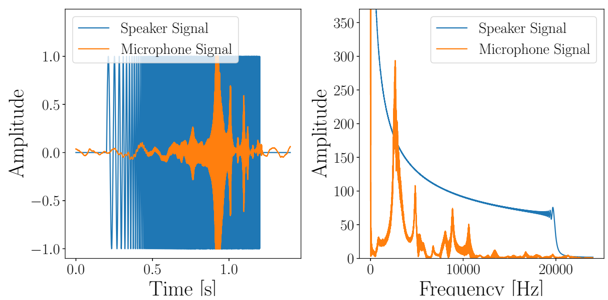

(a) Sound Waveforms (b) Amplitude Spectra

Fig. 2. The generated sweep (blue) emitted by the speaker is quite different from the sound recorded by the microphone (orange). This is due to the limitations of the audio components, but also because of the acoustic properties of the actuator. For example, some of the peaks in the recorded spectrum correspond to the resonance frequencies of the air chamber.

into the actuator's air chamber. But they also offer decent performance in a comparatively large frequency range.

The sensorized PneuFlex actuator is then mounted on an RBO Hand 2 (Fig. 4(a)). We removed the other three fingers for better access to all sides of the sensorized actuator, as we want to evaluate the sensor's ability to measure contact anywhere on the actuator. The sensorized RBO Hand has an attachment for a manual handle for human operators and for a Franka Emika Panda robot arm (Fig. 1(c)).

2) Generation and Recording of the Active Sound: Next, we generate the active sound emitted by the embedded speaker which travels through the actuator, where it is modulated depending on the actuator's contact state. The sound we use is a logarithmic frequency sweep from 20 Hz to 20 kHz with a duration of 1 s (Fig. 2). The sweep contains all frequencies in sequence, while the logarithmic profile emphasizes the lower-frequency sounds. We expect these to be more relevant for the distinction of contact states, because in our previous work [1], low frequencies played an important role for classification performance. Because we have no reliable acoustic models to predict the change of acoustic properties of a soft actuator during interaction [3], we instead use the full-range frequency sweep, which contains all possibly relevant frequencies. The 1 s duration was chosen as to have enough time for even low frequencies to observe the contact-dependent modulation effects. While this limits the achievable sensing rate, we first aim to prove the viability of the sensor. We are confident that the sensing rate can be improved in the future.

The microphone at the base of the actuator records the active sound after it traveled through the structure. The recording is synchronized with the playback and trimmed to 1.4 s samples (1 s of sweep with 0.2 s of silence before and after to compensate for small inaccuracies in the cutting of the sounds). In the frequency domain, the contact-dependent changes to the active sound are easily recognizable. We extract the amplitude spectrum of each sample using a discrete Fourier Transform. The result is a 33k-dimensional feature vector containing the amplitude of each frequency in the signal from 0 Hz to 24 kHz (Fig. 2(b)). In our approach, we disregard the phase spectrum to be independent of the exact timing of sounds. In the future, the phase spectrum might be used to extract even more information from the contact sounds. While our work shows high performance without weighting or reducing frequency bins, this represents another source for potential future optimization.

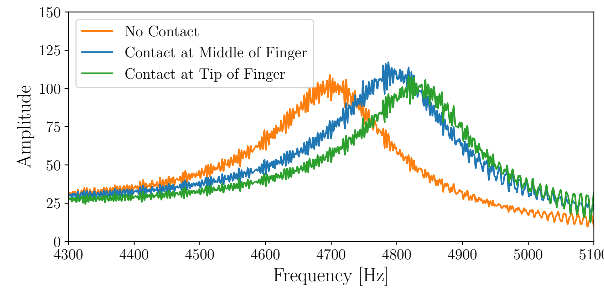

Fig. 3. The amplitude spectra from three different contact states are visibly different. Contact changes the acoustic properties of the actuator and each state has a distinguishable amplitude spectrum.

3) Mapping from Sound to Contact: The last step is the supervised learning of the mapping between the sound and contact state. We record sets of training data that contain example recordings of the active acoustic sensor in various contact states and their corresponding contact locations. We use a k-nearest neighbor (KNN) predictor from the scikit-learn library [27] because it is simple and does not need a lot of data. The default implementation uses five neighbors and the Euclidean distance metric. We did not need to change the parameters, because even with these default values the prediction was very good. This indicates that the actual learning problem is rather simple, which can also be seen by manually inspecting the amplitude spectra of different contact locations: Figure 3 shows a section of three spectra recorded in different contact states, using the same active sound. The different contact locations create a noticeable frequency shift. To predict the contact location for a new sound, the sensor model simply identifies the most similar, existing spectra. With these steps, we have sensorized the PneuFlex actuator for active acoustic contact sensing and have set up the necessary processing steps for the measurement of contact states from sound.

