

ARTIFICIAL INTELLIGENCE

Soft robot perception using embedded soft sensors and recurrent neural networks

Thomas George Thuruthel^{1*}†, Benjamin Shih^{2†}, Cecilia Laschi¹, Michael Thomas Tolley²

Recent work has begun to explore the design of biologically inspired soft robots composed of soft, stretchable materials for applications including the handling of delicate materials and safe interaction with humans. However, the solid-state sensors traditionally used in robotics are unable to capture the high-dimensional deformations of soft systems. Embedded soft resistive sensors have the potential to address this challenge. However, both the soft sensors—and the encasing dynamical system—often exhibit nonlinear time-variant behavior, which makes them difficult to model. In addition, the problems of sensor design, placement, and fabrication require a great deal of human input and previous knowledge. Drawing inspiration from the human perceptive system, we created a synthetic analog. Our synthetic system builds models using a redundant and unstructured sensor topology embedded in a soft actuator, a vision-based motion capture system for ground truth, and a general machine learning approach. This allows us to model an unknown soft actuated system. We demonstrate that the proposed approach is able to model the kinematics of a soft continuum actuator in real time while being robust to sensor nonlinearities and drift. In addition, we show how the same system can estimate the applied forces while interacting with external objects. The role of action in perception is also presented. This approach enables the development of force and deformation models for soft robotic systems, which can be useful for a variety of applications, including human-robot interaction, soft orthotics, and wearable robotics.

INTRODUCTION

Perception is an essential component of an intelligent autonomous system. It is one of the basic necessities for closed-loop control and representation of the environment. Robotic perception involves the kinematic estimation of the self, contact modeling, and mapping of the surroundings. With traditional rigid robotics, solutions to proprioception and tactile sensing involve highly specialized sensors precisely developed and arranged to ensure maximum state observability. This is feasible because of the availability of accurate models and reliable technological development. With the rise of soft robotics and the complexities involved with modeling and development of soft robots (1, 2), we are presented with new challenges in perception. The high dimensionality of soft robots and soft sensors complicates the selection of type, number, and placement of sensors. With the availability of analytical models, statistical metrics can be formulated for this problem (3). However, the modeling of soft sensors is challenged by inconsistencies in their manufacture and nonlinearities in their dynamics (4, 5, 6).

The development of technologies for sensing in soft robotics is a growing field with diverse potential solutions (5). There are subtle differences between each of these technological solutions that give each unique advantages depending on the required task. An ideal soft sensor must provide state information along the body of a soft system with minimal effect on the dynamics of the system. Embedded sensing is the most viable solution for strain, stress, contact, and roughness estimation (7, 8, 9, 10, 11, 12, 13). Unlike external sensing (e.g., vision), embedded soft strain sensors are not restricted by occlusion and coordinate transformation problems. In cases where the sensor must be embedded in the soft system, high omni-directional compli-

ance is required. Conductive nanocomposites are among the most commonly used materials for soft strain sensors (14, 15). However, the modeling of these embedded sensors is quite difficult because of high nonlinearities and creep (16, 17, 18), although precise manufacturing may help alleviate the latter (14). Another prominent strain sensor design has its basis in metals that are liquid at room temperature, encased in a nonconductive elastomer (19, 20). Although they do not exhibit notable creep characteristics, these sensors are difficult to manufacture and are susceptible to leakage. Higher accuracy can be obtained with fiber Bragg gratings (21), stretchable optical waveguides (12), and magnetic sensors (22). However, these options have reduced omni-directional compliance. For this work, we use strain sensors consisting of layers of polydimethylsiloxane (PDMS) impregnated with conductive carbon nanotubes (cPDMS). We chose this particular sensor design because of its ease of manufacture and its scalability in number. From the viewpoint of modeling, these sensors exhibit many of the nonlinearities and the creep phenomenon typically observed in other soft sensors. Hence, an approach viable for these sensors should be easily transferred to soft sensors with other designs (as discussed above).

Once a consistent sensor is embedded in a system of concern, the next step is to obtain meaningful information about the system states from the raw sensor readings. Unlike traditional sensing technologies, soft sensors conform to the structure of the surrounding dynamical system. Consequently, formulating kinematic and contact models based on these sensors requires an understanding of the sensor dynamics as well as the system dynamics. Because of their omni-directional compliance, these sensors could potentially have singular configurations (sensor values do not change at certain system configurations) and nonunique mappings (i.e., similar sensor readings for different system configurations). Furthermore, in the case of interactions with the surroundings, contact modeling is a highly complex mathematical problem, currently limited to theoretical studies (5). Because of the complexity in modeling, most work has adopted empirical or semi-analytical approaches. A purely analytical

¹The BioRobotics Institute, Scuola Superiore Sant'Anna, Pisa, Italy. ²Department of Mechanical and Aerospace Engineering, University of California, San Diego, San Diego, CA, USA.

*Corresponding author. Email: t.thuruthel@sssup.it

†These authors contributed equally to this work.

framework would require advancements in technologies for precise and repeatable manufacturing of the sensors, as well as the dynamical system in question. This work circumvents these challenges by providing a general framework for automatically generating these models experimentally using machine learning algorithms. Recent work has also begun to explore the viability of using learning-based approaches for model synthesis (23, 24).

One of the biggest challenges in perception for soft robotics is multimodal sensing, i.e., the capacity of a soft robot to perceive multiple physical parameters. One example is a soft manipulator that is able to simultaneously estimate its kinematic configuration and the external forces applied to it. Previous works have proposed complex manufacturing and sensor design solutions for multimodal sensing (25, 26, 27); however, these works are limited to proof of concept only and lack modeling processes of their embedded sensors. Moreover, they are designed based on previous knowledge about the type of contacts (typically only tip contacts), sensor properties, and applications. Many sensory modalities—like stress, strain, and pressure—can be theoretically observed with multiple strain sensors of the same kind (5). Another overlooked phenomenon is the contribution of action to perception, an additional source of information used by humans (28). Sensing in the presence of active (internal forces) and passive (due to external constraints) forces is a task that has not yet been investigated in the context of soft robotics.

A compelling solution to the soft perception problem can be found in our own sensory system. The human proprioceptive system, in contrast to the traditional approach in robotics, is characterized by a highly redundant, diverse, and unstructured sensor architecture (29). Either by design or because of fundamental limitations of biological systems, there is constant transformation in the sensory system caused by growth, damage, and fatigue. Hence, the neuronal modeling centers are constantly adapting using visual data as teaching signals and frames of reference (30, 31, 32).

