Embedded air channels transform soft lattices into sensorized grippers

Annan Zhang*, Lillian Chin*, Daniel L. Tong, Daniela Rus

***Abstract*— Sensing plays a pivotal role in robotic manipulation, dictating the accuracy and versatility with which objects are handled. Vision-based sensing methods often suffer from fabrication complexity and low durability, while approaches that rely on direct measurements on the gripper often have limited resolution and are difficult to scale. Here, we present a soft robotic gripper made out of two cubic lattices that are sensorized by embedding air channels within the structure. The lattices are 3D printed from a single build material, simplifying the fabrication process. The flexibility of this approach offers significant control over sensor and lattice design, while the pressure-based internal sensing provides measurements with minimal disruption to the grasping surface. With only 12 sensors, 6 per lattice, this gripper can estimate an object's weight and location and offer new insights into grasp parameters like friction coefficients and grasp force.**

I. Introduction

Sensing in manipulation is critical to understanding what is being grasped. Traditionally, manipulation tasks have relied on external vision systems to report where objects are and plan accordingly [1]. However, there are significant issues with vision-only approaches. Occlusion often occurs in manipulation scenarios, especially during and after the grasping process. Furthermore, some properties are difficult to verify visually, such as the weight or stiffness of an object. Adding sensing within a robot's hand can address these issues. Whether by putting a camera in the hand itself [2], [3], measuring actuation effort [4], or adding force and strain sensors to the finger [5], in-hand sensing solutions allow for direct measurement of the grasping process. While the extra information is useful, these systems add significant complexity to the fabrication of the grippers while simultaneously introducing new issues. Vision-based systems are often not durable enough to withstand the fatigue of grasping [6], while purely tactile or proprioceptive sensors do not have the resolution needed to determine granular object information [7]. There is a need for a robust sensing methodology that provides sufficiently detailed information for manipulation while minimizing added complexity to the fabrication process.

In this work, we address this need by creating a soft robotic grippper made out of internally sensorized cubic lattices (Fig. 1). These lattices are 3D printed, allowing for a one-step fabrication of both structure and sensing elements. This easy fabrication allows for a high density of sensors (6 per finger), while remaining simpler than vision-based systems. Furthermore, the internal nature of these air-based sensors creates a more durable sensorized gripper while maintaining a fast response time. A parallel jaw gripper made of two of these lattices is able to do complex tasks like weight estimation and object localization.

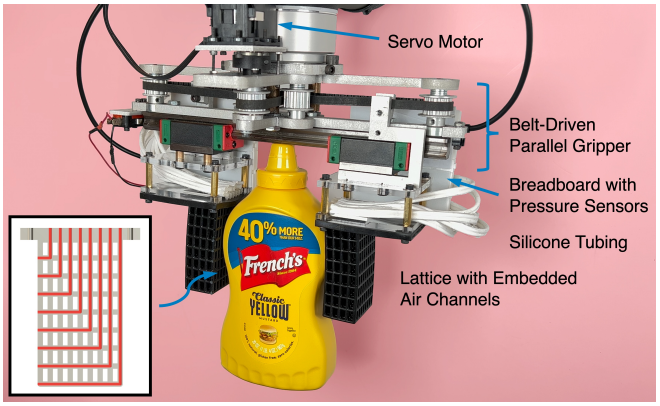

Fig. 1: Soft elastomer lattices are 3D-printed with L-shaped channels (highlighted in the inset) and mounted on a parallel gripper. Closed off and connected with pressure sensors, these channels provide sensory feedback during deformation.

We achieve this performance through the fluidic innervation technique [8], which relies on variations in internal air channel pressures to comprehend the forces acting on the lattice. While this technique showed immense promise for internal proprioception, no previous studies have explored its application in contact scenarios. In this paper, we offer an extensive analysis of this sensorization method within the context of grasping and its implications for vision-less manipulation tasks. Our contributions are:

1) Extension of the fluidic innervation technique to account for external contact
2) Characterization of the impact of lattice deformation and transient effects on the sensor readings
3) Design and fabrication of a manipulator capable of discerning an object's weight and grasp location

II. Related Work

In-hand object sensing for manipulation can broadly be categorized into (1) vision-based tactile sensing and (2) direct state measurements on the gripper itself.