IV. Experimental Validation

To evaluate the active acoustic contact sensor, we first show that the active sensing significantly improves the localization of static contacts, compared to passive or dynamic sensing. Subsequently, we confirm the sensor's robustness to different objects and background noises. Finally, we demonstrate the sensor's suitability to a robotic environment by testing it on a Panda robot arm. [1]

A. Localizing Contact with a Classification Rate of 93 %

We demonstrate the usefulness of the active acoustic sensor by evaluating the classification rate of static contacts at different locations across the actuator. When comparing

[1]The code and data are available online: http://dx.doi.org/10.14279/depositonce-9711

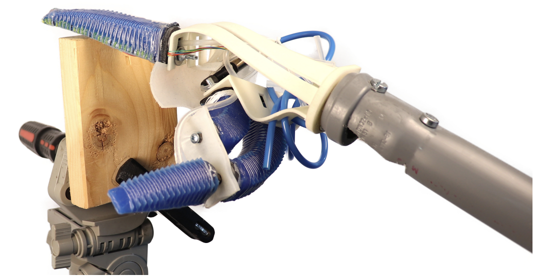

(a) An acoustic sensor finger on an RBO Hand 2 is in contact with a wooden object.

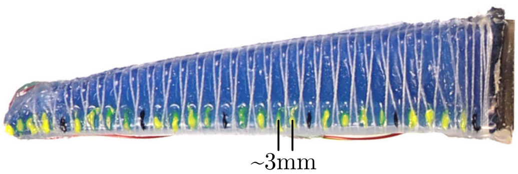

(b) We test the spatial accuracy of the acoustic sensor with 30 contact locations at a distance of ca. 3 mm.

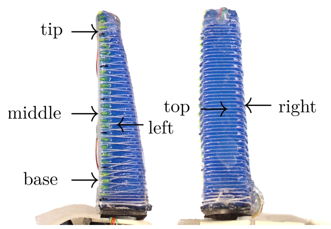

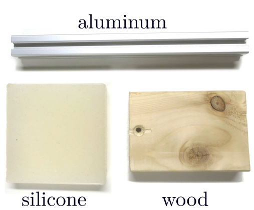

(c) Six contact locations all around the finger show our method turns the whole actuator into a sensor.

(d) The sensor is unaffected by different object shapes and materials.

Fig. 4. In our experiments, we evaluate the active acoustic sensor by localizing contact on all sides of the actuator and analyzing its spacial accuracy and robustness to different objects.

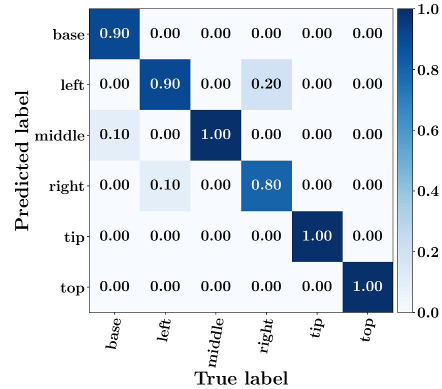

(a) Active Sensing

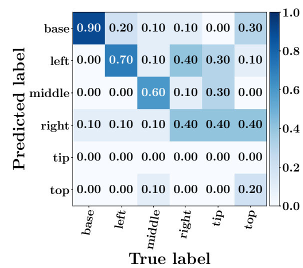

(b) Passive Sensing

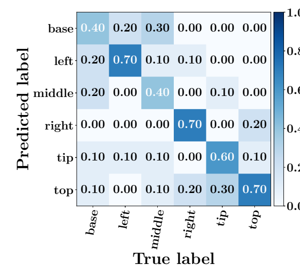

(c) Dynamic Sensing

Fig. 5. Acoustic sensing of six contact locations: The high values on the diagonal of the confusion matrix for active sensing show that the active acoustic contact sensor improves upon both passive sensing and dynamic sensing with classification rates of 93 %, 47 % and 58 % respectively.

to results of using no sound (passive) and the sound during impact (dynamic), we expect to see that the active approach improves the accuracy of contact location prediction.