Highly redundant, randomized, and scattered sensor distribution is a simple yet powerful solution to the problems of sensor placement and determination of number and type of sensors. For soft sensors with higher compliance than the soft bodies they are sensing, an injudicious number of sensors can be embedded in the body being monitored with minimal perturbations to the system's dynamics. Previous work has investigated this concept by randomly placing commercial bend sensors and selecting the best combination for static modeling using information theory (24). The challenge in modeling unknown sensor and system dynamics can be solved with any machine learning approach as long as reliable training signals are available for developing the model. These ideas have been explored for object identification (33, 34), bend angle prediction (35), tactile gesture recognition (36), force sensing (23), and localization (23).

The three main areas of interest in soft robot perception are concerned with the estimation of body kinematics, external wrenches (i.e., applied combinations of forces and torques), and contact point estimation. Because of the strong coupling between the kinematics and the statics of conventional soft robots, all of these problems are interconnected (37). The problem of kinematic estimation can be stated as follows: Given the current sensor deformation states $s_d(t)$ and the control input $\tau(t)$, the objective is to provide a model that predicts the position $y(t)$ of the system. The cardinality of the required sensor space increases with the number of contacts and the dynamic actuation range. For example, for a soft robot with a single actuated degree of freedom (DoF) and no contacts with the sur-

roundings, a single deformation sensor is sufficient for static modeling. In this case, even local strain information along the length of the robot is sufficient for full observability. Additional sensors may be required for dynamic modeling because passive DoFs can get excited during motion. Once the robot comes in contact with the surroundings, the kinematics of a soft robot itself changes. Consequently, additional sensors are required to detect contact and to update the kinematic model accordingly.

External force sensing is a diversified problem with varying complexities and challenges depending on the system design. Broadly, they can be divided into direct and indirect estimation methods. Direct force sensing refers to approaches where the sensor is directly placed at the area of contact (20). Hence, modeling direct force sensors becomes independent of the system in which they are embedded. However, this approach imposes restrictions on the type and placement of sensors. Indirect force sensing infers contact forces based on information transmitted along the soft system, an approach that is more flexible in the type of sensors and their placement. This type of force sensing that uses configuration-level information without force sensors located proximal to the end effector is commonly referred to as “intrinsic force sensing” (38). Static models that predict the external wrenches applied to a continuum robot given deformation sensor states s_d and tension sensor states s_t have been proposed for both estimation (39) and control (40). The sensible wrench space depends on the configuration of the system and the cardinality of the tension sensors (38). Khan *et al.* (41) showed that, with estimates of the compliance matrix, external forces could be measured indirectly using only strain sensors.

This paper describes the use of a bioinspired sensory architecture with a modeling recipe based on machine learning that can address many of the current challenges in soft robot perception (fig. S1). We demonstrate that this methodology can be used to perform model-free, real-time multimodal sensing. First, we demonstrate a kinematic state estimator that could detect external contacts and modify the kinematics accordingly. Second, we show how we could use the same sensor architecture for indirect external force sensing. Compared with the state of the art, we could relax numerous assumptions commonly made in previous efforts: (i) Our system has both active and passive elements, and the modeling was done in continuous time without assumptions of static equilibrium. (ii) We took into consideration the drift and hysteresis effects typically found in current soft sensors by representing our problem as a time sequence prediction problem. (iii) The whole system could be made “sensitive” without restrictions on the location and duration of contacts. Therefore, unlike previous works on direct force modeling system, we could develop a force-sensing module that could be trained at regions anywhere along the robot. We propose a simple fabrication, integration, and learning methodology for rapid prototyping. In addition, we demonstrate how redundancy in the sensory system not only helped multimodal sensing but also provided graceful degradation in response to sensor failure.

The next section presents the materials and methods used in this paper. The performance of the kinematic model is presented first in the results section. With three partially independent cPDMS strain sensors, we trained and tested the model for three different conditions: (i) free motion of the finger, (ii) external contact at the tip, and (iii) external contact at a fixed location along the continuum finger. The performance was benchmarked by applying the same learning approach to a finger embedded with commercial flex sensors in the place of the

soft cPDMS sensors. The subsequent section presents the results of the force estimation model. The experimental setup was not varied in this case except for adding a load cell to the external contact environment for obtaining the ground truth. The case of predicting forces applied at the tip is then presented. Last, we present simulation studies that investigate how the redundant architecture could be exploited by the learned network to be more robust to noise even in the case of the complete loss of some sensors.

MATERIALS AND METHODS

The proposed approach was validated primarily on a pneumatically actuated planar soft finger with three embedded soft resistive strain sensors (Fig. 1). The soft finger was composed of a series of channels and chambers surrounded by a soft elastomer. On pressurization, the finger deformed according to the internal stress distribution along the elastomer (42). A single pneumatic actuator was used for driving the finger. cPDMS, with a resistance that increases with strain, encased in a nonconductive elastomer, served as the soft strain sensor. The sensors were manually manufactured with varying lengths and implanted in the finger by randomly placing them roughly aligned with finger length during the curing process of the finger. The main human knowledge required during sensor placement was to ensure that the sensors were not placed in a location that would not strain during actuation (e.g., along the neutral axis of bending). For the training and real-time testing period, the actuators were commanded to random reference pressures varying every second. A low-level proportional derivative (PD) controller tracked the reference pressure independently (43). A motion capture system acted as our ground truth, tracking the motion of the tip of the finger during the training phase. For force modeling, a commercial, single-axis load cell provided the ground truth. A type of recurrent neural network called a long short-term memory (LSTM) network was used for learning the time series mapping because of its ability to train long time-lagged data (44). The reference pressure inputs and the current impedance values of the three sensors were the only inputs to the network, and the outputs were the Cartesian coordinates of the fingertip and the forces applied by the finger at the point of training.




Fig. 1. Soft actuator design. (A) Side view of the computer-aided design (CAD) of the soft actuator with infrared-reflective balls for tracking the motion of the tip. Embedded sensors are used to estimate the coordinates of the tip and the forces applied by the actuator when in contact. The plots we present in this paper describe the position of the marker at the tip relative to the marker at the base. (B) Physical actuator with embedded soft sensors.