Vision-based approaches have become quite popular due to leveraging advances in computer vision to obtain high resolution tactile information [7]. The most common approach is to point the camera towards a contact surface. This surface either has large compliance and three-part illumination, allowing for detection of object surface features (GelSight-style sensors, after [2]), or is physically textured

*These authors contributed equally to this work. All authors are with the MIT Computer Science & Artificial Intelligence Laboratory, Massachusetts Institute of Technology, 32 Vassar St, Cambridge, MA 02139, USA.
Correspondence: zhang@csail.mit.edu

with ridges and markers so the camera can track movement (TacTip-style sensors, after [3]). This sensing methodology has been used in many impressive manipulation applications, such as estimating shear forces [9], localizing a grasped object [10], and dynamically swinging an object upright [11]. Despite the richness of this data, vision-based techniques suffer from significant durability issues. Since the surface must be soft enough for large deformations to be visible, the surface can easily tear. This lets outside light bleed into the camera aperture, creating out-of-distribution data that is difficult to account for [6]. Another approach is to point the camera towards the internal kinematic chain of the finger [12], [13]. This avoids many of the durability issues of the surface-based model. However, since these systems require the camera to be physically embedded within the finger itself, this introduces complexity to their fabrication.

For approaches measuring the gripper state, most hardware innovation has happened within the realm of soft robotics. It is difficult to incorporate sensors within a rigid structure, so most proprioceptive sensing approaches rely on using motor feedback [4], [14] or force-torque sensors attached to the base of the end effector [15]. Soft robots have more design flexibility for incorporating sensors, but these sensors must comply with the object's flexible geometry. Some sensorization approaches include off-the-shelf force-sensitive resistors and bend sensors [16], optical waveguides [5], capacitive silicone sensors [17] and liquid metal strain sensors [18]. These methods all have the difficulty of ensuring sufficient resolution. Rigid approaches can only give bulk information about the grasping process, while soft robotic approaches require manually placing sensors within a structure. Scaling up these methods requires significant fabrication effort and are limited by the physical space on the gripper.

Our system addresses the issues of both of these approaches by having a one-print fabrication method that can embed many sensors. While the fluidic innervation method cannot compare to the resolution of vision-based systems, it still has sufficient sensitivity to achieve similarly complex tasks (Sec. VI). Our sensorized grippers also avoid many of the durability issues of vision-based systems since the internal channels do not rely on a soft interface with the grasped object. This lets us use a higher durometer material (Shore hardness of 68) that protects the air channels from being burst.

III. Hardware Design

In this section, we discuss how we build a sensorized gripper through fluidic innervation. We describe how we design the lattices to accommodate a large amount of sensor channels, the electronics necessary to read the pressures of those channels, and the assembly of the lattices into a gripper.

A. Sensorized Lattice Design

The basis of our sensorized lattice design is our prior work [8], where we 3D printed lattices out of an elastomeric polyurethane (EPU 40, Carbon Inc.) with empty air channels inside their struts. After sealing them off, the internal air channels experience a volume change as the lattices deform. By the ideal gas law ($pV = nRT$), this change in volume corresponds to a change in pressure, which can then be measured using an off-the-shelf pressure transducer.

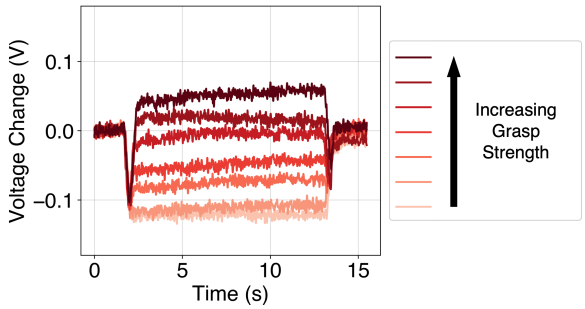

Fig. 2: Sensor response to grasping a box with different strengths. Initially, the voltages decrease for all grasp strengths due to bending of the lattice. Later, while the box is held in place, stronger grasps compress the lattice more and lead to higher voltages.

For our gripper's fingers, we choose a $8 \times 8 \times 4$ cubic lattice with a unit cell size of $9 \times 5 \times 6\,\mathrm{mm}$ and a strut diameter of $2.7\,\mathrm{mm}$. When designing a channel, we need to consider that the pressure is constant across its entire length. A change in pressure is only expressed if the channel experiences deformation along that length (i.e. changes to V in $pV = nRT$). This means that how a channel is routed greatly affects what information it can capture. For example, a vertical channel can only capture location information horizontally, as a force at the top of the channel will create the same response as the same force at the bottom of the channel. To capture both horizontal and vertical information within a single channel, we choose our sensor channels to be in the shape of an L. This asymmetric L shape is particularly useful since our gripper is composed of two lattices placed opposite one another. The opposite-facing L's complement one another and give us more information about an object grasped between the two jaws. The channels have an inner diameter of $1.5\,\mathrm{mm}$ and are designed in Rhino Grasshopper. Each lattice has 6 channels within the first layer away from the gripper surface (Fig. 1 inset, light red). The L's start from the outermost edge of the lattice and nest inwards. Although the lattice is printed with all 8 possible channels, preliminary experiments show that the two innermost channels have too little volume to provide a meaningful reading. We thus only connect the outer 6 channels, leaving a blind spot on the top left of the lattice. The opposite-facing lattice provides sensor information for that area.

