

# Multi-Task Sensorization of Soft Actuators Using Prior Knowledge

Vincent Wall

Oliver Brock

**Abstract**—The space of all possible deformations of soft robotic actuators is extremely large. It is impossible to explicitly measure each internal degree of freedom, regardless of the number and types of sensors. It is, however, possible to measure a smaller subset of task-relevant deformations using only a few well-placed sensors. But for a different task, the soft actuator’s deformation behavior might differ significantly. Instead of finding a new sensor placement for the new task, which would result in a separate hand for every task, we propose a method that maintains the original sensors and uses prior knowledge about each task to extend the applicability of the existing sensorized actuators to new tasks. We demonstrate our approach by the example of a PneuFlex actuator of the RBO Hand 2. When sensorizing the actuator for a single task, the sensor model does not transfer well to other tasks. Using our multi-task method, we train new sensor models that use prior knowledge about the tasks. The new models improve measurement accuracy for the new tasks without having to change the sensor hardware.

## I. INTRODUCTION

Soft actuators have many internal degrees of freedom due to their compliant design. While this makes them ideal for grasping and manipulation applications that require safety and robustness [1], it also makes sensorization a challenge. The space of all possible deformation states of a soft actuator is extremely big, and no amount of sensors could measure it exactly. However, many of the possible deformations do not typically occur. For a given task, only a small set of deformations can be expected to appear. This drastically reduces the sensorization difficulty as only the subspace of expected deformations needs to be assessed. When the task is known, we can find a small set of sensors that measure the task-relevant deformation features of the actuator [2].

For different tasks, however, the behavior of the soft actuator may be very different [3]. The expected deformations may not lie in the same subspace. Instead of building new sensorized actuators for each new tasks, it would be preferable to reuse existing sensor hardware. But given the complex deformation behaviors of soft actuators, it is likely that the combined deformation subspaces of multiple tasks can no longer be described with a simple sensor layout.

In this paper, we present a method that achieves multi-task sensorization for soft actuators while maintaining simple sensorization hardware. By using prior knowledge about which

of the different tasks the actuator is currently performing, we can reduce the space of expected deformations back to a simple subspace. This allows us to interpret the sensor data accordingly, without requiring complex sensor layouts. This way we can use the same sensor hardware for multiple tasks with different deformation behavior.

We demonstrate our approach on the PneuFlex actuator of the RBO Hand 2 [4], [3]. We identify three manipulation tasks and show that sensor models that successfully predict deformations for one task, perform significantly worse for other tasks. This confirms that the deformation behaviors lie in separate subspaces. Subsequently, we employ our multi-task sensorization method to maintain the sensor hardware found for task one, while improving the deformation prediction by a factor of 3 and 1.5 for tasks two and three, respectively. This shows that we can reuse existing sensorized actuators for new tasks by using prior knowledge about the actuator’s expected deformation subspace.

## II. RELATED WORK

The goal of sensorizing soft actuators is to measure the relevant interactions with an object or the environment. This feedback allows to compensate motion uncertainty and react to unexpected contact events [5]. This requires novel, flexible sensor technologies. Additionally, we need to identify where to place those sensors on the actuator, as we can never measure the complete state of a soft actuator.

### A. Soft Sensor Technologies

With the rising interest in soft robotics, the field of flexible sensors has also gained attention recently. A comprehensive overview of various soft sensor technologies is given in [6] and [7]. The collection of technologies ranges from strain and contact sensing to deformation and curvature sensing.

In this paper, we use the PneuFlex actuator [4]. This pneumatic actuator is made of highly flexible silicone. We have found that liquid metal strain sensors are well-suited for these actuators because they are very flexible and easily customizable [2]. However, the multi-task sensorization method we present in this paper is independent of the sensor type and applies to any sensorized actuator that has more internal degrees of freedom than can be sensed directly.