The LSTM network performs a mapping from $[s_d(t), \tau(t), c(t)] \rightarrow \mathbf{y}(t)$ for kinematic estimation and a mapping from $[s_d(t), \tau(t), c(t)] \rightarrow \mathbf{F}(t)$ for force estimation. Here, $s_d(t)$ is the sensor impedance readings consisting of the sensor resistance and reactance, $\tau(t)$ is the input pressure to the actuator, and $c(t)$ is the current state of the LSTM network. $\mathbf{p}(t)$ is the kinematic parameter to be estimated (tip position for our case), and $\mathbf{F}(t)$ is the estimated force. Note that input components are the same for both the kinematic and force estimation. Therefore, it is possible to perform both kinematic and force sensing at the same time.

Fabrication of the sensors

Biological skin contains a dense, distributed network of a large variety of interspersed sensors (29). In our system, we mimicked the biological anatomy with three soft cPDMS sensors. The sensors were placed arbitrarily along the finger length when the finger was being cured (fig. S7).

We made the soft sensors from patterned cPDMS traces (fig. S6). First, we dispersed multiwall carbon nanotubes (MWCNTs) (30 to 50 nm diameter, Cheap Tubes Inc.) in PDMS base (Sylgard 184, 3M) to achieve 14% MWCNT loading by mass by mixing in a speed mixer for 20 min. We poured a thin layer of silicone elastomer (Dragon Skin 10, Smooth-On Inc.) onto a rigid substrate to form the lower layer of the sensing skin. We thoroughly coated the cPDMS onto the surface of the base silicone layer to ensure that there were no gaps. Then, we placed silicone-insulated wires with exposed leads onto the uncured cPDMS and poured a second layer of silicone elastomer on top to seal the conductive material inside. The cPDMS layer and the second layer of silicone elastomer were then cured at room temperature for about 6 hours.

From this sheet of cPDMS sensors, we then cut out the desired geometry of the sensor manually. The resulting thickness was about 3 mm. The ability to learn the material properties enabled us to trade precision in the fabrication for time and to pay less attention to uniform thickness or orientation of the sensors within the finger. This makes the sensor fabrication process more general and scalable than previous approaches to fabricating cPDMS sensors (34) because we can avoid the step of sensor masking.

In Results, the subsections on kinematic modeling, force modeling, and graceful degradation were done with a 1-DoF actuator that has a single pneumatic chamber. The section on multi-DoF system was done with both a 2-DoF parallel actuator (fig. S8A) and a 2-DoF serial actuator (fig. S8B), both of which had two pneumatic chambers but inflated in different ways due to the alignment of the chambers. In the 1-DoF actuator and the 2-DoF serial actuator, all three sensors had dimensions of about 110 mm by 25 mm by 3 mm, with intentional, minor differences in each sensor. In the 2-DoF actuator, the three sensors had dimensions of 70 mm by 20 mm by 4 mm, 120 mm by 26 mm by 4 mm, and 103 mm by 27 mm by 4 mm.

Fabrication of the actuators

We fabricated the actuators by casting silicone (Dragon Skin 20, Smooth-On Inc.) with three-dimensional-printed molds (VeroClear, Stratasys Objet350 Connex3) (34). After molding the chambers of the actuator, we embedded and bonded the sensors to the bottom of the actuator and sealed off the chambers by submerging in additional silicone. With the chambers fully sealed, we then inserted silicone tubing to enable inflation of the chambers. The wires for the sensors were then soldered to a printed circuit board, which provided mechanical

stability for the connections and header pins to connect the electronics. The overall dimensions of the actuator are 120 mm by 35 mm by 25 mm.

Experimental setup

Our analog to the visual feedback used by many animals is a motion capture system (OptiTrack Prime 13, NaturalPoint Inc.). We used two OptiTrack cameras to provide ground truth validation for the position of the finger in the real world. We placed infrared tracking balls for the motion capture system at the base and tip of our pneumatic actuator with embedded sensors. All the coordinates of the fingertip were measured with respect to the base coordinates. We mounted the actuator onto the front of a metal stand (80/20, McMaster-Carr) such that two reflective markers affixed to the base and tip of the actuator were visible from the front (Fig. 1). For the force sampling, we used a compression load cell (FX1901, TE Connectivity).

Then, we connected the finger to a volumetric control system designed to apply commanded pressures to the internal chamber of the actuator (43). A randomly generated sequence of pressure inputs was sent serially to the control board, which regulated the pressure inside the finger using a low-level PD controller running at 1000 Hz. The pseudorandom sequence was in the form of square wave with a range varying from 0 to 3.5 bars and a time period of 1 s. The learned model used these reference pressure values along the actual sensor readings for kinematics and force predictions. We assumed that the actual pressure values inside the finger were similar to the reference value. For real-time testing of the learned model, a new random sequence was sent to the volumetric control system.

The sensors were connected to an LCR meter (Keysight, E4980AL), which provided high-precision measurements. The measurements were made using AC signals of 300 kHz. Both the resistance and reactance of the sensor were measured and used for prediction. Because the meter only had a single input, we first connected the embedded sensors to a multiplexer circuit. The three sensor measurements were obtained at 10 Hz. The marker coordinates and force readings were also resampled to 10 Hz. A schematic of the whole setup is shown in fig. S9.

Sampling for kinematics and force modeling

The samples for learning the kinematic model and the force model were obtained with the same setup. For the kinematic model, the marker information from the motion capture system and the corresponding sensor data were required for different kinematic configurations. To obtain this, the finger was occasionally brought in contact with a fixed line contact at two different locations (see fig. S10). Because the external contact was fixed, the finger was still completely free to move in the other direction. The contacts were designed to touch the finger at the tip and at a point near the center of the continuum finger. The timing, duration, and location of the external contact were randomized to avoid biases. The sampling was continuous, and the data were not shuffled for learning. This is important to keep the temporal information intact. For force modeling, the external contact at the tip was integrated with a load cell (see fig. S11).

LSTM for nonlinear, time-varying material characteristics

LSTM networks are a class of recurrent neural networks widely used for time series predictions (44). We used the LSTM network provided by the MATLAB deep learning toolbox for creating our network. For all the trained networks, for both the cPDMS sensor and the commercial flex sensor, we used the same network parameters. An LSTM layer size of 100 was taken with a dropout layer preceding the

LSTM layer. The dropout rate was kept high at 0.5 for the graceful degradation test and at 0.1 for all the other tests. A fully connected layer that multiplied the output of the LSTM layer by a weight matrix and then added a bias vector provided the final output from the network. L_2 regularization was also used to prevent overfitting. The training data were normalized and split into two continuous blocks in the ratio 80:20 for training and testing. The network parameters were optimized using the Adam algorithm (45). The mini batch size was 512.