B. Readout Electronics

To read the pressures of the channels, we follow the same methodology as [8]. Each of the embedded air channels is sealed off on one end with a silicone sealant and connected via silicone tubing to differential pressure transducers on the other end. The transducers are connected to the analog inputs of an Arduino Micro and sampled at 100 Hz.

We choose differential pressure transducers as they allow

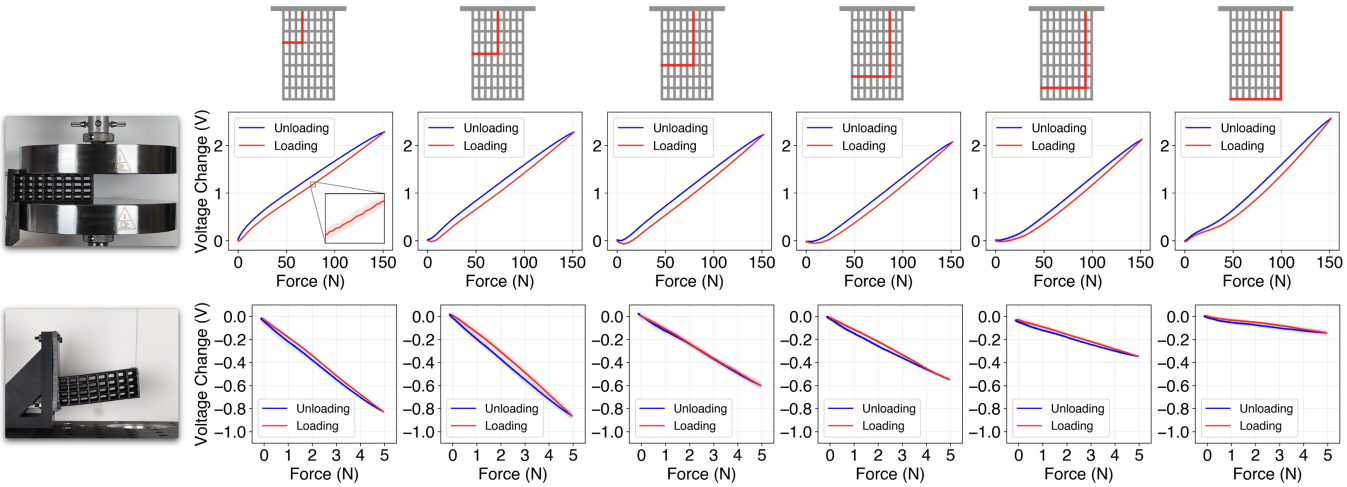

Fig. 3: Characterization of the sensorized lattice on tensile testing machine. First row of plots shows the sensor response during compression. Second row of plots shows sensor response during bending. Column-wise, the plots belong to the sensor highlighted in red in the top row cartoons. Barely visible error bands represent the standard deviation across 5 runs.

us to hook up a secondary line of tubing to serve as the "ground" for the measured sensor value. We use this secondary tube as compensation for thermal drift in the main line. The fluidic innervation technique is sensitive to temperature effects as, by $pV = nRT$, a change in temperature dramatically affects the pressure readings. By cutting secondary tubing that is roughly the same volume as the internal channel's plus the signal tubing's volume, we can braid that tube with the signal one and compensate for any thermal effects. Noting that the outermost channel (112 mm) is 2.7 times the length of the innermost channel (42 mm) and that the channel diameter (1.5 mm) is 1.5 times the tubing diameter (1 mm inner diameter, 3 mm outer diameter), we choose tubing lengths of 6, 8, 10, 12, 14, 16 inches for the signal line and 9.75, 13, 16.25, 19.5, 22.75, 26 inches for the thermal line.