Setup: For this experiment, the sensorized PneuFlex on the RBO Hand 2 is used with the manual handle (Fig. 4(a)). To demonstrate the sensor's ability to measure contact anywhere on the actuator, we define six contact locations distributed across its hull: Three on the palmar side (tip, middle, and base) and one each on the other sides (left, top, and right) (Fig. 4(c)). A wooden block is mounted on a tripod as the contact object. A human operator places the actuator in contact with the object, varying both the contact location on the actuator and the side of the object. This variation ensures that we learn to recognize the contact location and not just the actuator's pose. Static contact is held while the microphone records. The speaker either emits the 1 s sweep (active) or no sound (passive). For better comparison to our first paper [1], we also include dynamic contacts. For this, we record the impact sound of the actuator making contact while maintaining all other parameters, e.g. recording duration, etc. For the three sensor types, we record 25 samples of six contact locations for a total of 150 sounds each. The recording order is randomized to prevent any temporal effects in the data. Each dataset is split into 90 samples for training and 60 samples for testing, with equal class distributions.

In this paper, we only record data using the uninflated actuator, as it would be used, for example, in pre-grasp interactions or tactile exploration. Nonetheless, we acknowledge that sensing at different inflation levels is essential for a sensorized soft actuator. For this, we refer to our previous work [1], which showed that dynamic acoustic sensing works well regardless of inflation.

Results: The confusion matrix in Figure 5 shows both true and predicted contact location and the corresponding predictions. The high diagonal values for the active sensor indicates good classification. The near-perfect predictions (the classification rate is at 93%) demonstrate the impressive performance of the active sensing approach. In comparison, the passive and dynamic sensors are only able to achieve classification rates of 47% and 58%, respectively. This shows that the active sound transmits relevant information about the contact which is not available to the sensor in the other cases.

Interestingly, the passive sensing of static contacts, without any added sound, is still significantly better than the random chance baseline (17%). It seems the classifier picks up on small regularities in the sounds created from holding the contact and ambient noise. Nevertheless, active sensing clearly outperforms the other two approaches and demonstrates that successful acoustic localization of static contacts is possible anywhere on the actuator.

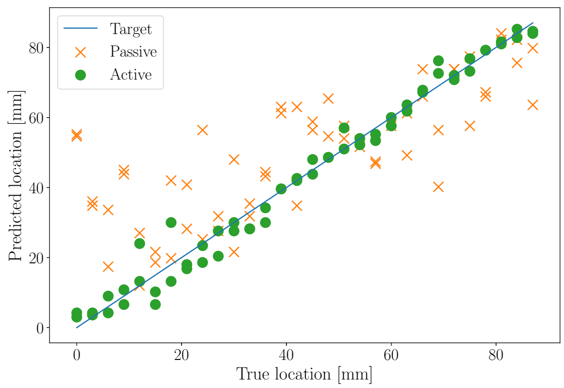

Fig. 6. Spatial accuracy of the acoustic sensor: Active sensing (green) shows high accuracy, predicting close to the target (blue) with a regression RMSE of 3.7 mm. Passive sensing (orange) shows the correct tendency, but is less accurate with a regression RMSE of 18 mm.

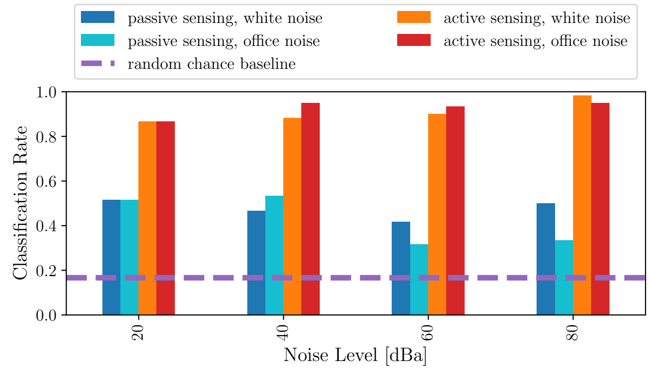

Fig. 7. A loud environment does not affect our contact sensing performance: Four experiments were recorded at four noise levels from very quiet (20 dBa) to very loud (80 dBa). The classification rate for active sensing (red and orange) does not drop for higher volumes and only slightly for passive sensing (light and dark blue).