RESULTS

Kinematic modeling

To demonstrate the potential of the proposed methodology for full-body kinematic estimation, we performed a fundamental test from which scalability was evident. The test involved the finger following a random actuation pattern while being obstructed at unknown times by two fixed-point contacts. The height of the obstacles was fixed, but their placement along the x axis varied. One contact was enforced at the tip of the finger and the other at an arbitrary location along the length of the finger. The same experiment was repeated with a finger where the soft cPDMS sensor was replaced by a commercial flex sensor. No adjustment to the learning approach was required in this case because it is agnostic to the type of sensors. Kinematic estimation with the flex sensor was more accurate for the no-contact case because of the absence of any temporal nonlinearities (fig. S2B), but the high axial stiffness of the flex sensor reduced the effective compliance of the finger, thereby reducing the reachable workspace of the finger (fig. S2A).

The training performance of the LSTM network is shown in table S1, and the real-time test performance is shown in Fig. 2. The number of samples required for the cPDMS sensor was higher than for the flex sensor. The reason higher samples were required for the cPDMS sensor is solely because drift is a slow dynamic process. Because normalization of the data was performed using the sample data, it is necessary to make sure that the sensor reached the boundaries of working limit. Otherwise, during the testing phase, the learned network would get saturated when the sensor readings went outside the sample ranges. The sampling rate was 10 Hz, so the whole sampling period for training lasted about 50 min. As expected, the prediction using the flex sensor was more accurate even during the training and the real-time testing phase without contact. However, the prediction performance deteriorated upon contact. The soft sensors, on the other hand, performed consistently for all three cases. The trajectory of the fingertip and the predicted positions




Fig. 2. Real-time performance.

are shown in Fig. 3B. A notable characteristic of the cPDMS sensor was the slight phase lag of the predictions. This could be due to the slower dynamics of the soft sensor (18) compared with the dynamics of the finger itself. Therefore, kinematic information from the body would become observable through the sensor only after a delay. This phase delay was not observed in the stiffer commercial flex sensor. The real-time test results were measured for a period of 20 s for each scenario. The error plots for both the test are shown in Fig. 4. For scaling the current setup to accommodate more points of contact, we would need to embed more sensors and devise more training scenarios.

Even after being constrained at the tip from one side, the soft cPDMS sensors still responded to actuation inputs, because of internal stress induced by the pneumatic pressure. This was not evident with the flex sensors because of the increased axial rigidity of the finger itself (fig. S3). Similarly, the independence of the three flex sensors was affected because of this unresponsiveness to contact (fig. S4). Thus, the predictions with the flex sensor were more prone to errors when in contact.

Force modeling

The force prediction model was learned by using the same methodology, but we replaced the position signals with the forces applied by the tip of the finger. The tip forces were measured in the x -axis direction (i.e., parallel to the direction of travel of the fingertip in its resting state) using a single-axis load cell, and 9500 samples were obtained for training. The inputs to the LSTM network remained the same as the kinematic model. The sampling rate was also kept the same at 10 Hz. The average force prediction error for the first 40 s of the real-time test was found to be 15.3% with respect to the total range. The

prediction and error plot of the same test is shown in Fig. 5. An additional uncalibrated test with a human hand was performed to ensure that the learning was not specific to the setup. The learned model performed with an average error of 0.05 ± 0.06 N in estimating the magnitude of error and detecting the onset of contact (implicitly); however, the system had exhibited a delay in detecting the cessation of contact. Similar phase lags observed in the kinematic estimator were also observed for this case.

Movie S8 shows that the setup could theoretically be used for force prediction anywhere along the finger but did not measure the exact magnitude of the applied force at every location with the current number of sensors. The actuator can respond to contact anywhere along the body because a single sensor runs along the body. This is one of the advantages of indirect force sensing. Direct force sensors, on the other hand, would not be responsive to contact in other locations. With more sensors, the methodology could be extended to estimate the exact magnitudes along the arm. Note that we do not need to explicitly mention the location of the contact for learning. Adding the fact that learning algorithms have very good generalization ability, we suspect that a few cases of contact would be sufficient to develop an approximate full-body force estimate.

Graceful degradation

Biological systems typically exhibit redundancies in their sensing modalities, which allow the organism to function despite damages to subcomponents of the system. This concept of graceful degradation suggests that we can use redundancy in the soft sensor network to maintain functional performance despite damage to the individual sensors. For the task of predicting the position of the tip of the finger




Fig. 3. Predicted motions of the fingertips. (A) With the cPDMS sensors. The case of applying contact around the center of the finger is shown. The tip was still free to move after the constraint was applied, but the kinematics changed. (B) With the cPDMS sensors. The case of applying contact around at the tip of the finger is shown. (C) With the flex sensor. Both cases of contact—one at the tip and the other near the center of the finger—are shown. The first constraint was at the tip, and the second constraint was near the center of the finger.

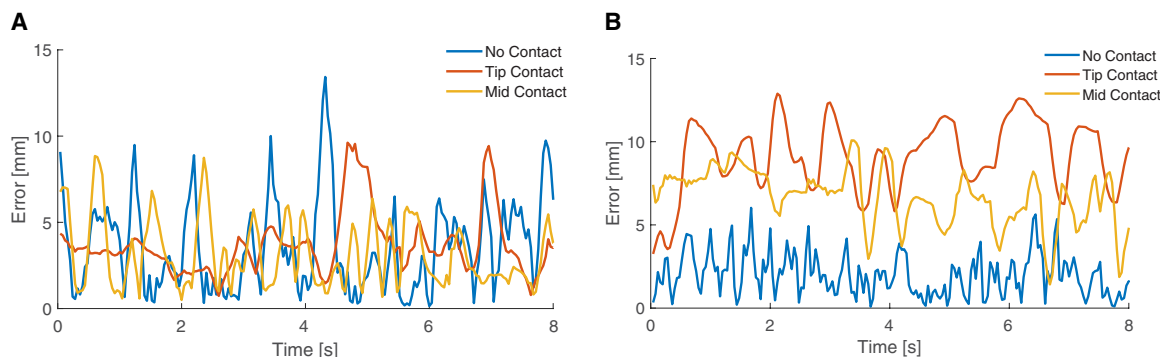


Fig. 4. Error plots for tracking. (A) With the soft cPDMS sensor. (B) With the commercial flex sensor.