We use 1 inch H2O transducers (All Sensors) for the left-hand lattice, and 0.5 inch H2O transducers (All Sensors) for the right-hand lattice. This asymmetry allows us to simultaneously capture subtle changes as well as large scale effects. Each of these transducers has their own offset voltage that ranges from 2.15–2.35 V when the pressure differential is equal to 0. When we perform a reading, we subtract this voltage offset. The exception to this is the experiment on diffusion and leakage in Sec. IV-B, where we examine the transducers signals from saturation (5 V) to equilibrium (offset voltage).

C. Gripper and Test Platform

To turn the sensorized lattices into a gripper, we mount the lattices as jaws of a belt-driven parallel gripper (introduced in [19]) that is mounted on a robot arm (UR5, Universal Robots). A servo (Dynamixel MX-28T, ROBOTIS) moves the belt with two carriages on a shared linear rail. These carriages move in opposite directions, each holding a sensorized lattice and its readout electronics (Fig. 1). The supplementary video shows a rotating view of the gripper setup.

To prevent the tubing from interfering with the gripper's full range of motion, we adjust the signal and thermal tubing lengths to 6, 7, 7, 8, 8, 9 inches. These lengths were chosen as a trade-off between minimizing slack within the gripper system and maximizing thermal compensation. To avoid excess movement, we glue the thermal tubing to the signal tubing with silicone adhesive (Sil-Poxy, Smooth-On).

IV. Characterization

The primary physical principle governing the air's behavior inside the channels is the ideal gas law $pV = nRT$. Our method is based on pressure transducers that measure changes in the pressure p. This section investigates how changes in the volume V and amount of air n affect the pressure readings p. Changes in T are assumed to be negligible due to the thermal compensation line. We separate this section into direct deformation of the lattice (change in V) and transient effects like unintended deformation due to the viscoelasticity of the material (change in V) and air escaping due to diffusion and leakage (change in n).

A. Deformation of the Lattice

Manual tests show that directly compressing the channels and bending the lattices lead to significant voltage readings (see supplementary video). To understand the dominant deformations that occur during manipulation, we grasp an empty acrylic box with increasing strength. We show the response of one representative sensor (second outermost on left lattice) in Fig. 2. The different grasp strengths are achieved by driving the servo to 7 pre-defined angles (from 103.7° to 130.1° in steps of 4.4°). For clarity, we filter the data in this plot with a 20 Hz low-pass filter.

From this experiment, we note the overall trend that the signal first goes down, plateaus at some fixed amount, then almost restores back to the original voltage value. We also note that as grasp strength increases, the plateau value increases as well. We attribute the initial negative effect to the lattices bending backwards from the grasped object's

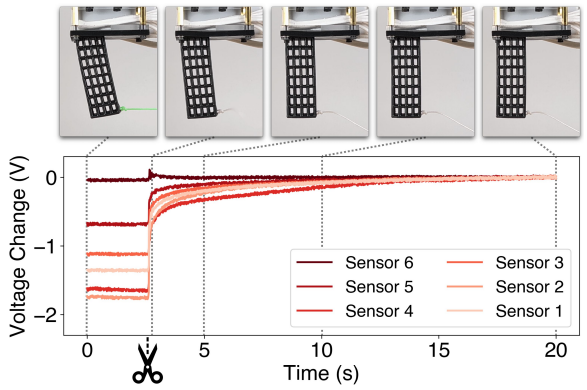

Fig. 4: Sensor response to viscoelastic material behavior. The lattice is held by a string (highlighted in green in the leftmost plot), which is cut at around 2.6 seconds. An exponential fit to the mean (not shown) has $\tau = 3.4\,\mathrm{s}$. Sensors are numbered in ascending order from inside to outside.

reaction force. From [8] we know that sensors on the outer surface during bending give negative values. Furthermore, since we use the same box for each trial, the gripper is always bent away by the same amount, explaining why all trials go to the same initial negative value. Likewise, the plateau increase corresponds well to compression. The amount the lattice is compressed increases as we increase the grasp force, which makes sense both intuitively and from [8]. We thus conclude bending and compression are the primary effects as the lattice deforms during a grasp and perform mechanical characterizations on these effects.

1) Compression: To characterize the lattice under compression, we use a universal testing machine (Instron, Illinois Tool Works) with parallel plates (Fig. 3). We fix the lattice with tape and compress to $1.7\,\mathrm{mm}$ at a rate of $0.1\,\mathrm{mm/s}$ for a total of 5 cycles, measuring across all sensors. The resulting plots confirm the intuitive result that compression leads to higher pressures and positive reading. Their very similar response curves with difficult to read error bands indicate a high level of repeatability. This is partially due to our choice of tubing lengths that equalize the total volume of the channels and tubing for all sensors. The discrepancy between the loading and unloading curve can be attributed to elastic hysteresis of the lattice material itself. This effect is characteristic for elastomers and caused by energy dissipation during deformation.

