eFlesh: Highly customizable Magnetic Touch Sensing using Cut-Cell Microstructures

Venkatesh Pattabiraman Zizhou Huang Daniele Panozzo Denis Zorin Lerrel Pinto Raunaq Bhirangi

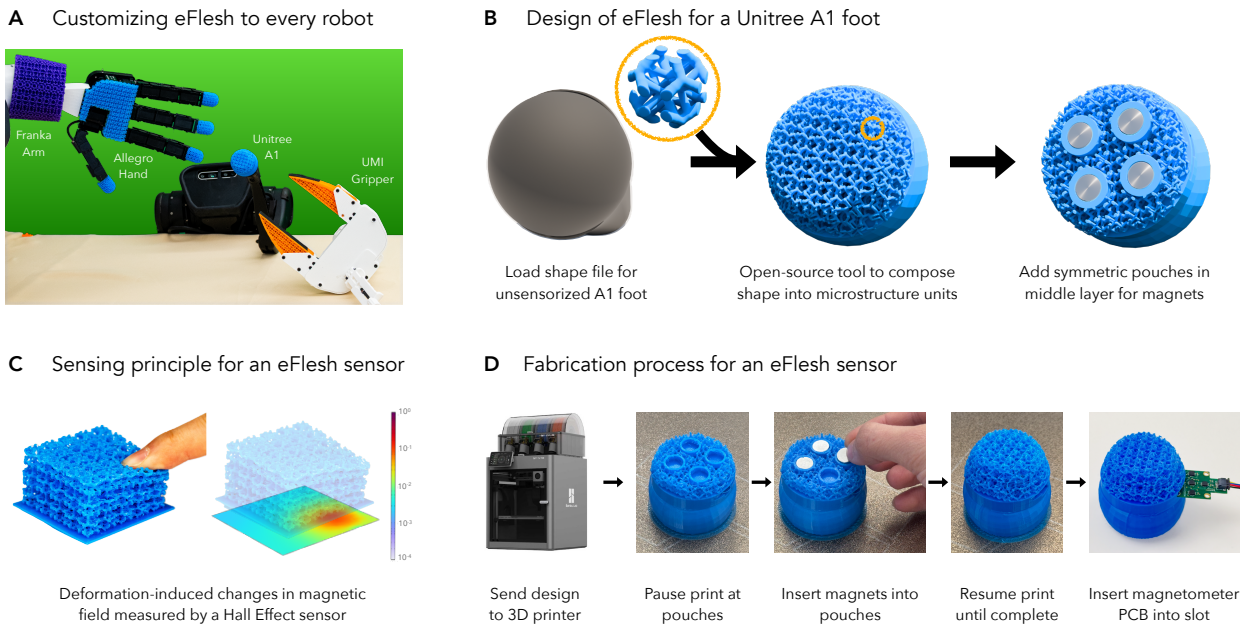

***Abstract*—If human experience is any guide, operating effectively in unstructured environments—like homes and offices—requires robots to sense the forces during physical interaction. Yet, the lack of a versatile, accessible, and easily customizable tactile sensor has led to fragmented, sensor-specific solutions in robotic manipulation—and in many cases, to force-unaware, sensorless approaches. With eFlesh, we bridge this gap by introducing a magnetic tactile sensor that is low-cost, easy to fabricate, and highly customizable. Building an eFlesh sensor requires only four components: a hobbyist 3D printer, off-the-shelf magnets (< $5), a CAD model of the desired shape, and a magnetometer circuit board. The sensor is constructed from tiled, parameterized microstructures, which allow for tuning the sensor's geometry and its mechanical response. We provide an open-source design tool that converts convex OBJ/STL files into 3D-printable STLs for fabrication. This modular design framework enables users to create application-specific sensors, and to adjust sensitivity depending on the task. Our sensor characterization experiments demonstrate the capabilities of eFlesh: contact localization RMSE of 0.5 mm, and force prediction RMSE of 0.27 N for normal force and 0.12 N for shear force. We also present a learned slip detection model that generalizes to unseen objects with 95% accuracy, and visuotactile control policies that improve manipulation performance by 40% over vision-only baselines – achieving 91% average success rate for four precise tasks that require sub-mm accuracy for successful completion. All design files, code and the CAD-to-eFlesh STL conversion tool are open-sourced and available on https://e-flesh.com.**