Fig. 5. Force prediction at the fingertip. The raw load cell readings are filtered with a simple moving average filter with a 1-s window. External hand contact without the load cell is also shown.

without contact, our sensory architecture was redundant. Therefore, with appropriate training, the learned model could be made more robust to the loss of sensory information. We achieved this by increasing the dropout rate during training while using a small LSTM network. We expected, however, that this would reduce the accuracy of the model.

Here, we show the results of fundamental tests to observe the predictive power of the pretrained network in the face of abrupt loss of sensory information. All results were obtained by using the training data itself. For practical reasons, we were unable to physically remove sensors from the setup for real-time testing. To virtually simulate the removal of sensors, we set each row of our inputs to zero during the simulated test phase. The loss in accuracy in response to sensor removal is shown in Fig. 6. For both the cPDMS sensor and the flex sensor, a gradual decrease in accuracy could be observed upon virtual removal of each sensor and each combination of sensors. For the cPDMS sensors, each of the sensors appeared to contribute equally to the predictions. This could also be observed from the error distribution in the workspace (Fig. 7A). The pressure information played a vital role in prediction for the cPDMS sensors. This is because the pressure information played a vital role in compensating for sensor drift. A model trained with a short training sample without pressure information showed notable drift in the test phase, whereas a model trained with the same data and pressure information was able to compensate for the sensor drift (fig. S5). Small vari-

ations in the performance of the model among sensors can be attributed to their signal-to-noise ratio and the variabilities in training. The same comparison for the contact scenario led to drastic performance degradation and clear error distributions in the workspace, indicating how each sensor contributed uniquely for different tasks (Fig. 7B).

Multi-DoF system

To validate the performance of the sensing methodology on a more complex system, we replicated the tests on two 2-DoF actuators. The 2-DoF parallel actuator (fig. S8A) is a nonplanar 2-DoF actuator, and the 2-DoF serial actuator (fig. S8B) is a planar 2-DoF actuator.

The methodology remained the same, with identical actuator geometry for better comparison with the 1-DoF system. We demonstrated that this methodology worked for different sensor geometries by also testing variable sizes of sensors on the 2-DoF parallel actuator. Ten thousand samples were obtained for each of the actuators for the kinematic estimation. The data were divided in the ratio 80:20 for training and testing, and validation performances were measured on the test set. The nonplanar actuator had a prediction accuracy of 2.43 ± 2.18 mm, whereas the planar 2-DoF actuator had a prediction accuracy of 2.17 ± 1.65 mm.

DISCUSSION

This paper presents a generalizable, model-less technique for real-time perception for a soft actuator using embedded soft sensors and recurrent neural networks. We followed a bioinspired approach for both hardware and software components. This allowed us to achieve an accurate kinematic model of a soft finger even with highly nonlinear sensors. Although more accurate predictions were obtained with commercial flex sensors, their relative rigidity and inextensibility made them undesirable for high-dimensional deformations. With our proposed methodology, we demonstrated how full-body kinematic models could be learned. In addition, by following the same methodology, the system learned models of externally applied forces using the stress-strain relationship of the soft body. We validated the approach for a fundamental test using a fabrication approach

for which scaling to large numbers of sensors could be assured. We were able to accomplish this with irregularly shaped strain sensors. Because of the continuous distribution of the sensing module and the learning process, we could easily adjust the location of sensing. Because of our reliance on a pure learning-based model, it was

possible to fuse the measured information from the sensors with the commanded information to the pressure regulators to achieve more accurate models. The role of action in perception is also a phenomenon observed in biological systems (29). The methodology is highly generalizable with the ability to interchange sensors, mode of sensing, and the system itself without any changes to the learning algorithm. Last, we explored how sensor redundancy can make the system more robust to unexpected changes to the system.

Learning-based approaches are very useful for modeling with minimal knowledge about the system; however, they also have inherent drawbacks. For example, these approaches do not afford the designer any physical intuition about the system. Thus, further analysis—for example, to determine the optimal shape, placement, and number of sensors—would be difficult, as would describing and correcting for sources of error. For this study, we used LSTM networks for training our time-dependent model. Although it is possible to train the same model with feed-forward neural networks using appended input sequences, this makes the model memory inefficient (46). Moreover, there is an additional step where the user has to hand-tune the fixed-input time lags. Compared with other recurrent neural network architectures, LSTM was chosen mainly because of the ease in training LSTM networks for long time-lag tasks (44). However, with the advancements in training recurrent neural networks, it would be meaningful to investigate other recurrent architectures.




Fig. 6. Test accuracy with virtual sensor removal. The performance is affected only a little when one of the sensor information is lost for the “no contact” case. The accuracy is considerably affected even with one sensor removal for the “with contact” case.