2) Bending: To characterize the lattice under bending, we mount a 3D-printed fixture onto the Instron's table to hold the lattice in place. We then attach a braided nylon string to the end of the lattice, fixing the other end of the string to the movable upper grip fixture of the Instron. While this string is elastic, it is two orders of magnitude stiffer than the lattice, making its extension negligible compared to the deformation of the lattice[1]. The string is pulled to

[1]We measure a string stiffness of $k_S = 25\frac{\mathrm{N}}{\mathrm{mm}}$ and a total stiffness of the lattice-string system of $k_T = 0.4\frac{\mathrm{N}}{\mathrm{mm}}$ on the Instron. Modeling the total system as two springs in series, the contribution of the string to the total extension is then given as $\frac{\Delta x_S}{\Delta x_T} = \frac{k_T}{k_S} = 1.6\%$.

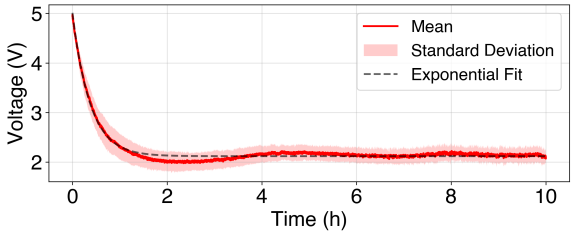

Fig. 5: Leakage and diffusion over a 10-hour period after pressurization, averaged across all sensors. The exponential fit has $\tau = 22\,\mathrm{min}$.

an extension of $15\,\mathrm{mm}$ at $0.5\,\mathrm{mm/s}$ for a total of 5 cycles. The resulting plots confirm our earlier intution that bending leads to negative readings. Since the forces are orders of magnitude lower than those for compression, the hysteretic effects are less prominent but still present. They are also minimized since the structure is deformed more than the material itself. This structure-based movement also explains why the bending affects the sensors unequally. The longest sensor along the edge matches this bulk movement, while the smaller sensor close to the pivot point is stretched more.

Overall, the high repeatability underscores the robustness of the sensorization method. The very clean data requires minimal filtering to reveal distinct mechanical behaviors. This significantly eases our analysis of more complex manipulations in Sec. VI.

B. Transient Effects

In addition to the above deformation effects, we must also consider transient effects on the sensor response. Time-dependent effects like viscoelasticity, diffusion, and leakage are present in any elastomer lattice with silicone tubing. The goal is to get a timescale for these effects to ensure they do not excessively affect our readings.

1) Viscoelastic Material Response: To evaluate the viscoelastic restoring response of the elastomeric lattice, we record the lattice's response to a sudden cut (Fig. 4). Similar to the bending setup, we use a nylon string to pull on the lattice when mounted in the gripper. We wait until the signal plateaus and then cut the string sharply with scissors. The lattice will slowly restore back to its original shape, allowing us to measure its viscoelastic response. Fitting an exponential to the mean of the sensor responses gives a time constant of $\tau = 3.4\,\mathrm{s}$ and an RMSE of $0.026\,\mathrm{V}$. This lets us conclude that the effect is on the scale of seconds and not negligible, but can be mitigated by either re-normalizing the readings or waiting for a few seconds between grasps.

2) Leakage and Diffusion: Since silicone is air permeable and the tube fittings are not perfectly airtight, there will inevitably be diffusion and leakage of air in our system. While we rely on this to equilibriate the channels after installation, this effect must be slow enough not to affect the signals during grasping. To measure these effects, we record the transducers' response over a 10-hour period after we pressurize the signal line. Since each of the transducers have their own offset voltage (Sec. III-B), we align the moment

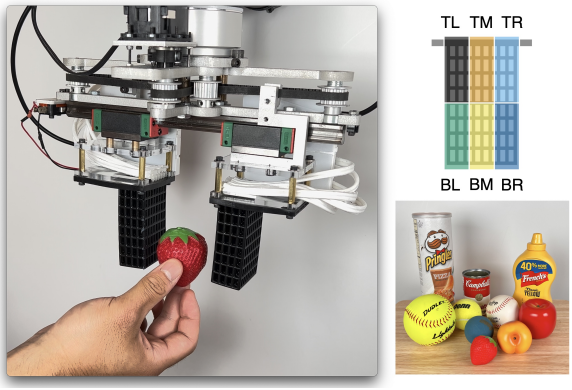

Fig. 6: Experimental setup for sensing different grasp locations. (Top Right) The grasp surface is divided into 6 regions. (Bottom Right) 10 objects from the YCB object set [20]. (Left) Each object is held in place by hand in each region and grasped with 3 different strengths.