Liquid metal strain sensors are based on metal alloys that are liquid at room temperature, e.g. eutectic gallium-indium (EGaIn). These sensors can be used to detect strain through change in resistance [8]. The sensor can be embedded in the actuator during the fabrication [9], or prepared separately and attached as needed [10]. This gives a lot of freedom when sensorizing a soft actuator. A flexible, manual fabrication

All authors are with the Robotics and Biology Laboratory, Technische Universität Berlin, Germany

We gratefully acknowledge financial support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy - EXC 2002 ”Science of Intelligence” - project 17, the European Commission (SOMA, H2020-ICT-645599), and by the Alexander von Humboldt Foundation and the Federal Ministry of Education and Research (BMBF).

method was presented by Farrow and Correll [11]. It allows for highly customized sensor layouts, without the need for special equipment. We adopt this technique to create the sensor layouts for our PneuFlex actuator.

### B. Using Sensorized Soft Actuators

Even when the preferred sensor technology is chosen, we still need to identify where to place the sensors on the actuator. Because the complete space of possible deformations of a soft actuator is too big to measure completely, any sensor placement will need to consider the desired task and the task-relevant subspace of actuator deformations.

Task-specific sensor placements are used to measure the surface shape of an actuator [12] and to identify objects by their shape and size by measuring the passive shape adaptation of the actuator [13], [11]. Strain sensors on pneumatic actuators have also been used to implement force and position control [14]. These applications identified a specific sensor placement to achieve a specific task. More generally, we have shown that for a given task it is possible to find a small set of sensors to measure the subspace of task-relevant actuator deformations [2]. However, the task imposes the reduced deformation space that allows to use only few sensors. When the task changes, the same sensorization may no longer be able effective, because the deformation behavior exists in a different subspace.

In this paper, we present a method that does not require to change the sensor hardware in order to adapt a soft actuator to new tasks. This is possible because we use prior knowledge about the different subspaces of deformation behaviors that correspond to the different tasks.

## III. SELECTING TASKS AND TASK-RELEVANT DEFORMATIONS

The method we are presenting is meant to enable an existing sensorized actuator to be used for novel tasks. This is only necessary if the deformation behavior of the actuator is significantly different in these tasks. Then the old sensor model can no longer be used. In this section, we describe three common manipulation tasks that result in such varying deformations of the actuator. Additionally, we identify a set of relevant deformations for all three tasks. Using a set of strain sensors, we will attempt to measure these "interaction features". The accuracy by which the sensor model is able to predict these features will be our metric to evaluate the applicability of a given sensorization for a specific task.

### A. Manipulation Tasks with Dissimilar Deformations

We are interested in tasks that are common in grasping and manipulation applications and involve interactions of the actuator with an object or the environment. Furthermore, the tasks should involve different types of deformation of the actuator, as this will make it difficult for a sensorization that was designed for one task, to transfer to another.

In our lab, we use the RBO Hand 2 for several manipulation applications. Its soft PneuFlex actuators are highly

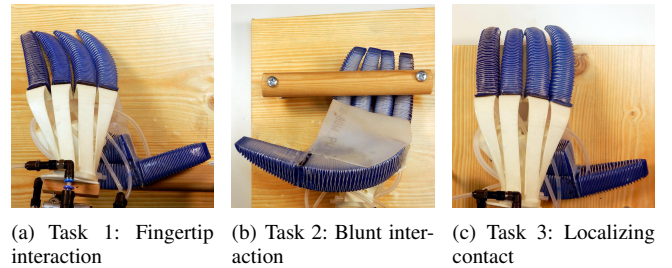


Fig. 1. Three common manipulation tasks for the RBO Hand 2; Because the hand interacts in different ways in each task, the deformation behavior of the fingers is not the same.

compliant and underactuated, which allows exploiting contact with the environment in various ways. The following three tasks commonly occur during such applications and benefit from sensor feedback (Figure 1):

*Task 1 - Fingertip interaction:* This task consists of interactions of the hand with the environment, in which the fingertips make contact. This often occurs when sliding objects or during a pinch grasp [15].

*Task 2 - Blunt interaction:* The interactions during this task involve higher forces, often localized around the central part of the finger. This occurs during pushing and pulling operations, and when lifting or holding a heavy object [16].