I. Introduction

Physical interaction with the world is nearly unimaginable without our sense of touch. While roboticists have long recognized the necessity of exteroceptive tactile feedback for complex manipulation [10, 8], most recent work has either entirely forgone the use of tactile sensing [3, 4]. With eFlesh, we seek to tackle this problem by leveraging recent advances in computer graphics and additive manufacturing to create a low-cost, customizable sensor that anyone can fabricate with just a hobbyist 3D printer, off-the-shelf magnets, and less than 5 minutes of active involvement.

Building on recent advances in additive manufacturing and computer graphics, we introduce eFlesh, a novel magnetic tactile sensing platform that fuses state-of-the-art approaches to parameterized microstructural design with embedded magnetic sensing. Our approach leverages the 3D boundary cell families proposed by Tozoni et al.[13], extending their utility beyond passive mechanics to active, multipurpose tactile sensing. Arbitrary surface geometries can be sensorized as composite microstructured tiles, as illustrated in Fig. 1A, while magnetometer circuit boards are embedded within or beneath the structure to transduce deformation into rich magnetic signals (Fig. 1C). This combination of material programmability, scalability, and ease of fabrication positions eFlesh as a powerful, versatile solution for integrating tactile feedback into

custom robotic systems. In this paper, we present a concrete methodology for the design and fabrication of eFlesh, and demonstrate its versatility across a range of robotic tasks and sensing applications. Our main findings can be summarized as follows:

1) We present eFlesh, a magnetic tactile sensor that can be fabricated into any 3D-printable convex shape with off-the-shelf magnets and Hall sensors.
2) We characterize the response of eFlesh and demonstrate a contact localization RMS error of 0.5 mm, and force prediction RMS errors of 0.27 N for normal force and 0.12 N for shear force.
3) We integrate eFlesh with learning based approaches for slip detection, where we achieve a success rate of 95% on unseen objects, and visuotactile learning, where we outperform vision-only baselines by 41%, while achieving $> 90\%$ success rates on set of contact-rich, precise manipulation tasks requiring sub-mm precision for success.
4) We propose a novel technique to circumvent the interference problem that has plagued magnetic sensing, by placing magnets of alternating polarities next to each other, and demonstrate that with this technique, inter-sensor as well as external magnetic interference correspond to signal from less than 1mm of sensor deformation.

All design files, code, and the CAD-to-eFlesh STL conversion tool will be open-sourced.

II. Results

eFlesh is a low-cost, modular, and accessible magnetic tactile sensing platform designed for both customizability and scalability. Our approach leverages parameterized, durable microstructures to decouple sensor design from complex fabrication constraints. The fabrication is fully compatible with consumer-grade 3D printers, lowering the barrier to entry for non-expert users. In addition to the ease of fabrication, we employ machine learning to demonstrate the richness and robustness of eFlesh signal. Our experiments also show that eFlesh maintains high signal consistency across instances – enabling scalable deployment scenarios like contact localization, force estimation, and generalizable policy learning.

To accommodate diverse sensor geometries, we adopt a variation of the infill strategy proposed in [13]. Given an arbitrary convex sensor shape, we first generate a voxelized lattice composed of microstructure units, trimmed to fit the desired form. Next, based on the number and size of magnets to embed, we generate symmetric, press-fit pouches in the sensor's mid-plane to hold them. This modular design framework greatly simplifies the process of sensorizing a variety of surfaces, as illustrated in Fig. (B and D).