Fig. 7: Experimental setup for sensing different weights. (Left) For each of 7 pre-defined grasp strengths, the gripper makes 3 attempts to grasp a box filled with weights. A digital scale confirms a successful lift-off. (Right) The experiment is repeated with 15 weights from 0 to 1400 g.

when the pressures of different sensors cross the upper bound of the sensor range and compute a mean and standard deviation across all six sensors (Fig. 5). Our exponential fit has an RMSE of 0.056 V and a time constant of $\tau = 1340\,\mathrm{s}$, or 22 min. Since a grasp happens over the course of seconds, we neglect this effect in our further analysis. We would need to reconsider this assumption for tasks that require holding an object for an extended period of time.

V. Experimental Setup

In this section, we apply our system to grasping tasks and discuss our evaluation setup for the sensors' performance at understanding a grasped object's location and mass.

A. Sensing Different Grasp Locations

Localizing an object within the gripper is trivial for vision-based tactile sensors but can be challenging for strain- and force-sensing approaches due to their lower resolution (Sec. II). To evaluate our gripper's performance in this task, we divide the inner surface of the lattice into 6 regions (Fig. 6 top right). We manually hold an object in each of the 6 regions and close the gripper with 3 different grasp strengths. We repeat this for 10 different objects from the YCB object set [20] (Fig. 6 bottom right). The supplementary video shows the execution of the experiment for one of these objects. Our goal is to see if we can tell the grasp location purely from our sensor data alone.

B. Sensing Different Object Weights

Detecting an object's weight is important for any gripper as it affects slip and potential downstream use in dynamic tasks. For the gripper to hold an object between its jaws, the static Coulomb friction model puts a lower bound on the normal force the gripper exerts on the object, dependent on the mass of the object. Since Fig. 3 establishes a relationship between sensor readings and the normal force, we seek a relationship between the sensors' readings and the mass by finding the lowest possible normal force. To find this relationship, we repeatedly grasp an acrylic box that weighs 52 g and load it with weights ranging from 0 to 1400 g (Fig. 7). We use the same hard-coded servo angles as in Fig. 2 to achieve different grasp strengths and measure the sensor readings when exerting the minimal normal force required for object lift-off. We use a USB-connected scale (Scout SPX2201, OHAUS) to verify successful lift-off.

VI. Experimental Results

A. In-Hand Localization

Fig. 8 shows the 2-dimensional t-distributed stochastic neighbor embedding (t-SNE) of the 12-dimensional sensor data collected during this experiment, colored by location, using the default parameters from the scikit-learn Python library. Note that t-SNE is an unsupervised learning technique, i.e., the location labels are not used during the computation. The plot reveals neat clusters, which suggests that the data is highly structured and easily separated into the 6 grasp locations despite the approximate placement by hand, the 10 different objects, and the 3 different grasp strengths.

Indeed, a *very* simple neural network with one hidden layer and 256 neurons (4,870 trainable parameters) trained to classify the data into the 6 grasp locations achieves an accuracy of $96.1\% \pm 0.5\%$ (averaged over 100 trials) after just one epoch and less than 2 seconds of training. To evaluate the generalization on an unseen object, we retrain the neural network with all data except from the mustard bottle, which is the object most different from all others in terms of shape and grasp signature (see Fig. 1 for the grasp pose). The neural network still achieves an accuracy of $77.1\% \pm 1.6\%$, even though the mustard data is highly out of distribution. This shows that despite having significantly lower-dimensional data than vision-based tactile sensor, we are able to achieve similar localization results with minimal effort.

B. Weight Estimation

Fig. 9 shows the results of the weight experiment. On the left, we show the number of successful grasps out of 3 for each combination of grasp strength and mass. For each

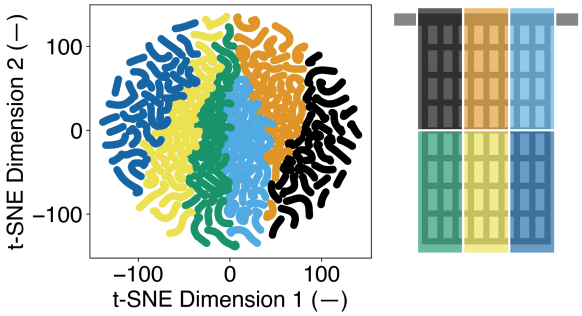

Fig. 8: Results of localization experiment. (Left) t-SNE of the 12D sensor data from location experiment. Clusters are clearly separated despite variations in objects and grasp strength. (Right) Corresponding locations on the lattice.