B. Localizing Contact with a Spatial Accuracy of 3.7 *mm*

The six contact locations in the previous section had a distance of roughly 2 cm. In this experiment, we record data for contact points 3 mm apart and train a sensor model to predict the exact contact location. The mean prediction error determines how close together contacts can be for their sounds to be still distinguishable.

Setup: We mark 30 equally-spaced contact points along the 87 mm long palmar side of the actuator (Fig. 4(b)). A human operator holds the actuator against the edge of a wooden block. We record two datasets, one active and one passive, with five samples for each of the 30 contact points for a total of 150 sounds, each. We use the default k-nearest neighbor regressor from the scikit-learn library [27] with 90 samples for training and 60 samples for testing.

Results: Figure 6 shows a cluster plot of the predicted contact points for both passive and active sensing. The blue line indicates the target if each contact were to be predicted perfectly. The active sensing prediction is relatively close to the target, with a root-mean-square error (RMSE) of 3.7 mm. Towards the tip of the actuator, the predictions appear to be more accurate. The passive sensing predictions are generally less accurate, with an RMSE of 18.0 mm. However, a correct tendency is visible in the slight slope of the data points.

While a more finely tuned regression method is likely to further improve this result, a sensor resolution of less than 4 mm demonstrates the high accuracy achievable with active acoustic sensing, which allows detailed measurements for grasping and manipulation tasks.

C. Sensing Success is Unaffected by Noise and Object Type

1) Robustness to Background Noise: To characterize the sensor, we analyze its robustness to changes in the environment, for example, loud external noises. In this experiment, we evaluate if a noisy environment affects the classification performance for passive and active contact location sensing.

Setup: We record six sets of 60 recordings with two types of background noise (white noise and generic office noise) at three volumes (40 dBa, 60 dBa, 80 dBa). A dataset with 150 recordings in a quiet room (20 dBa) is used as baseline. A KNN classifier is trained on 90 recordings of the baseline set and then used to predict all test sets.

Results: The classification rate of active sensing is consistent across all noise levels (Fig. 7). It even increases slightly with more background noise, which might be due to the noise drowning out other distractors. The type of noise does not appear to have a noticeable influence. The results for passive sensing deteriorate a little with an increasing volume of office noise; the effect of white noise is less pronounced. But even at 80 dBa, the passive sensing classification rate is roughly twice as high as the random chance baseline. These results show that background noise does not negatively affect the sensor performance of the active sensor. This aligns with the findings from our first paper [1], where background noise did not influence the dynamic acoustic sensing.

2) Robustness to Different Objects: Different object materials and geometries have different acoustic properties. Since our sensor uses the acoustic properties of the actuator, changes to the object might affect the sensors measurement results. We test this by training and testing the sensor on objects of three different materials: wood, silicone, metal. If the type of object influenced the sensor, we would expect to observe worse measurement performance when using a sensor model that was trained on one object being used to predict contact with a different object.

Setup: We use the same experimental setup as in Section IV-A and additionally record 2 more sets of 150 sounds, one with a silicone block as the contact object, and one with a metal strut (Fig. 4(d)).

Results: Figure 8 shows the cross-prediction results. There is no noticeable difference between the diagonal (trained and tested on the same object) and the off-diagonal values (tested on a different object than the training). This shows that active acoustic sensing is very robust to different objects.

It is, however, relevant to note that *classifying* the object material from the active sound signal is not as successful (material classification rate of 50.5% over the random chance

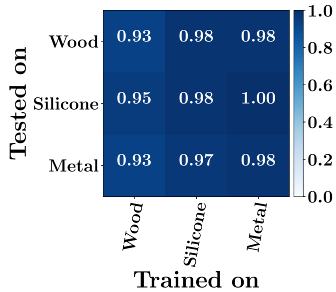

Fig. 8. Active contact sensing is not influenced by object type and material: Four contact location classifier were trained on the objects shown in Fig. 4(d). The matrix shows the classification rate of each classifier when predicting test sets of different materials. The classification rate remains very high regardless of material.

baseline of 33.3%). This is a disadvantage compared to our previous, dynamic acoustic sensing approach, which was capable of differentiating materials with 69.4% accuracy [1]. To achieve both, a combination of dynamic and active sensing might be necessary. For real-world applications, however, it is encouraging to see that active acoustic sensing can handle different types of objects.

