AnySkin: Plug-and-play Skin Sensing for Robotic Touch

Raunaq Bhirangi[1,2,†], Venkatesh Pattabiraman[1], Enes Erciyes[1], Yifeng Cao[3], Tess Hellebrekers[4], Lerrel Pinto[1]

[1]New York University, [2]Carnegie Mellon University, [3]Columbia University, [4]Meta AI Research

[†] Correspondence to: raunaqbhirangi@nyu.edu

Abstract

While tactile sensing is widely accepted as an important and useful sensing modality, its use pales in comparison to other sensory modalities like vision and proprioception. AnySkin addresses the critical challenges that impede the use of tactile sensing – versatility, replaceability, and data reusability. Building on the simple design of ReSkin, and decoupling the sensing electronics from the sensing interface, AnySkin makes integration as straightforward as putting on a phone case and connecting a charger. Furthermore, AnySkin is the first uncalibrated tactile-sensor to report cross-instance generalizability of learned manipulation policies. To summarize, this work makes three key contributions: first, we introduce a streamlined fabrication process and a design tool for creating an adhesive-free, durable and easily replaceable magnetic tactile sensor; second, we characterize slip detection and policy learning with the AnySkin sensor; third, we demonstrate zero-shot generalization of models trained on one instance of AnySkin to new instances, and compare it with popular existing tactile solutions like DIGIT and ReSkin. Videos and more details can be found on our https://any-skin.github.io/.

1 Introduction

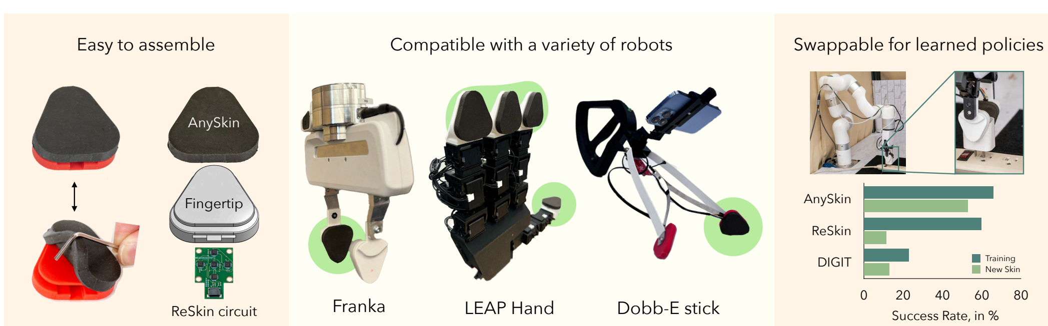

Figure 1: We present AnySkin, a skin sensor made for robotic touch that is easy to assemble, compatible with different robot end-effectors and generalizes to new skin instances.

Touch sensing is widely recognized as a crucial modality for biological movement and control [1, 2]. Unlike vision, sound, or proprioception, touch provides sensing at the point of contact, allowing agents to reason about forces. However, a closer examination of robotics literature reveals a different narrative. Prominent works and current state-of-the-art in robot learning primarily utilize vision with proprioception to train manipulation skills [3, 4, 5, 6], often ignoring touch. If touch is so vital from a biological perspective, why does it remain a second-class citizen in sensorimotor control?

In this work we present AnySkin, a new touch sensor that is cheap, convenient to use and has consistent response across different sensor instances. AnySkin builds on ReSkin [7], a magnetic-field based touch sensor, by improving its fabrication, separating the sensing mechanism from the interaction surface, and developing a new self-adhering, self-aligning attachment mechanism. This allows AnySkin to (a) have stronger magnetic fields, which significantly improves its sensor response, (b) be easy to fabricate for arbitrary surface shapes, which allows easy use on different end-effectors, (c) be easy to replace the sensor without adversely affecting the data collection process or the efficacy of models trained on previous sensors (Fig. 1).

We run a suite of experiments to understand the efficacy of AnySkin vis-a-viz other prominent touch sensors. Our main findings can be summarized below:

1. AnySkin can readily be used on a variety of robots (See Section 3).
2. AnySkin affords ML techniques for slip detection and precise policy learning (See Section 4).
3. AnySkin takes an average of 12 seconds to replace (See Section 4.3).
4. Models trained on one AnySkin transfer zero-shot to a different AnySkin with only a 13% reduction in performance on a plug insertion task compared to a 43% drop with ReSkin [7].

AnySkin will be fully open-sourced. Videos of fabrication, attachment, and robot policies are best viewed on our project website: `https://any-skin.github.io/`.

2 Related Work

Existing literature on tactile sensing explores a wide range of modalities, each with their advantages and limitations [8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18]. Magnetic tactile sensors [19, 20, 21] largely overcome the limitations of resistive, capacitative, optical and MEMS sensors due to three salient advantages: (a) separation of the sensing electronics from the sensing interface to improve robustness (b) compatibility with different form factors, and (c) an ability to capture shear forces [7]. In this work, we build on ReSkin sensors due to their lower cost and ease of fabrication. In Section 4, we also quantitatively demonstrate the improved consistency of AnySkin signal over ReSkin, and present a direct replaceability comparison with DIGIT [17] and ReSkin through policy learning experiments.