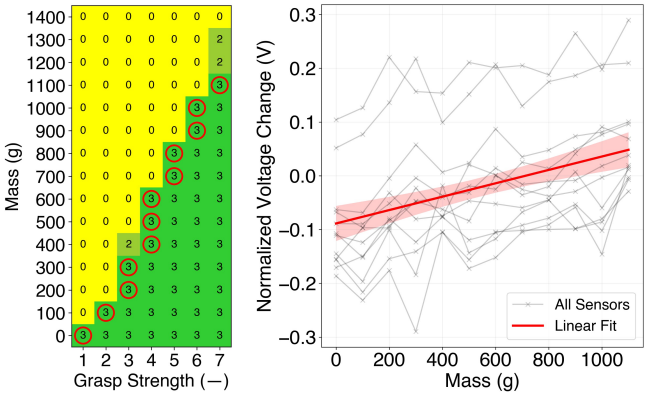

Fig. 9: Results of weight experiment. (Left) Number of successful grasps out of 3 for each mass-grasp strength combination. Red circle marks minimal grasp strength required for 3/3 grasps for any given mass. (Right) Average sensor response over the 3 grasps when exerting the minimal grasp strength for each mass. Linear fit to mean shows positive trend. Error band represents 95% confidence interval.

mass, we seek the minimum grasp strength required for 3/3 successful grasps (circled in red) and compute the average sensor signal across the 3 grasps for each of the sensors. This gives a voltage-mass relationship for each sensor that we plot in gray on the right. Since we use pressure sensors with different ranges for the two lattices, we first normalize the voltages to enable comparison. We then rescale the voltages with the ratio of the lengths of the signals lines used for characterization (Sec. III-B) and those used for the gripper (Sec. III-C) to enable comparison to the characterization results. A linear least squares fit across all sensors reveals a positive relationship between mass and sensor signals. We construct a 95% confidence interval with the conventional assumption that residuals are normally distributed.

We note that these results are in line with earlier observations. First, the top two lines that stand out in Fig. 9 right correspond to the outermost sensors of the left and right lattice. This is consistent with Fig. 3, where bending has the least effect on the outermost sensor. More importantly, the same trends seen in Fig. 2 apply here as well. The average sensor value in Fig. 9 right starts out negative, which means that for low weights, bending dominates. Since the lattice is bent by the same amount for every mass (due to grasping the same box), we conclude that the positive slope is solely caused by higher compression. This implies that we can take the slope of this line (voltage change due to increased mass) and combine it with the slope of the compression characterization plot (voltage change due to compression force) to reason about compression force. Specifically, this gives us an estimate of the minimally required compression force to grasp each mass. For slopes of $0.125\,\mathrm{V/kg}$ (Fig. 9 right) and $67.2\,\mathrm{N/V}$ (Fig. 3 first row average), we get $8.37\,\mathrm{N/kg}$. In other words, for every kilogram of added mass to the box, we need to increase the force by $8.37\,\mathrm{N}$ to ensure a secure grasp. Moreover, since this value is solely determined by the gravitational acceleration and the friction coefficient according to the static Coulomb friction model, we can provide an estimate of the friction coefficient. Our calculations gives us a friction coefficient of 0.59, a plausible value for the lattice-box contact. This experiment shows that the clean readings we get from our sensors enables us to conduct rich experiments about grasping performance, including extracting information like friction coefficient and grasping force that are difficult to get for any robotic gripper.

VII. Discussion

In this paper, we have explored the power of applying the fluidic innervation technique to a gripper. Conceptually, this method is very simple: 3D print empty channels within a cubic lattice structure, use those lattices as the jaws of a parallel jaw gripper, and measure their pressure. Despite that simplicity, this sensorized gripper performs on par with more intricate tactile sensors. We can provide feedback for tasks such as weight detection and object localization, which are traditionally addressed by vision-based sensors and dense neural networks for interpretation. Through just 12 signals, the method proposed here efficiently accomplishes a range of tasks and provides novel insights for measuring friction coefficients and grasping forces.

We plan to build on the versatility of fluidic innervation by adding complexity to the lattice geometries. This includes exploring the potential for programmable material properties like spatially varying stiffness and creating denser channel layouts for more comprehensive sensing. We hope that this approach can serve as a foundational technique for better identification of objects based on size, stiffness, and shape.