D. Acoustic Sensing Works Well in a Robotic Environment

Until now, we demonstrated that the principle of active acoustic contact sensing works very well. Next, we test if this holds true when deployed on a robot. Robotic environments contain several sources of potential interference for the acoustic sensor, for example, motor noises and vibrations or cross-talk from other acoustic sensors. We will first show the sensor working on a robot with multiple sensorized actuators, and then demonstrate the transfer of the sensor model to different workspace locations.

Robot setup: We attach four PneuFlex actuators, each equipped with the active acoustic contact sensor, to an RBO Hand 2. The hand is mounted on a Panda robot arm (Fig. 1(c)). A wooden contact object is mounted in a vice. We define six different contact locations on the index finger actuator. However, because the actuator sides are not easily reachable due to the other fingers, we instead use *tip*, *middle*, and *base* on both palmar and dorsal side of the finger. Each contact location is recorded 25 times, for a total of 150 samples. To test if there is interference with other acoustic sensors, we record one dataset where only the index finger is active, and another one with all four fingers active. As before, we train on 90 samples and test with 60.

Robot results: The classification rate for a single active finger is at 100%, which means that *all* contact locations from the test sets were correctly identified. Obviously, motor noises and vibration do not negatively affect the measurement. With four active fingers, the classification rate is still very high at 96.7%. The sounds from the other sensorized actuators do not affect the measurements much. This shows that multiple actuators, each with active acoustic contact sensors, can be operated in parallel, without much interference.

Transfer to other workspace locations: Finally, we evaluate if sensor models transfer between workspace locations, to make sure that the acoustic sensor learns the sound of contact and not the specific motor vibrations of each robot pose. This is an important prerequisite for using the sensor in robotic applications. For this, we place the contact object in three different workspace positions, each requiring different end-effector poses to make contact. We run two evaluations: First, we train the sensor separately at each of the three object locations with 90 training samples and 60 test samples. Then, we train on two of the object locations ($2*90$ training samples) and test on the third (60 test samples). This will show if the active acoustic sensor model transfers to novel object locations.

Transfer results: The results of the transfer test are very promising: For active sensing, the location classification rate for the three object positions when trained separately is 100% each time. Interestingly, the passive sensor, without the active sound, also predicts surprisingly well for the three positions separately (mean classification rate of 88.3%). However, when attempting to train on two positions and predicting the third, the classification rate drops to 40%. This indicates that the sensor likely did not learn to recognize the contact sound, but rather the specific motor noises of the robot arm for each pose. In contrast, the active sensor, when trained on two object positions and predicting on the third, still achieves a classification rate of 100%. This impressive result shows that the active acoustic contact sensor works very well in a robotic environment and that the sensor model transfers well between different object locations.

V. Conclusion

We presented an active acoustic sensor that measures static contacts on a soft actuator using sound. Embedded audio components emit and record a frequency sweep, which undergoes small changes while traveling through the actuator. These changes depend on the contact state of the actuator and can be distinguished with a simple k-nearest neighbor classifier. This effectively turns the complete actuator into a contact location sensor, without affecting the actuator's compliance. By using an active sound, the sensor no longer relies on impact sounds of dynamic contacts but instead measures the contact state at any time.

We demonstrated the active acoustic sensing approach on the soft PneuFlex actuator. The sensor successfully identified six contact locations with a classification rate of 93% and a regression accuracy of 3.7 mm. This is a significant improvement over passive and dynamic acoustic sensing with 47% and 58% classification rate, respectively. The approach is robust against background noise up to 80 dBa and transfers between different objects. The sensor works very well when mounted on an RBO Hand 2 on a Panda robot arm. It is unaffected by neighboring active sensors and contact location models transfer between different end-effector poses.

We believe this active acoustic sensor to be an impressive demonstration of the possibilities of using sound for sensing in soft robotics. By exploring other active sounds and more advanced learning methods, the range of measurable actuator states can certainly be expanded even further.

References

[1] G. Zöller, V. Wall, and O. Brock, "Acoustic Sensing for Soft Pneumatic Actuators," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Madrid, Spain, 2018, pp. 6986–6991.

[2] 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.

[3] R. A. Rompf, "Entwicklung eines akustischen Modells für einen weichen pneumatischen Aktuator," Master's thesis, Technische Universität Berlin, 2019.

[4] R. Deimel and O. Brock, "A compliant hand based on a novel pneumatic actuator," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2013, pp. 2047–2053.

[5] 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, pp. 171–187, 2011.