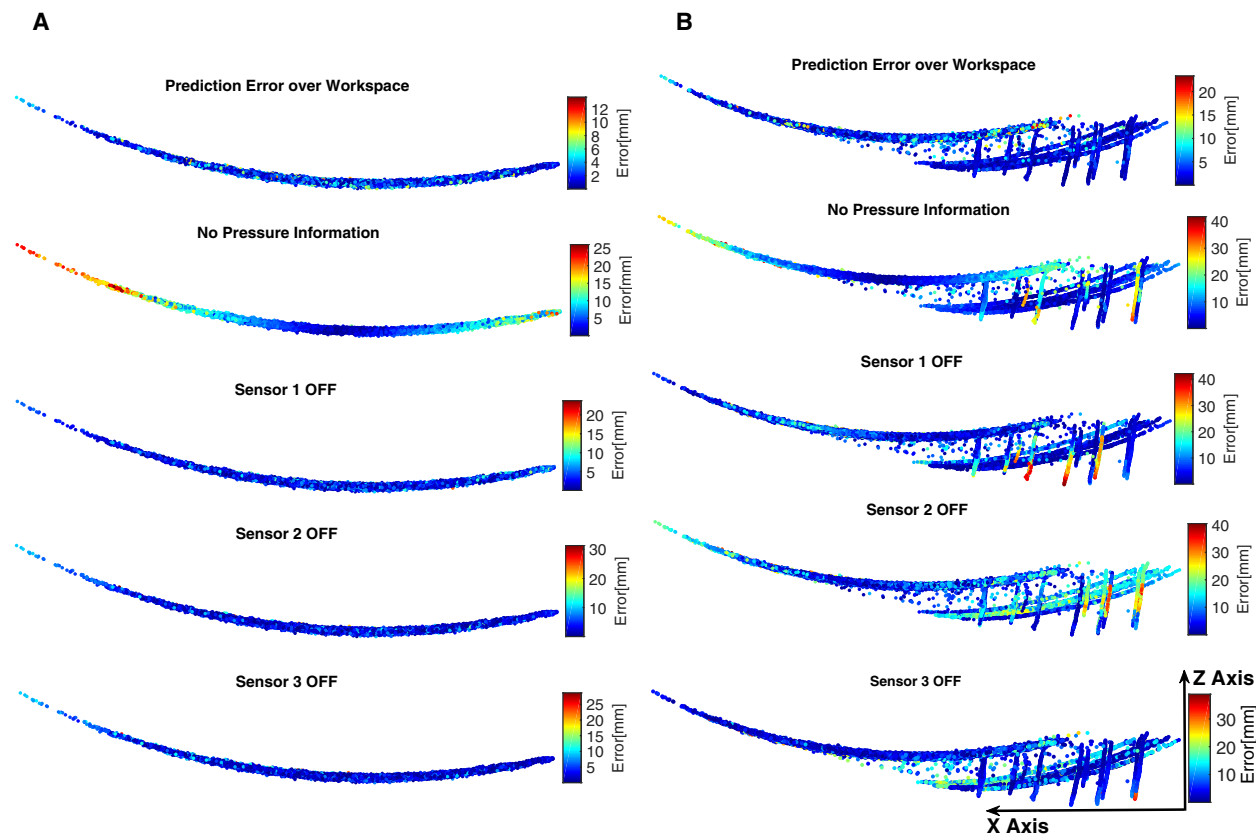


Fig. 7. Divisions of labor among the sensors. (A) For the case without contact, all the sensors have equal contribution to the underlying model. Hence, removing any one of them affects the prediction error slightly but equally in the workspace. For this case, removing the pressure information drastically reduces the accuracy, showing how motor action information is also important for accurate proprioception. (B) Division of labor among the sensors once in contact. Here, we can see division of labor among the sensors because there are no redundant sensors. Each sensor is “specialized” to a particular kinematic case as can be seen from the error distribution in the workspace.

The approach described here is demonstrated on a planar soft finger with only three embedded strain sensors. The main bottleneck for scaling to a larger number of sensors was the serial nature of the multiplexer circuit that we used to read the sensor signals from a single, high-precision inductance-capacitance-resistance (LCR) meter. This also led to signal mixing because of ghosting (or cross-talk) effects that became more prevalent at higher sampling frequency. Because of this, our sensing system had a maximum sampling rate of 10 Hz. Although 10 Hz was sufficient for this iteration of our experimental tasks, it will be insufficient for more dynamic tasks or actuators with additional DoFs. This issue could be solved with improved data acquisition methods with multiple analog-to-digital channels.

A source of error for the cPDMS sensor was the lag of the predictions behind the actual values. The lag was not observed with the commercial flex sensor, which means that it can be attributed to the slower dynamics of the soft cPDMS sensor (18). Nevertheless, this would affect the applicability of using embedded sensory feedback for highly dynamic tasks. Devoid of the visual feedback system, the human proprioceptive system is also susceptible to erroneous drifts (47) and biases (48). An important question to be investigated is whether the slow sensory response is intrinsic to all soft sensors or whether it is due to the internal mechanics of our cPDMS sensor. A potential avenue toward improvement of the intrinsic errors in our cPDMS sensors could be to experiment with non-polymer-based materials, which may exhibit less notable time lags compared with cPDMS because of the molecular structure of the conductive material. However, alternative sensor materials have their own challenges, as discussed in Introduction.

Although our complete methodology closely resembles the human perceptive system, our reference feedback loops were well structured when compared with the biological counterparts. The tracking system and the load cell that we used as ground truths provided physically meaningful outputs that could be easily learned in a supervisory manner using the LSTM network. However, in the human perceptive system, we map our sensor signals with multimodal raw signals coming from the ocular, vestibular, auditory, and muscular systems. For simplicity, we preprocess the raw images and force readings coming from the reference systems to obtain physically relevant variables. A faithful end-to-end replication of the human perceptive system, on the other hand, would require a direct mapping from the sensory space to the image space. This introduces additional complexities in the form of object recognition, calibration, and coordinate referencing. The constant presence of the visual, inertial, and auditory feedback is also important for adaptation in case of drastic system changes (29, 49). Our methodology currently relies on independent external sensing technologies for reference feedback, which have to be removed for real-world applications. However, if the entire system undergoes permanent physical changes—like growth, stiffening, and material deterioration—then the learned model would display biases. Hence, a potential future endeavor would be to integrate other sensing modalities—like vision, inertial measurement units, and force sensors—directly into the soft system.

SUPPLEMENTARY MATERIALS

robotics.sciencemag.org/cgi/content/full/4/26/eaav1488/DC1

Table S1. Training performance of the kinematic model.

Fig. S1. Overview of the modeling architecture and its parallel to the human perceptive system.

Fig. S2. Differences between a commercial flex sensor and the cPDMS sensor.

Fig. S3. Sensor response to tip contact.

Fig. S4. Intersensor dependencies.

Fig. S5. Contribution of pressure information for drift compensation.

Fig. S6. Schematic of sensor fabrication process.

Fig. S7. Sensor topology.

Fig. S8. Schematic of the motion of the 2-DoF actuators.

Fig. S9. Schematic of the experimental setup.

Fig. S10. Diagram showing how contact along the continuum of the actuator results in a deformation that propagates throughout the system.

Fig. S11. Diagram of how we obtain the force measurement at the tip of the actuator using a load cell.

Movie S1. Kinematic prediction with the cPDMS sensors—without contact.

Movie S2. Kinematic prediction with the cPDMS sensors—contact at tip.

Movie S3. Kinematic prediction with the cPDMS sensors—contact along the finger.