Acknowledgments

The authors thank Gregory Xie for help with microcontroller programming and Shruti Garg for systems integration work between the UR5, servo, microcontrollers, and ROS. This work was supported by the National Science Foundation EFRI grant #1830901 and the Gwangju Institute of Science and Technology.

References

[1] A. Billard and D. Kragic, "Trends and challenges in robot manipulation," *Science*, vol. 364, no. 6446, p. eaat8414, 2019.

[2] M. K. Johnson and E. H. Adelson, "Retrographic sensing for the measurement of surface texture and shape," in *2009 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2009, pp. 1070–1077.

[3] 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*. IEEE, 2009, pp. 1–6.

[4] B. Belzile and L. Birglen, "A compliant self-adaptive gripper with proprioceptive haptic feedback," *Autonomous Robots*, vol. 36, pp. 79–91, 2014.

[5] 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, Dec. 2016.

[6] O. C. Kara, N. Ikoma, and F. Alambeigi, "Hysense: A hyper-sensitive and high-fidelity vision-based tactile sensor," in *2022 IEEE Sensors*. IEEE, 2022, pp. 1–4.

[7] S. Zhang, Z. Chen, Y. Gao, W. Wan, J. Shan, H. Xue, F. Sun, Y. Yang, and B. Fang, "Hardware technology of vision-based tactile sensor: A review," *IEEE Sensors Journal*, 2022.

[8] R. L. Truby*, L. Chin*, A. Zhang, and D. Rus, "Fluidic innervation sensorizes structures from a single build material," *Science Advances*, vol. 8, no. 32, Aug. 2022.

[9] H. Song, T. Bhattacharjee, and S. S. Srinivasa, "Sensing shear forces during food manipulation: resolving the trade-off between range and sensitivity," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 8367–8373.

[10] S. Cui, R. Wang, J. Hu, J. Wei, S. Wang, and Z. Lou, "In-hand object localization using a novel high-resolution visuotactile sensor," *IEEE Transactions on Industrial Electronics*, vol. 69, no. 6, pp. 6015–6025, 2021.

[11] C. Wang, S. Wang, B. Romero, F. Veiga, and E. Adelson, "Swingbot: Learning physical features from in-hand tactile exploration for dynamic swing-up manipulation," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 5633–5640.

[12] Y. She, S. Q. Liu, P. Yu, and E. Adelson, "Exoskeleton-covered soft finger with vision-based proprioception and tactile sensing," in *2020 ieee international conference on robotics and automation (icra)*. IEEE, 2020, pp. 10 075–10 081.

[13] A. Zhang, R. L. Truby, L. Chin, S. Li, and D. Rus, "Vision-based sensing for electrically-driven soft actuators," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 11 509–11 516, 2022.

[14] T. Chen and M. Ciocarlie, "Proprioception-based grasping for unknown objects using a series-elastic-actuated gripper," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 6675–6681.

[15] M. Y. Cao, S. Laws, and F. R. y Baena, "Six-axis force/torque sensors for robotics applications: A review," *IEEE Sensors Journal*, vol. 21, no. 24, pp. 27 238–27 251, 2021.

[16] B. Homberg, R. K. Katzschmann, M. R. Dogar, and D. Rus, "Robust proprioceptive grasping with a soft robot hand," *Autonomous Robots*, vol. 43, pp. 681–696, 2019.

[17] L. Chin, M. C. Yuen, J. Lipton, L. H. Trueba, R. Kramer-Bottiglio, and D. Rus, "A simple electric soft robotic gripper with high-deformation haptic feedback," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 2765–2771.

[18] Z. Xie, F. Yuan, Z. Liu, Z. Sun, E. M. Knubben, and L. Wen, "A proprioceptive soft tentacle gripper based on crosswise stretchable sensors," *IEEE/ASME Transactions on Mechatronics*, vol. 25, no. 4, pp. 1841–1850, 2020.

[19] L. Chin, F. Barscevicius, J. Lipton, and D. Rus, "Multiplexed Manipulation: Versatile Multimodal Grasping via a Hybrid Soft Gripper," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, May 2020, pp. 8949–8955.

[20] B. Calli, A. Singh, A. Walsman, S. Srinivasa, P. Abbeel, and A. M. Dollar, "The ycb object and model set: Towards common benchmarks for manipulation research," in *2015 international conference on advanced robotics (ICAR)*. IEEE, 2015, pp. 510–517.