[6] Z. Kappassov, J.-A. Corrales, and V. Perdereau, "Tactile Sensing in Dexterous Robot Hands —Review," *Robotics and Autonomous Systems*, vol. 74, pp. 195–220, 2015.

[7] K. M. Digumarti, B. Trimmer, A. T. Conn, and J. Rossiter, "Quantifying Dynamic Shapes in Soft Morphologies," *Soft Robotics*, vol. 6, no. 6, pp. 733–744, 2019.

[8] P. Polygerinos, N. Correll, S. A. Morin, B. Mosadegh, C. D. Onal, K. Petersen, M. Cianchetti, M. T. Tolley, and R. F. Shepherd, "Soft Robotics: Review of Fluid-Driven Intrinsically Soft Devices; Manufacturing, Sensing, Control, and Applications in Human-Robot Interaction," *Advanced Engineering Materials*, vol. 19, no. 12, p. 1700016, 2017.

[9] V. Wall, G. Zöller, 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)*, Singapore, 2017, pp. 4965–4970.

[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] 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.

[12] K. C. Galloway, Y. Chen, E. Templeton, B. Rife, I. S. Godage, and E. J. Barth, "Fiber Optic Shape Sensing for Soft Robotics," *Soft Robotics*, vol. 6, no. 5, pp. 671–684, 2019.

[13] C. Chorley, C. Melhuish, T. Pipe, and J. Rossiter, "Development of a tactile sensor based on biologically inspired edge encoding," in *2009 International Conference on Advanced Robotics*, 2009, pp. 1–6.

[14] 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.

[15] M. Drumheller, "Mobile Robot Localization Using Sonar," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-9, no. 2, pp. 325–332, 1987.

[16] Liang-Ting Jiang and J. R. Smith, "Seashell effect pretouch sensing for robotic grasping," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2012, pp. 2851–2858.

[17] 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.

[18] O. Kroemer, C. H. Lampert, and J. Peters, "Learning Dynamic Tactile Sensing With Robust Vision-Based Training," *IEEE Transactions on Robotics*, vol. 27, no. 3, pp. 545–557, 2011.

[19] S. Luo, L. Zhu, K. Althoefer, and H. Liu, "Knock-Knock: Acoustic Object Recognition by using Stacked Denoising Autoencoders," *Neurocomputing*, vol. 267, pp. 18–24, 2017.

[20] E. Torres-Jara, L. Natale, and P. Fitzpatrick, "Tapping into Touch," L. Berthouze, F. Kaplan, H. Kozima, H. Yano, J. Konczak, G. Metta, J. Nadel, G. Sandini, G. Stojanov, and C. Balkenius, Eds., vol. 123. Lund University Cognitive Studies, 2005, pp. 79–86.

[21] Tomoaki Nakamura, Takayuki Nagai, and Naoto Iwahashi, "Multimodal Object Categorization by a Robot," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2007, pp. 2415–2420.

[22] J. Sinapov and A. Stoytchev, "Object Category Recognition by a Humanoid Robot Using Behavior-Grounded Relational Learning," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2011, pp. 184–190.

[23] C. Schenck, J. Sinapov, D. Johnston, and A. Stoytchev, "Which Object Fits Best? Solving Matrix Completion Tasks with a Humanoid Robot," *IEEE Transactions on Autonomous Mental Development*, vol. 6, no. 3, pp. 226–240, 2014.

[24] M. Ono, B. Shizuki, and J. Tanaka, "Touch & activate: adding interactivity to existing objects using active acoustic sensing," in *Proceedings of the 26th annual ACM symposium on User interface software and technology*, ser. UIST '13. St. Andrews, Scotland, United Kingdom: Association for Computing Machinery, 2013, pp. 31–40.

[25] A. Mujibiya, X. Cao, D. S. Tan, D. Morris, S. N. Patel, and J. Rekimoto, "The sound of touch: on-body touch and gesture sensing based on transdermal ultrasound propagation," in *Proceedings of the 2013 ACM international conference on Interactive tabletops and surfaces*, ser. ITS '13. St. Andrews, Scotland, United Kingdom: Association for Computing Machinery, 2013, pp. 189–198.

[26] K. Takaki, Y. Taguchi, S. Nishikawa, R. Niiyama, and Y. Kawahara, "Acoustic Length Sensor for Soft Extensible Pneumatic Actuators With a Frequency Characteristics Model," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 4292–4297, 2019.

[27] F. Pedregosa et al., "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.