In the remainder of this section, we present a number of experiments demonstrating the capabilities of eFlesh across a range of contact reasoning tasks. We start with controlled contact localization and force estimation tasks where we train machine learning models to demonstrate the richness and utility of the eFlesh signal. Following this, we explore two deployment-driven scenarios:

1) **Learned slip detection:** A Hello Robot Stretch uses eFlesh to predict when a grasped object is being pulled out of its grasp – demonstrating the sensor's slip detection capability and deployability in real-time interaction scenarios.
2) **Contact-rich visuotactile policy learning:** Based on the VisuoSkin framework presented in [10], we integrate eFlesh into a policy learning pipeline for four contact-rich tasks: USB insertion, plug insertion, credit card swiping and whiteboard erasing. These tasks highlight the sensor's utility in capturing fine-grained contact information critical for precise, closed-loop manipulation.

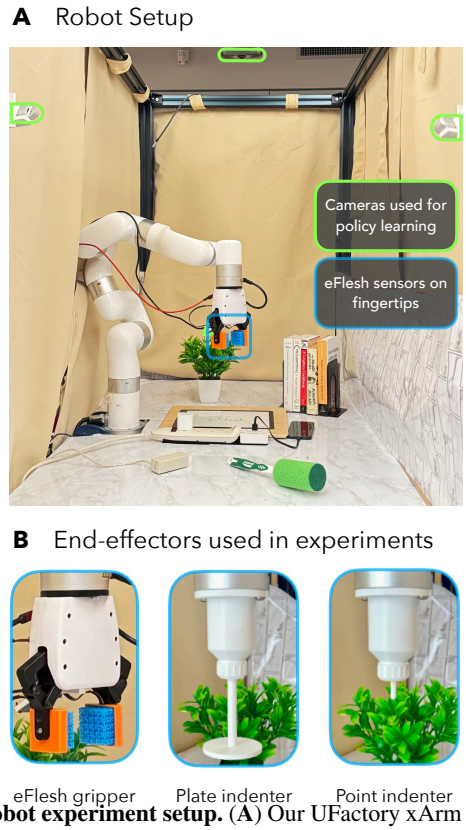

Fig. 1: **Robot experiment setup.** (**A**) Our UFactory xArm 7 robot and experimental environment. (**B**) End-effectors used in policy learning and sensor experiments.

A. Contact localization

An essential factor in selecting a tactile sensor is its contact resolution – its ability to distinguish between spatially proximal contact points. To evaluate the contact resolution of eFlesh, we train a neural network to predict the planar contact location (x, y) relative to the sensor's center, as well as the contact depth, z, relative to its surface, from the raw magnetic field measured by the magnetometers. We place the sensor on a flat table and use a 6mm-diameter 3D-printed hemispherical indenter mounted on a UFactory xArm 7 robot to probe the sensor surface at 1 mm intervals along x and y within a 30mm $\times$ 30mm grid as shown in Fig. 2A. For each indentation, we

uniformly sample the indentation depth, z, between 0.2 mm and 4.2 mm from the sensor surface, and record the change in magnetic field along with the corresponding contact location label, (x, y, z), obtained from robot proprioception. We repeat this process for five passes over the entire grid resulting in a total of 4,500 labeled samples.