*Task 3 - Localizing contact:* In this task, the hand is used to localize the edge of an object. By detecting the relative position between object and hand we can compensate perceptual and motion uncertainty [5].

### B. Selecting Relevant Interaction Features

Exact reconstruction of the actuator's shape is infeasible in all of these tasks. The space of possible deformations is too big. But to successfully perform the tasks, we only need to measure certain aspects of the interactions between actuator and object. In the following, we describe a 10-dimensional vector of interaction features and explain why these are relevant for the tasks (Figure 2):

*Overall deformation (3D):* The cumulative deformation in flexional, lateral, and twist direction describes the overall shape of the actuator. While this feature disregards where along the actuator deformation occurs, it is a good indication of grasp success and collisions.

*Forces and torques (6D):* Measuring interaction forces helps to prevent damage, detect collisions, and reason about object properties like weight. The reference frame is placed at the base of the actuator, as this is where forces and torques are transmitted to the mounting scaffold.

*Contact location (1D):* This feature describes where along the side of the actuator a contact occurs. Knowing the contact location allows compensating for other sources of uncertainty, e.g. in perception or motion execution.

## IV. SINGLE-TASK SENSOR PLACEMENT

Before introducing our method that extends existing sensor layouts to additional tasks, we first need to find the sensor placement for a single task. For this, we adopt the sensorization method previously presented in [2].

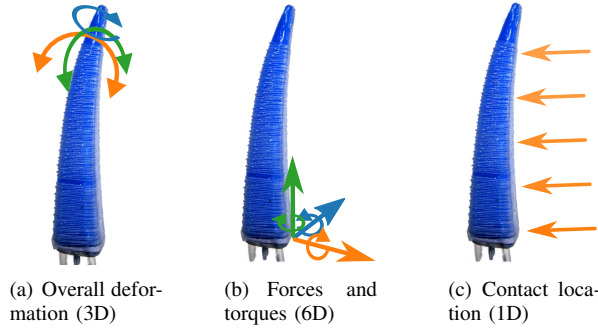


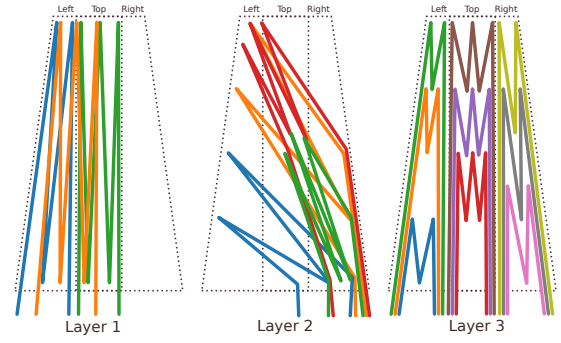
Fig. 2. The 10-dimensional vector of relevant interaction features illustrated on the PneuFlex actuator

#### A. Layout of Liquid Metal Strain Sensors

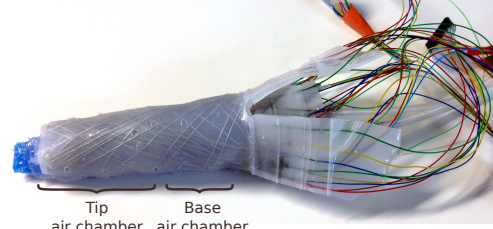
The first step of the single-task sensorization is to create a redundant sensor layout, from which the most useful sensors can be selected. We distribute 16 liquid metal strain sensors across the hull of the actuator. By covering sections of the hull with multiple different sensors, we introduce redundancy. For easier manufacturing, the strain sensors are divided into three separate layers (Figure 3). Layer 1 consists of three sensors that cover the full length of the actuator. These sensors were chosen because they have the potential to measure the overall shape of the actuator. Layer 2 consists of four sensors that are placed diagonally at different angles and sections of the actuator. These sensors were chosen because they have the potential to measure twist, as well as detecting a difference between tip and base deformations. Layer 3 consists of nine sensors, each focused on a different part of the actuator’s hull. These sensors were chosen because they have the potential to localize the contact location by detecting variances in deformation along the actuator. All three layers are attached to the hull of the PneuFlex actuator (Figure 3). In addition to the 16 strain sensors, we measure the pressure in the base and tip air chamber with two pressure sensors (MPX4250DP). They are needed for actuation, anyway, and have the potential to be used to measure forces or contact locations through relative pressure changes.