Movie S4. Kinematic prediction with the commercial flex sensors—without contact.

Movie S5. Kinematic prediction with the commercial flex sensors—contact at tip.

Movie S6. Kinematic prediction with the commercial flex sensors—contact along the finger.

Movie S7. Force sensing experiment with the cPDMS sensors.

Movie S8. Experiment with the cPDMS sensors showing that the same learned model is sensitive to contact anywhere along the arm.

REFERENCES AND NOTES

1. D. Rus, M. T. Tolley, Design, fabrication and control of soft robots. *Nature* **521**, 467–475 (2015).
2. C. Laschi, B. Mazzolai, M. Cianchetti, Soft robotics: Technologies and systems pushing the boundaries of robot abilities. *Sci. Robot.* **1**, eaah3690 (2016).
3. A. W. Mahoney, T. L. Bruns, P. J. Swaney, R. J. Webster, On the inseparable nature of sensor selection, sensor placement, and state estimation for continuum robots or “where to put your sensors and how to use them”, in *2016 IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2016), pp. 4472–4478.
4. M. Amjadi, K.-U. Kyung, I. Park, M. Sitti, Stretchable, skin-mountable, and wearable strain sensors and their potential applications: A review. *Adv. Funct. Mater.* **26**, 1678–1698 (2016).
5. H. Wang, M. Totaro, L. Beccai, Toward perceptive soft robots: Progress and challenges. *Adv. Sci.* **5**, 1800541 (2018).
6. P. Polygerinos, N. Correll, S. A. Morin, B. Mosadegh, C. D. Onal, K. Petersen, M. Cianchetti, M. T. Tolley, R. F. Shepherd, Soft robotics: Review of fluid-driven intrinsically soft devices; manufacturing, sensing, control, and applications in human-robot interaction. *Adv. Eng. Mater.* **19**, 1700016 (2017).
7. N. Wettels, V. J. Santos, R. S. Johansson, G. E. Loeb, Biomimetic tactile sensor array. *Adv. Robot.* **22**, 829–849 (2008).
8. I. M. Van Meerbeek, C. M. De Sa, R. F. Shepherd, Soft optoelectronic sensory foams with proprioception. *Sci. Robot.* **3**, eaau2489 (2018).
9. Y. Mengüç, Y.-L. Park, H. Pei, D. Vogt, P. M. Aubin, E. Winchell, L. Fluke, L. Stirling, R. J. Wood, C. J. Walsh, Wearable soft sensing suit for human gait measurement. *Int. J. Robot. Res.* **33**, 1748–1764 (2014).
10. D. Kim, J. Kwon, S. Han, Y.-L. Park, S. Jo, Deep full-body motion network (DFM-Net) for a soft wearable motion sensing suit. *IEEE/ASME Trans. Mechatron.* **1** (2018).
11. R. Patel, R. Cox, N. Correll, Integrated proximity, contact and force sensing using elastomer-embedded commodity proximity sensors. *Autonom. Robot.* **42**, 1–16 (2018).
12. H. Zhao, K. O'Brien, S. Li, R. F. Shepherd, Optoelectronically innervated soft prosthetic hand via stretchable optical waveguides. *Sci. Robot.* **1**, eaai7529 (2016).
13. R. D. P. Wong, R. B. Hellman, V. J. Santos, *Next-Generation Robots and Systems* (International Society for Optics and Photonics, 2014), vol. 9116, 911605 pp.
14. T. Yamada, Y. Hayamizu, Y. Yamamoto, Y. Yomogida, A. Izadi-Najafabadi, D. N. Futaba, K. Hata, A stretchable carbon nanotube strain sensor for human-motion detection. *Nat. Nanotechnol.* **6**, 296–301 (2011).
15. C. Mattmann, F. Clemens, G. Tröster, Sensor for measuring strain in textile. *Sensors* **8**, 3719–3732 (2008).
16. L. Wang, F. Ma, Q. Shi, H. Liu, X. Wang, Study on compressive resistance creep and recovery of flexible pressure sensitive material based on carbon black filled silicone rubber composite. *Sens. Actuat. A Phys.* **165**, 207–215 (2011).
17. Q. Zheng, J. Zhou, Y. H. Song, Time-dependent uniaxial piezoresistive behavior of high-density polyethylene/short carbon fiber conductive composites. *J. Mater. Res.* **19**, 2625–2634 (2004).
18. J. T. Muth, D. M. Vogt, R. L. Truby, Y. Mengüç, D. B. Kolesky, R. J. Wood, J. A. Lewis, Embedded 3D printing of strain sensors within highly stretchable elastomers. *Adv. Mater.* **26**, 6307–6312 (2014).
19. Y.-L. Park, C. Majidi, R. Kramer, P. Bérard, R. J. Wood, Hyperelastic pressure sensing with a liquid-embedded elastomer. *J. Micromech. Microeng.* **20**, 125029 (2010).
20. D. M. Vogt, Y.-L. Park, R. J. Wood, Design and characterization of a soft multi-axis force sensor using embedded microfluidic channels. *IEEE Sens. J.* **13**, 4056–4064 (2013).