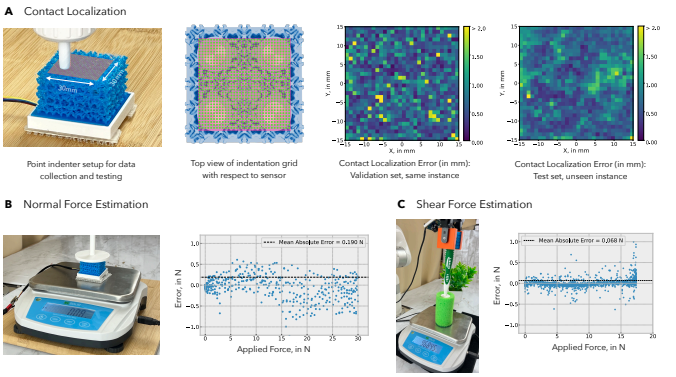

Fig. 2: **eFlesh Sensor Characterization Experiments.** (**A**), (Left to Right) The UFactory xArm robot arm equipped with a point indenter end-effector probing an eFlesh sensor instance for contact localization; Top view of the 30mm $\times$ 30mm indentation grid overlaid on the sensor; Heatmap visualizations of contact localization errors on the validation set of train instance and an unseen eFlesh instance. (**B**), xArm with a plate indenter compressing an eFlesh sensor on a weighing scale for normal force estimation, alongside the resulting error plot of our trained estimation model. (**C**), xArm equipped with an eFlesh gripper pressing a grasped foam stick against a weighing scale, alongside the resulting shear force estimation error plot.

Our neural network is a simple MLP with two hidden layers with 128 nodes each and ReLU activations, mapping the 15-dimensional change in magnetic field (B) to the 3-dimensional contact location: (x, y, z). We deliberately avoid random sampling when splitting the dataset for training and validation. Instead, we use data from the first four passes over the indentation grid for training, and reserve the fifth pass for validation. This setup better reflects real-world deployment conditions where sensor response tends to drift over time – a well-documented challenge in soft sensing systems [7, 1]. Under this protocol, the model achieves an $\mathrm{RMSE}_{x,y}$ of $0.5\,\mathrm{mm}$ and an $\mathrm{RMSE}_{\mathrm{depth}}$ of $0.16\,\mathrm{mm}$ on the validation set, demonstrating that eFlesh can localize contact with sub-mm precision. We find that this is slightly worse than a random train-validation split which results in a validation $\mathrm{RMSE}_{x,y}$ of $0.4\,\mathrm{mm}$ and a validation $\mathrm{RMSE}_{\mathrm{depth}}$ of $0.12\,\mathrm{mm}$, justifying our temporal train-validation split. To better contextualize these results, Fig. 2A also visualizes contact localization error at every point of the indentation grid for the validation set, as well as the test set comprised of indentation data from an unseen eFlesh instance. These plots illustrate the consistency of low contact localization error over the entire surface of the sensor as well as the generalization of this result to unseen instances of eFlesh.

B. Normal Force estimation

To further characterize the force response of eFlesh, we train a separate neural network to predict normal force from raw magnetic field signals. We swap the hemispherical point indenter for a flat plate indenter shown in Fig. 1B and place a weighing scale underneath the sensor (Fig. 2B). To collect data for the experiment, we apply compressive forces ranging from no force to 30 N while recording synchronized measurements from both the sensor and a weighing scale. For training and validation, we adopt a temporal split of the dataset mirroring the protocol used in the contact localization experiment: the first 7,200 data points are used for training, while the remaining 1,800 points form the validation set. We find a normal force prediction RMSE of 0.27 N, corresponding to a pressure of 125 Pa, validating eFlesh's ability to capture normal force/pressure with a high degree of precision. Fig. 2B also shows the variation of force prediction error as the amount of applied force is increased, illustrating that maximum error remains under 1N across the entire range of applied forces.

C. Shear force estimation

A key aspect of human dexterity is the ability to detect impending slip and adjust manipulation strategies accordingly [2]. Tactile sensors can enable the same capability in robots if they can capture shear forces acting on the surface of the sensor before an object in contact slips out of grasp. In this experiment, we evaluate eFlesh's capability to estimate planar (shear) forces applied to the sensor surface. To do so, we use a foam cleaning stick (Fig. 2C), which enables the application of a broader range of shear forces to the sensor surface without causing slip. The stick is grasped using a parallel jaw gripper, with both gripper tips equipped using eFlesh. We then generate data by randomly sampling vertical displacements of the foam stick, pressing it against a weighing scale. During this process, we record synchronized raw magnetic field signals from the sensors and ground-truth force measurements from the scale. The applied shear forces span a range from 0 N to 17.5 N. We train a neural network to predict the planar shear force from the raw magnetic data, using the same temporal split strategy as in previous experiments to account for sensor drift: the first portion of the dataset is used for training, and the remainder is held out for testing. The trained model achieves a root mean squared error (RMSE) of 0.12 N, demonstrating eFlesh's effectiveness in accurately estimating shear forces.