#### B. Predicting Interaction Features from Sensor Data

The next step is to use the sensor information to predict the task-relevant deformations. For this, we train a sensor model that maps sensor data to interaction features. The previous sensorization method used polynomial regression [2]. However, the interaction features describing the *force/torque* and *contact location* appear to be too complex for such a simple approach. Instead, we use a basic multilayer perceptron (MLP) which is better suited to encode these non-linear interactions. Using the scikit-learn framework [17], we performed a parameter search and identified the following network parameters: 3 hidden layers of 64 nodes, a learning rate of 0.01 for the “adam” solver, and early stopping with a validation fraction of 0.1. The sensor data (input data) and the interaction features (target data) are both normalized to zero mean and unit variance.



(a) The 16 strain sensors across 3 layers cover different parts of the actuator’s hull. Each color represents a separate sensor.



(b) PneuFlex actuator with 16 strain sensors

Fig. 3. The redundant strain sensor layout and the sensorized prototype

To identify which sensors are most useful, we perform an iterative construction of sensor sets. This is similar to the recursive feature elimination of the original approach [2] but reverses the direction. It starts by identifying the best 1-sensor layout, then adds the next best sensor and so on. This emphasizes few-sensor layouts, which is what we are usually interested in. Each sensor layout is evaluated with 3-fold cross-validation, using the mean squared error (MSE) between ground truth and model-prediction, averaged over the ten-dimensional normalized interaction features.

As a result of this procedure, we obtain the order of usefulness of all 18 sensors (16 strain, 2 pressure) for a given task. This order tells us the task-specific, most effective  $n$ -sensor layout, where  $n$  is the number of sensors we choose to use.

#### V. MULTI-TASK METHOD EXTENSION USING PRIOR-KNOWLEDGE

In the previous section, we explained the sensorization process for a single task. One possibility to approach the sensorization for additional tasks would be to simply repeat the process for each task. However, this would result in a new sensor layout for each task. It would be preferable not to build new hardware for each task. An alternative would be to use the existing layout and train a sensor model to predict deformations in both the old and the new tasks. But as the experimental validation will show, the deformation spaces of different tasks can be very different. Thus, a simple sensor layout will not be able to describe both tasks at the same time. Instead, we present a method that maintains an existing sensorization from the first task but uses prior knowledge to distinguish between different deformation subspaces.

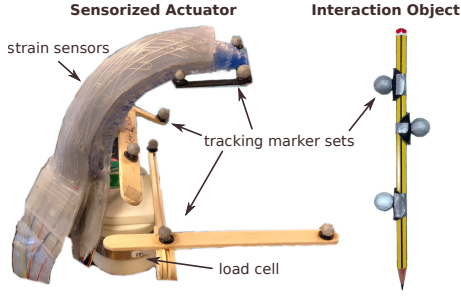


Fig. 4. Experiment setup with sensorized PneuFlex prototype, tracking markers, load cell, and interaction object

#### A. Using Prior-Knowledge in the Sensor Model

The prior knowledge we use consists of the information which of the tasks is currently active. The active task information is included as three additional input features to the multilayer perceptron regressor. The three signals are one-hot encoded inputs that indicate which of the three tasks the sensor measurement originated from. Using the extended sensor input we retrain the MLP on data from both the original Task 1, as well as the new additions Task 2 and 3. During this phase we keep the sensor sets fixed, that means we do not change the hardware that was found for Task 1.

### VI. EXPERIMENTAL VALIDATION

#### A. Gathering Ground Truth Data

The first step in validating our method is to record ground truth data of sensor measurements and corresponding interaction features. We use the data to train and evaluate the different sensor models.