21. B. A. L. Gwandu, X. W. Shu, Y. Liu, W. Zhang, L. Zhang, I. Bennion, Simultaneous measurement of strain and curvature using superstructure fibre Bragg gratings. *Sens. Actuat. A Phys.* **96**, 133–139 (2002).
22. S. Ozel, N. A. Keskin, D. Khea, C. D. Onal, A precise embedded curvature sensor module for soft-bodied robots. *Sens. Actuat. A Phys.* **236**, 349–356 (2015).
23. S. Han, T. Kim, D. Kim, Y.-L. Park, S. Jo, Use of deep learning for characterization of microfluidic soft sensors. *IEEE Robot. Autom. Lett.* **3**, 873–880 (2018).
24. V. Wall, G. Zöllner, O. Brock, A method for sensorizing soft actuators and its application to the RBO hand 2, in *2017 IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2017), pp. 4965–4970.
25. R. L. Truby, M. Wehner, A. K. Grosskopf, D. M. Vogt, S. G. M. Uzel, R. J. Wood, J. A. Lewis, Soft somatosensitive actuators via embedded 3D printing. *Adv. Mater.* **30**, 1706383 (2018).
26. J. Morrow, H.-S. Shin, C. Phillips-Grafflin, S.-H. Jang, J. Torrey, R. Larkins, S. Dang, Y.-L. Park, D. Berensen, 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)* (2016), pp. 5024–5031.
27. B. Shih, J. Mayeda, Z. Huo, C. Christianson, M. T. Tolley, 3D printed resistive soft sensors, in *2018 IEEE International Conference on Soft Robotics (RoboSoft)* (IEEE, 2018).
28. N. Sornkarn, T. Nanayakkara, Can a soft robotic probe use stiffness control like a human finger to improve efficacy of haptic perception? *IEEE Trans. Haptics* **10**, 183–195 (2017).
29. U. Proske, S. C. Gandevia, The proprioceptive senses: Their roles in signaling body shape, body position and movement, and muscle force. *Physiol. Rev.* **92**, 1651–1697 (2012).
30. K. Fiehler, F. Rösler, Plasticity of multisensory dorsal stream functions: Evidence from congenitally blind and sighted adults. *Restor. Neurol. Neurosci.* **28**, 193–205 (2010).
31. G. Cappagli, E. Cocchi, M. Gori, Auditory and proprioceptive spatial impairments in blind children and adults. *Dev. Sci.* **20**, e12374 (2017).
32. R. J. van Beers, A. C. Sittig, J. J. van der Gon Denier, How humans combine simultaneous proprioceptive and visual position information. *Exp. Brain Res.* **111**, 253–261 (1996).
33. B. S. Homberg, R. K. Katschmann, M. R. Dogar, D. Rus, Haptic identification of objects using a modular soft robotic gripper, in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE, 2015), pp. 1698–1705.
34. B. Shih, D. Drotman, C. Christianson, Z. Huo, R. White, H. I. Christensen, M. T. Tolley, Custom soft robotic gripper sensor skins for haptic object visualization, in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE, 2017), pp. 494–501.
35. K. Elgeneidy, N. Lohse, M. Jackson, Bending angle prediction and control of soft pneumatic actuators with embedded flex sensors—A data-driven approach. *Mechatronics* **50**, 234–247 (2018).
36. C. Larson, J. Spjut, R. Knepper, R. Shepherd, A deformable interface for human touch recognition using stretchable carbon nanotube dielectric elastomer sensors and deep neural networks. arXiv:1706.02542 [cs.HC] (8 June 2017).
37. T. George Thuruthel, Y. Ansari, E. Falotico, C. Laschi, Control strategies for soft robotic manipulators: A survey. *Soft Robot.* **5**, 149–163 (2018).
38. K. Xu, N. Simaan, An investigation of the intrinsic force sensing capabilities of continuum robots. *IEEE Trans. Robot.* **24**, 576–587 (2008).
39. A. Bajo, N. Simaan, Finding lost wrenches: Using continuum robots for contact detection and estimation of contact location, in *2010 IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2010), pp. 3666–3673.
40. R. E. Goldman, A. Bajo, N. Simaan, Configuration and joint feedback for enhanced performance of multi-segment continuum robots, in *2011 IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2011), pp. 1126–1132.
41. F. Khan, R. J. Roesthuis, S. Misra, Force sensing in continuum manipulators using fiber Bragg grating sensors, in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (2017), pp. 2531–2536.
42. R. F. Shepherd, F. Ilievski, W. Choi, S. A. Morin, A. A. Stokes, A. D. Mazzeo, X. Chen, M. Wang, G. M. Whitesides, Multigait soft robot. *Proc. Natl. Acad. Sci. U.S.A.* **108**, 20400–20403 (2011).
43. T. Kalisky, Y. Wang, B. Shih, D. Drotman, S. Jadhav, E. Aronof-Spencer, M.T. Tolley, Differential pressure control of 3D printed soft fluidic actuators, in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE, 2017), pp. 6207–6213.
44. S. Hochreiter, J. Schmidhuber, Long short-term memory. *Neural Comput.* **9**, 1735–1780 (1997).
45. D. P. Kingma, J. Ba, Adam: A method for stochastic optimization. arXiv:1412.6980 [cs.LG] (22 December 2014).
46. R. J. Williams, D. Zipser, A learning algorithm for continually running fully recurrent neural networks. *Neural Comput.* **1**, 270–280 (1989).
47. A. Tsay, G. Savage, T. J. Allen, U. Proske, Limb position sense, proprioceptive drift and muscle thixotropy at the human elbow joint. *J. Physiol.* **592**, 2679–2694 (2014).
48. R. J. van Beers, A. C. Sittig, J. J. Denier van der Gon, The precision of proprioceptive position sense. *Exp. Brain Res.* **122**, 367–377 (1998).
49. P. B. Matthews, Mammalian muscle receptors and their central actions. *Am. J. Phys. Med. Rehabil.* **53**, 143 (1974).

Acknowledgments: We thank the members of the Bioinspired Robotics and Design Lab for helpful discussions, T. Bewley for loaning us his motion capture system, Z. Huo for constructing the test stand, M. Yip for helpful discussions, and K. Morris and E. Lathrop for loaning us data acquisition devices. **Funding:** This work was supported by the Office of Naval Research (grant number N000141712062). **Author contributions:** T.G.T., B.S., and M.T.T. conceptualized the experiment, developed the methodology, validated the results, wrote the software, and wrote the manuscript. All authors contributed to the figures. B.S., C.L., and M.T.T. supervised and administered the project. **Competing interests:** The authors declare that they have no competing interests. **Data and materials availability:** All data needed to evaluate the conclusions in the paper are present in the paper or the Supplementary Materials. Other data for this study can be found in the database (<https://github.com/mlsensor/SciRo>).

Submitted 4 September 2018

Accepted 21 December 2018

Published 30 January 2019

10.1126/scirobotics.aav1488

Citation: T. George Thuruthel, B. Shih, C. Laschi, M. T. Tolley, Soft robot perception using embedded soft sensors and recurrent neural networks. *Sci. Robot.* **4**, eaav1488 (2019).

Soft robot perception using embedded soft sensors and recurrent neural networks

Thomas George Thuruthel Benjamin Shih Cecilia Laschi Michael Thomas Tolley

Sci. Robot., 4 (26), eaav1488. • DOI: 10.1126/scirobotics.aav1488

View the article online

<https://www.science.org/doi/10.1126/scirobotics.aav1488>

Permissions

<https://www.science.org/help/reprints-and-permissions>

Use of this article is subject to the [Terms of service](#)