D. Slip detection using deep learning

A closely related and equally important capability for robots operating in unstructured environments is the ability to detect when an object has slipped from their grasp. In this experiment, we demonstrate that eFlesh can be used to reliably detect object slip using a simple, generalizable linear classifier trained on interactions with a small set of objects. Our setup consists of a Hello Robot Stretch with integrated eFlesh as shown in Fig. 3A. A human operator slowly tugs on the grasped object for 1-2 seconds, after which a human annotator labels the sequence as "force" or "no-force" based on corresponding

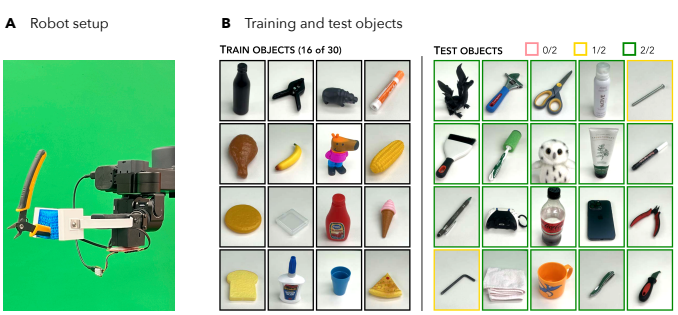

Fig. 3: **Slip Detection.** (**A**) Hello Robot Stretch setup used for slip data collection and classifier evaluation. (**B**) Object sets: (Left) A diverse ensemble used for training and (Right) a held-out validation set of objects to evaluate our model on.

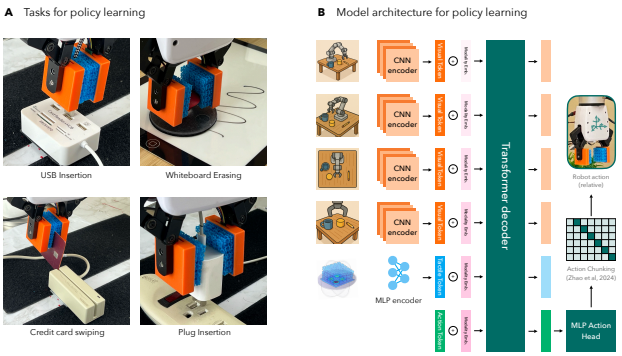

Fig. 4: **Robot Policy Learning Tasks and Model Architecture.** (**A**) We evaluate the performance of eFlesh on four precise tasks - USB Insertion, Whiteboard Erasing, Credit card swiping and Plug Insertion. (**B**) Visuo-Skin model architecture, used for our robot policy learning experiments. Images from the four cameras are encoded by a CNN encoder while the eFlesh signal is encoded by an MLP encoder. The encoded features are passed through a transformer decoder policy along with an action token, whose output is used to predict an action chunk. This action chunk is supervised using demonstration data, and policies employ exponential smoothing during deployment to avoid jerky motions [15].

videos. We estimate three statistics from a sliding temporal window of the raw magnetometer signals: (1) the norm of the x and y components of the magnetic field corresponding to each of the five magnetometers, (2) the maximum change in sensor signals within the window, and (3) the standard deviation of the signal in the window, and use these statistics to train a linear binary classifier. The training dataset consists of four trajectories each of 30 everyday objects varying in shape, size, weight, and surface texture. To evaluate generalization, we test the model on a separate set of 20 unseen objects (Fig. 3B). Despite its simplicity, the classifier achieves a high classification accuracy of 95% on this held-out set.