**Hardware Setup:** Figure 4 shows the setup of the experiment. The sensorized PneuFlex prototype is equipped with three sets of tracking markers: at the base, between the two air chambers, and at the fingertip. A motion capture system (Motion Analysis) records marker positions. The actuator is mounted on top of a load cell (ATI Mini40) that measures forces and torques at the actuator’s base. To measure where along the actuator contact occurs, the interaction object is also equipped with tracking markers. Finally, sensor data from the 16 strain and the 2 pressure sensors are recorded at 100Hz using a data acquisition unit (LabJack U6).

**Experimental Protocol:** To test if the sensor models generalize over different actuator inflations, we record data at nine pressure levels: [40, 120, 235] kPa for the base chamber, combined with [20, 105, 200] kPa for the tip chamber. For *Task 1 - Fingertip interaction* we push, pull, and twist the tip of the actuator and recorded 100k samples (ca. 16 min). For *Task 2 - Blunt interaction* we simulate the expected interactions by pushing the midsection of the actuator with a blunt object. We apply forces of various magnitudes (0–8 N) and directions ( $-90^\circ$  to  $90^\circ$ ). We recorded 50k samples (ca. 8 min) for Task 2. For *Task 3 - Localizing contact* the tip of an object is used to contact the actuator at different positions along its palmar side. The force of the contact is smaller compared to the blunt interaction task ( $<5$  N). We

recorded 40k samples (ca. 6 min) for Task 3. Additionally, we recorded separate test sets for each of the tasks with 25k samples (ca. 4 min).

**Data Pre-Processing:** The 10-dimensional vector of interaction features is determined from the measured ground truth of motion tracking and load cell. The flexional, lateral, and twist of the *overall deformation* (3D) is calculated as the pitch, yaw, and roll angles, respectively, of the rotation between the base frame and the fingertip frame, measured by the tracking markers. The *forces and torques* (6D) are directly measured with the load cell. The *contact location* (1D) is calculated as the intersection point between the tracked object and actuator, whenever the absolute force value indicated any contact. Example segments of corresponding sensor data and interaction features for each of the three tasks are shown in Figure 5.

#### B. Comparing the Deformation Behaviors of Tasks

With the recorded data we can train task-specific sensor models and evaluate their accuracy in predicting the observed deformations. If it is indeed possible to measure the selected subspace of task-relevant interaction features using only a few sensors, we should see small prediction errors even with a small number of sensors in the layout. We would also expect that error decreases with more sensors, as more information about the actuator is available. Figure 6 shows the mean squared errors (MSE) of the model prediction on the test data set. The shaded regions indicate the standard error of the mean of the ten interaction features. We trained a separate sensor model for each task and evaluated their predictions for data samples from all three tasks.

The sensor models perform reasonably well for tasks they were trained for, e.g. Task 1 in Figure 6(a). The fact that the models predict the test data with a low MSE shows that a lower-dimensional subspace of deformations exists and can be measured with only a few sensors.

The second observation is the steady decline of the prediction error with an increasing number of sensors. With more sensor measurements available the prediction gets more accurate. This points to an important trade-off: Prediction accuracy versus number of sensors. While each additional sensor improves the prediction, they also influence the actuator’s compliance, increase fabrication time, and introduce additional failure points. This is a fundamental challenge in the sensorization of soft actuators, which our method cannot solve. However, by returning an importance-sorted list of sensors the method enables an informed decision. Instead of making this decision now, we will use the importance order defined for the single-task case of Task 1 and assume it as our *fixed* sensor order. This way we evaluate the usefulness of prior knowledge for any  $n$ -sensor layout.

Finally, we analyze the prediction accuracy for tasks the sensor model was not trained on. This test shows how similar the deformation behavior is between tasks. In the case of comparable deformation subspaces, the sensor model should predict the other tasks equally well. However, the predictions of untrained tasks in Figure 6 is noticeably worse. The only

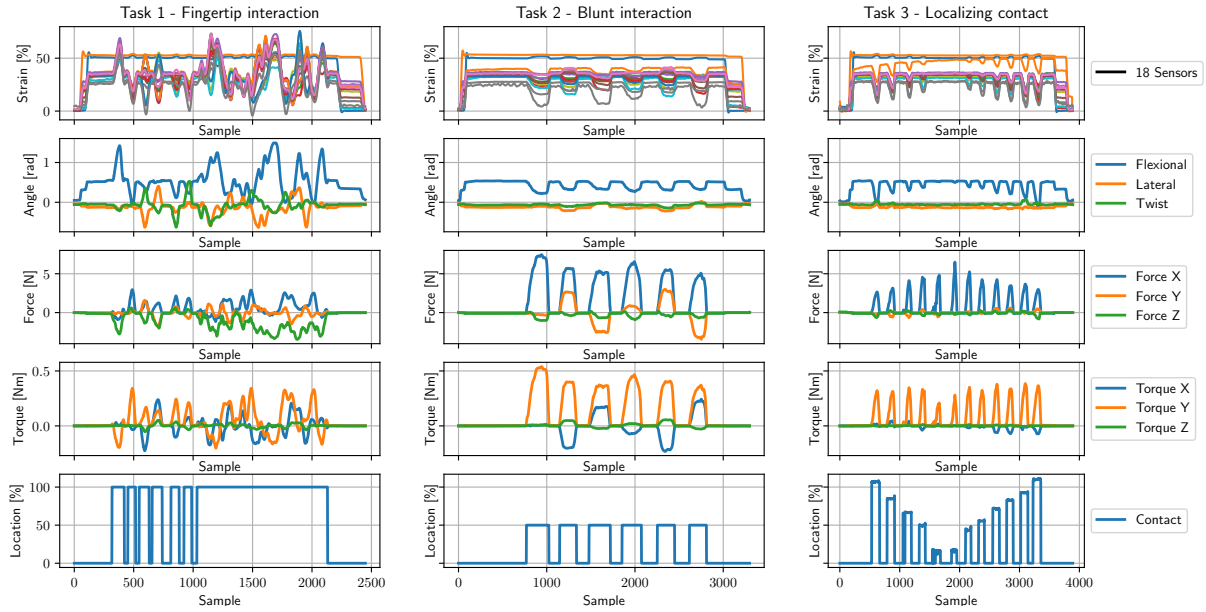


Fig. 5. Three example segments of sensor data and corresponding interaction features: The deformation behavior of the actuator is different from task to task. This is visible in the distinct patterns of the interaction features and makes it unlikely that sensor models can be transferred between tasks.

exception is Task 3 in Figure 6(b), where the trained Task 2 and untrained Task 3 have similar prediction errors. But since the opposite is not true, i.e. Task 2 is not predicted equally well in Figure 6(c), the models are not interchangeable. Altogether, it can be seen that sensor models that were designed for one specific task, are not necessarily applicable to measure the same interaction features for a different task.

### C. Extending an Existing Sensorization to New Tasks

The final step of the validation of our sensorization method is to evaluate the multi-task performance of sensor models that use prior information about the tasks. We claim that by providing data about which task is active, the sensor model can focus on the task-specific subspace of deformation behaviors and consequently improve its prediction accuracy.

To test this we first evaluate the extensions to Tasks 2 and 3 separately, and finally for all three tasks combined. Each retraining of the sensor model with prior knowledge is repeated three times. The plots in Figure 7 show the average of the prediction MSE. The shaded area indicates the standard error of the mean across the three repeats.

Comparing the difference in prediction error between the single-task model and the prior knowledge model, the improvement is clearly visible. For the extension to Task 2 (Fig. 7(a)), the MSE is reduced by a factor of 3. In the case of Task 3 (Fig. 7(b)), the improvement is less drastic but still notable at a factor of ca. 1.5. In both plots, the performance of Task 1 stays almost the same. This means, the addition of the other tasks does not influence the model's ability to predict Task 1 data, which is good. These results validate our method and show that it is not necessary to change the sensors of an existing sensorized actuator to use it for novel tasks. By including prior knowledge about the tasks, the sensor model can distinguish between the different subspaces of expected deformation behavior of each task.