E. Visuotactile Policy Learning

Finally, a key application area for tactile sensors is in learning force-aware policies for robots for precise, contact-rich manipulation. We study the effectiveness of eFlesh in visuotactile robot policy learning through four contact-rich manipulation tasks shown in Fig. 4A. Drawing from the VisuoSkin [10] framework, we employ a transformer-based architecture for learned policies using behavior cloning (Fig. 4B). Here, we describe the robot environment, set of tasks, and the evaluation protocol.

1) Environment Setup: We use a Ufactory xArm 7 robot with its standard two-fingered gripper and an eFlesh sensor attached to either fingertip as shown in Fig. 4A. We use a Meta Quest 3 along with the OpenTeach [9] teleoperation framework to collect demonstrations for behavior cloning. To improve dataset diversity and policy robustness, we adopt a strategy similar to VisuoSkin [10], introducing uniformly sampled angular perturbations to the commanded robot velocity at each timestep during teleoperation. We record visual data from four camera views: three third-person perspectives and one egocentric wrist-mounted camera attached to the robot's gripper, all captured at 30 Hz. Simultaneously, we collect tactile data from the eFlesh sensors in the form of raw magnetometer signals sampled at 100 Hz. All data streams are synchronized and resampled to 10 Hz for downstream training and deployment.

2) Effect of eFlesh on performance: We observe that the trained visuotactile policies are able to successfully complete the aforementioned precise tasks, at a high average success rate of 91%. Given the low margins of error needed for these tasks, this underscores both the signal consistency and the ability of eFlesh to reliably capture the subtle sensory cues to then learn meaningful seeking behaviors in order to complete the tasks.

III. Conclusion

In this work, we introduced eFlesh, a 3D-printable, magnetic sensor that offers a low-cost and accessible approach to tactile sensing. By leveraging only a standard 3D printer, a small number of embedded magnets, and a magnetometer circuit board, eFlesh enables sensorizing a wide range of robots. This design opens the door to democratizing the fabrication of tactile sensors, significantly improving accessibility and customizability for in-the-loop design iteration.

The broader significance of eFlesh lies in its potential to equip robots with richer sensing, particularly as we deploy them in unstructured, dynamic environments. Vision alone is often insufficient in a number of tasks requiring fine-grained spatial awareness, especially in occluded and cluttered settings. eFlesh provides a lightweight, conformal, and minimally disruptive tactile solution with high sensitivity and a large dynamic range, in addition to shape and response customizability.

A key advantage of eFlesh over traditional magnet-based sensors [12] and existing soft sensing techniques [2, 8, 14] is its ease of fabrication and scalability. Unlike previous systems that rely on manual assembly, specialized equipment, or complex, multi-step fabrication, eFlesh is entirely 3D-printable and modular. This lowers the barrier to entry as well as prototyping time for users.

References

[1] Raunaq Bhirangi, Tess Hellebrekers, Carmel Majidi, and Abhinav Gupta. Reskin: versatile, replaceable, lasting tactile skins. In *5th Annual Conference on Robot Learning*, 2021. URL https://openreview.net/forum?id=87_OJU4sw3V.

[2] Raunaq Bhirangi, Venkatesh Pattabiraman, Enes Erciyes, Yifeng Cao, Tess Hellebrekers, and Lerrel Pinto. Anyskin: Plug-and-play skin sensing for robotic touch. *arXiv preprint arXiv:2409.08276*, 2024.

[3] Cheng Chi, Zhenjia Xu, Chuer Pan, Eric Cousineau, Benjamin Burchfiel, Siyuan Feng, Russ Tedrake, and Shuran Song. Universal manipulation interface: In-the-wild robot teaching without in-the-wild robots. In *Proceedings of Robotics: Science and Systems (RSS)*, 2024.

[4] Haritheja Etukuru, Norihito Naka, Zijin Hu, Seungjae Lee, Julian Mehu, Aaron Edsinger, Chris Paxton, Soumith Chintala, Lerrel Pinto, and Nur Muhammad Mahi Shafiullah. Robot utility models: General policies for zero-shot deployment in new environments. *arXiv preprint arXiv:2409.05865*, 2024.

[5] Siddhant Haldar, Zhuoran Peng, and Lerrel Pinto. BAKU: An efficient transformer for multi-task policy learning. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL https://openreview.net/forum?id=uFXGsiYkkX.

[6] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. doi: 10.1109/CVPR.2016.90.

[7] Tess Hellebrekers, Oliver Kroemer, and Carmel Majidi. Soft magnetic skin for continuous deformation sensing. *Advanced Intelligent Systems*, 1(4):1900025, 2019.

[8] Binghao Huang, Yixuan Wang, Xinyi Yang, Yiyue Luo, and Yunzhu Li. 3d-vitac: Learning fine-grained manipulation with visuo-tactile sensing. In *8th Annual Conference on Robot Learning*, nov 2024.

[9] Aadhithya Iyer, Zhuoran Peng, Yinlong Dai, Irmak Guzey, Siddhant Haldar, Soumith Chintala, and Lerrel Pinto. OPEN TEACH: A versatile teleoperation system for robotic manipulation. In *8th Annual Conference on Robot Learning*, 2024. URL https://openreview.net/forum?id=cvAIaS6V2I.

[10] Venkatesh Pattabiraman, Yifeng Cao, Siddhant Haldar, Lerrel Pinto, and Raunaq Bhirangi. Learning precise, contact-rich manipulation through uncalibrated tactile skins. *arXiv preprint arXiv:2410.17246*, 2024.

[11] Nur Muhammad Mahi Shafiullah, Anant Rai, Haritheja Etukuru, Yiqian Liu, Ishan Misra, Soumith Chintala, and Lerrel Pinto. On bringing robots home. *arXiv preprint arXiv:2311.16098*, 2023.

[12] Tito Pradhono Tomo, Massimo Regoli, Alexander Schmitz, Lorenzo Natale, Harris Kristanto, Sophon Somlor, Lorenzo Jamone, Giorgio Metta, and Shigeki Sugano. A new silicone structure for uskin—a soft, distributed, digital 3-axis skin sensor and its integration on the humanoid robot icub. *IEEE Robotics and Automation Letters*, 3(3):2584–2591, 2018.

[13] Davi Tozoni, Zizhou Huang, Daniele Panozzo, and Denis Zorin. Cut-cell microstructures for two-scale structural optimization. *Computer Graphics Forum*, 43, 07 2024. doi: 10.1111/cgf.15139.

[14] Boxin Xu, Luoyan Zhong, Grace Zhang, Xiaoyu Liang, Diego Virtue, Rishabh Madan, and Tapomayukh Bhattacharjee. Cushsense: Soft, stretchable, and comfortable tactile-sensing skin for physical human-robot interaction. *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5694–5701, 2024. URL https://api.semanticscholar.org/CorpusID:269605928.

[15] Tony Z Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning fine-grained bimanual manipulation with low-cost hardware. *arXiv preprint arXiv:2304.13705*, 2023.

APPENDIX

A. Limitations and Future Work

Magnetic sensing in uncontrolled environments remains a challenge. One known limitation is susceptibility to magnetic interference, especially in proximity to electromagnetic or ferromagnetic objects, including other instances of the magnetic sensor itself. While eFlesh effectively mitigates interference from most everyday sources and nearby eFlesh sensors, it may still encounter stronger magnetic fields in industrial settings such as power infrastructure or server environments. In future work, we aim to investigate materials with high magnetic permeability and explore passive shielding techniques to further reduce these effects.

Another open question concerns scaling and generalization. Our current results show strong promise in small-scale experiments, as well as signal consistency across instances conspicuously absent across the spectrum of existing tactile sensors. This data reusability along with robustness to magnetic interference makes eFlesh uniquely amenable to data collection and deployment at scale across a diverse set of robotic platforms and environments. To this end, future work will involve large-scale data collection in realistic settings along the lines of prior work in home robotics [11, 4], to validate generalization and robustness. We also envision learning tactile representations from these datasets to extract richer features from magnetic signatures, further enhancing applicability in novel scenarios.