Finally, the third plot (Fig. 7(c)) shows the results when both Task 2 and 3 are added to the model. Again, Task 2 is predicted nearly as well as Task 1. However, Task 3 only improves for more than ten sensors. If the prior knowledge model had truly learned to differentiate between the tasks, the prediction in the three-task case should be comparable to the two-task cases. The fact that the performance is worse—even the Task 1 prediction deteriorates slightly—likely indicates that the network size is too small, to account for all tasks at the same time. Increasing the number of hidden layers and nodes should alleviate this problem.

## VII. CONCLUSION

We presented a method for the multi-task sensorization of soft actuators that uses existing sensor hardware for new tasks by using prior knowledge about the expected actuator behavior. This is possible even though the space of actuator deformations is high-dimensional because each task corresponds to a smaller subspace of expected deformations. By including prior knowledge about which task is active, the sensor model used to predict the deformation is aware of what subspace to expect and interprets the sensor data accordingly. We demonstrated our approach for the PneuFlex actuator of the RBO Hand 2. For three common grasping and manipulation tasks (Fingertip interaction, Blunt interaction, and Localizing contact) we identified task-relevant interaction features (3D overall deformation, 6D forces and torques, and 1D contact location). A multilayer perceptron regressor maps data from up to 16 liquid metal strain sensors and 2 pressure sensors to the interaction features. The sensor models trained with prior knowledge achieve a 3-fold and 1.5-fold reduction of the prediction errors for two new tasks while maintaining the existing sensor hardware.

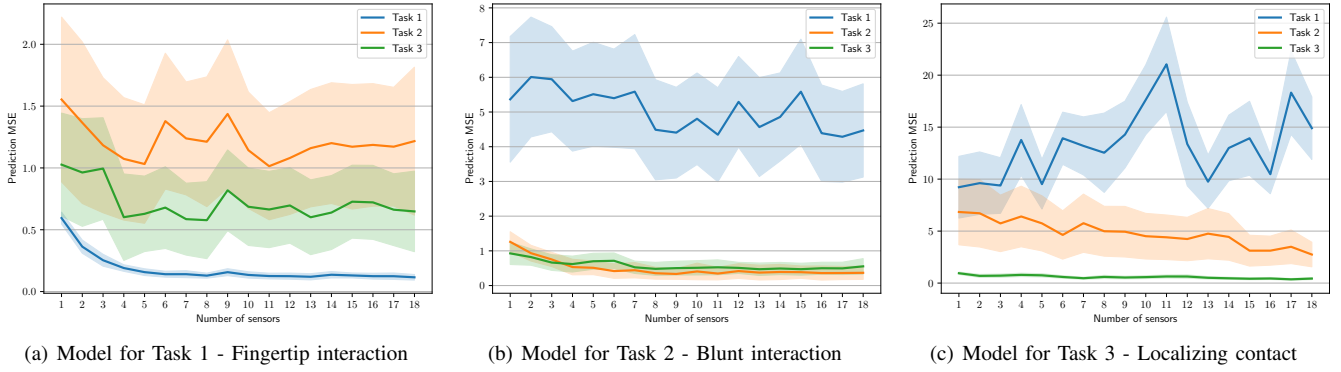


Fig. 6. Prediction performance of single-task models (standard error of the mean indicated as shadow): Each model is trained for one task and evaluated on all three. While the prediction error for the designated task is low, it is noticeably higher for the other tasks. This shows that sensor models can be found for one task, but do not transfer well to other tasks.

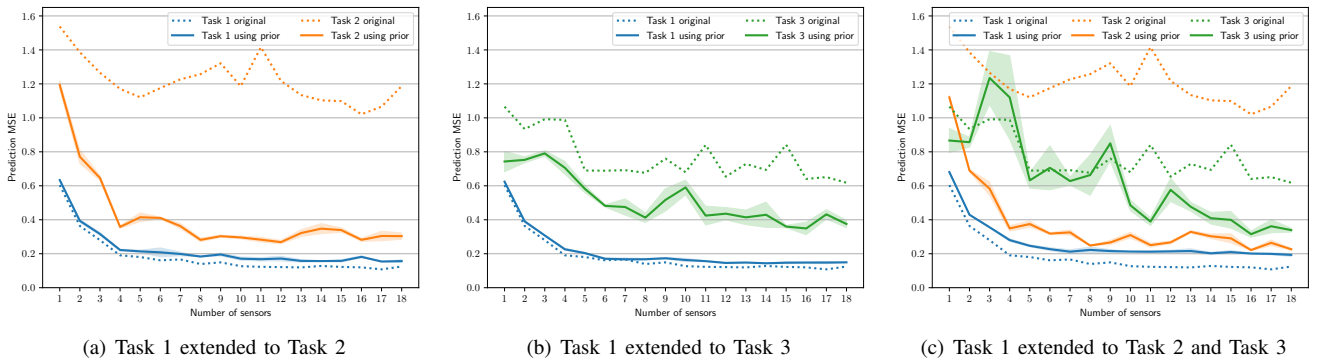


Fig. 7. Prediction performance of the prior knowledge models (original performance is shown with dotted lines, standard error of the mean indicated as shadow): By including information about the task-specific deformation subspace, the sensor model improves prediction of new tasks without impacting the original task. Only when extending to both new tasks simultaneously, the MLP regressor seems to be too small to accurately describe all cases.

## REFERENCES

- [1] J. Hughes, U. Culha, F. Giardina, F. Guenther, A. Rosendo, and F. Iida, "Soft manipulators and grippers: A review," *Frontiers in Robotics and AI*, vol. 3, p. 69, 2016.
- [2] 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)*, 2017, pp. 4965–4970.
- [3] 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, March 2016.
- [4] —, "A Compliant Hand Based on a Novel Pneumatic Actuator," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2013, pp. 2047–2053.
- [5] E. Páll, A. Sieverling, and O. Brock, "Contingent contact-based motion planning," in *International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 6615–6621.
- [6] M. Amjadi, K.-U. Kyung, I. Park, and M. Sitti, "Stretchable, skin-mountable, and wearable strain sensors and their potential applications: A review," *Advanced Functional Materials*, vol. 26, no. 11, pp. 1678–1698, Mar. 2016.
- [7] 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, Dec. 2017.
- [8] Y.-L. Park, B.-R. Chen, and R. J. Wood, "Design and fabrication of soft artificial skin using embedded microchannels and liquid conductors," *IEEE Sensors Journal*, vol. 12, no. 8, pp. 2711–2718, Aug. 2012.
- [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)*, Sept. 2015, pp. 2324–2329.
- [10] 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.
- [11] N. Farrow and N. Correll, "A soft pneumatic actuator that can sense grasp and touch," in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sept. 2015, pp. 2317–2323.
- [12] J. Hughes and F. Iida, "Localized differential sensing of soft deformable surfaces," in *Robotics and Automation (ICRA), 2017 IEEE International Conference on*. IEEE, 2017, pp. 4959–4964.
- [13] B. S. Homberg, R. K. Katzschmann, M. R. Dogar, and D. Rus, "Haptic identification of objects using a modular soft robotic gripper," in *International Conference on Intelligent Robots and Systems (IROS)*, 2015, pp. 1698–1705.
- [14] J. Morrow, H. S. Shin, C. Phillips-Grafflin, S. H. Jang, J. Torrey, R. Larkins, S. Dang, Y. L. Park, and D. Berenson, "Improving Soft Pneumatic Actuator fingers through integration of soft sensors, position and force control, and rigid fingernails," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, May 2016, pp. 5024–5031.
- [15] 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.
- [16] A. Sieverling, C. Eppner, F. Wolff, and O. Brock, "Interleaving motion in contact and in free space for planning under uncertainty," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017, pp. 4011–4017.
- [17] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.