Finally, the modular nature of eFlesh suggests exciting possibilities for integration into multi-functional systems combining sensing and actuation. By allowing sensing capabilities embedded directly into the structure of robots, eFlesh represents a step forward in material-based robotics that closely integrates form and function in the construction of robots.

B. Materials and Methods

1) Fabrication of eFlesh: The cuboidal instance of eFlesh used in all of the experiments presented above is composed of three layers of 5 $\times$ 5 microstructure grids stacked on top of each other. The Young's modulus of each layer is constant across the grid plane and increases linearly from the bottom layer to the top layer as $0.001E_f$, $0.0015E_f$ and $0.002E_f$ respectively, where E_f represents the Young's modulus of the filament (we use TPU 95A). Each microstructure cell is a cube of size 8mm, resulting in overall sensor dimensions of 40mm $\times$ 40mm $\times$ 24mm. After generating the microstructure grid, we add four lip-sealed pouches in the middle layer, distributed as shown in Fig . These pouches are designed to press-fit tolerance for N52 neodymium magnets that each have a diameter of 9.525mm and a thickness of 3.175mm.

This workflow is adaptable to any arbitrary 3D geometry, specified as a `.obj/.stl` file. Once the Young's modulus is specified, cell size is primarily constrained by the minimum printable thickness of the beams that make up each cell. On deciding a cell size, the number of layers can be computed by dividing the object's height by this cell size. We first create a cuboidal block of dimensions of the input mesh, thus encompassing it, composed of cut-cells of the specified cell size and Young's Modulus. Next, we trim the region that extends beyond the convex hull of the input shape, to preserve the surface shape. The user can then specify their cylindrical magnet dimensions and position, following which we parametrically add lip-sealed pouches within the microstructure grid.

We fabricate the sensors using a Bambu X1 Carbon 3D printer with a standard 0.4mm nozzle, printing with TPU filament with a Shore hardness of 95A, requiring no print supports. The model is processed in OrcaSlicer, an open-source slicing tool, which is configured to pause the print one layer before the magnetic pouch covers. At this point in the print, the user inserts magnets into the pouch packets, with the north poles facing upwards in this case, but in general, polarities may be altered based on the considerations outlined in the previous sections. Printing is then resumed to encase the magnets in the pouch as well as complete the rest of the sensor structure.

2) Neural Architectures and Training: In this section, we provide an overview of the training details and model architectures used for the sensor characterization and robot policy learning experiments.

3) Spatial and Force Resolution: We use a multi-layered perceptron (MLP) with two hidden layers, each with 128 nodes and ReLU activations. The network is trained using mean squared error (MSE) loss, with the Adam optimizer, a learning rate of `1e-3` and a batch size of 64. Target outputs - x, y, z coordinates are normalized for training using their corresponding ranges of 30 mm for x and y and 4 mm for z. RMSEs are reported after unnormalizing the predictions. Training is performed on a single NVIDIA RTX 3080 GPU over a maximum of 1000 epochs, requiring under 5 GPU minutes.

4) Visuo-Tactile Robot Learning: Drawing from Visuo-Skin [10], we use a multi-sensory transformer architecture [5], where we fuse the visual and tactile observations as a modified ResNet-18 [6] encoding and a 2-layer MLP encoding, respectively, projected to the same dimensionality.

Visual data is resized to 128x128 and the eFlesh data corresponds to 15-dimensional vectors for each fingertip, and the time-synchronized visuo-tactile dataset is subsampled by a factor of 5. We tokenize each of the two fingertips' eFlesh data separately prior to feeding them into the transformer decoder. Robot policies are then trained on the subsampled dataset up to 36,000 checkpoints, requiring 3 hours on a single NVIDIA RTX 8000 GPU.