3 AnySkin: Components

AnySkin builds on ReSkin [7], a tactile skin composed of a soft magnetized skin coupled with magnetometer-based sensing circuitry. By detecting distortions in magnetic fields, ReSkin measures deformations caused by normal and shear forces [19, 7]. AnySkin uses the same circuitry as ReSkin, while introducing key design and fabrication changes to improve durability, repeatability, and replaceability: (1) Magnetizing skins post-curing using a pulse magnetizer, (2) introducing physical separation between magnetic elastomer and magnetometer circuit, (3) utilizing finer magnetic particles to achieve a more uniform particle distribution, (4) implementing a self-aligning design for fixed relative positioning of elastomers and circuitry. While some of these have been used in isolation in prior work [22, 23, 24], in this work, we quantitatively compare sensor response in Section 4.1.

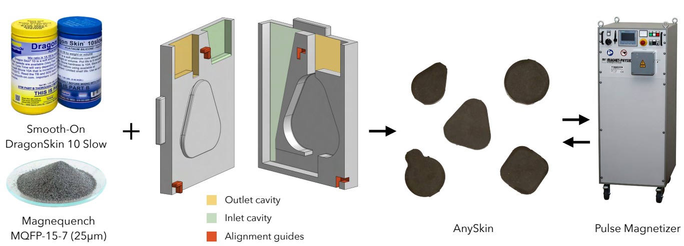

(a) Fabrication of AnySkin

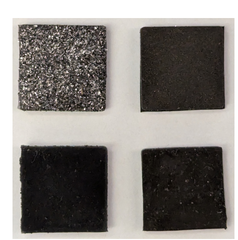

(b) Surfaces of skins

Figure 2: (a) AnySkin is made by mixing DragonSkin 10 Slow and MQFP-15-7(25μm) magnetic particles in a 1:1:2 ratio, and curing it in the two-part molds shown above. Cured skins are magnetized using a pulse magnetizer. (b) Particle concentrations in skins made with MQP-15-7 and MQFP-15-7.

4 Experiments and Results

4.1 Comparison between ReSkin and AnySkin signal

To quantitatively demonstrate the effect of each of the fabrication changes listed in Section 3 towards improving the consistency of AnySkin, we present a set of experiments analyzing the raw signal from the four different skins shown in Table 1, tracking the progression from ReSkin to AnySkin. All statistics presented are computed over five instances of each type.

Table 1: AnySkin shows lower variability across instances. Statistics computed over 5 samples of each type (PM: Pulse magnetizer, FP: finer particles, SA: self-aligning – AnySkin).

	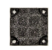				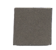			
Experiment	ReSkin		+PM		+FP		+SA	
	B_{xy}	B_z	B_{xy}	B_z	B_{xy}	B_z	B_{xy}	B_z
Average strength, in μT	1062	302	**1818**	**5212**	**1602**	**5784**	283	1265
Normalized σ (instances)	0.54	0.87	0.34	0.12	0.21	0.15	**0.12**	**0.10**
Normalized σ (1mm displaced)	1.38	1.43	0.25	0.11	0.18	0.07	**Self-aligning**	

We first see a significant increase in the raw magnetic field by using the pulse magnetizer for magnetization. Next, finer particles results in both a lower variability across instances as well as durability due to reduced leakage of particles resulting from stress concentration on the surface as shown in Fig. 2b. Finally, the self-aligning design of AnySkin removes the possibility of misalignment between skin and circuitry which introduced significant variability in the signal. This design also adds a physical separation between the electronics and the sensing interface which improves durability while maintaining the signal strength of ReSkin.

4.2 Slip Detection

We quantify AnySkin's ability to detect slip through a controlled experiment. Our setup consists of a Kinova Jaco arm and an Onrobot RG-2 gripper with integrated AnySkin. An object held by a human operator is grasped and lifted up slowly for 1 second. We use a set of 40 daily objects – 30 for training and 10 evaluation – with varying shapes, weights and materials. We collect 6 trajectories for each object by changing the grasping force, width and location. After the data collection is complete, a human annotator labels the sequence as slip or no-slip from the corresponding videos. Objects used, along with the data collection setup can be found on our website. The data collection frequency for tactile data is 100 Hz. We subsample the signal by 15 along the temporal axis and take the first difference. We use an LSTM [25] to train our slip prediction models purely from tactile data. Our model is able to detect slip on unseen objects with 92% accuracy.

4.3 Ease of replaceability

We compare the replacement time of AnySkin against other skins like DIGIT and ReSkin through a user study with 10 non-expert users. We find that ReSkin takes 236 ± 64 seconds, DIGIT takes 58 ± 22 seconds, while AnySkin takes a mere 12 ± 5 seconds to replace.

However, the most important consequence of the signal consistency and replaceability of AnySkin is its ability to enable policy generalization across different instances of the skin. We demonstrate cross-instance generalizability of AnySkin across three precise manipulation tasks shown in Fig. 3 and compare the cross-instance generalizability of policies trained on DIGIT, ReSkin and AnySkin on the plug insertion task. Our setup consists of a UFactory xArm 7DOF robot and four cameras – three fixed to the frame and one egocentric camera attached to the robot's wrist, with AnySkin sensor on the gripper tip. A Meta Quest 3 and the accompanying joystick controller are used to teleoperate the robot using Open-Teach [26], an open-source teleoperation framework, and collect 96 demonstrations per task. For each task, learned policies are evaluated on locations of the target object unseen in training data.

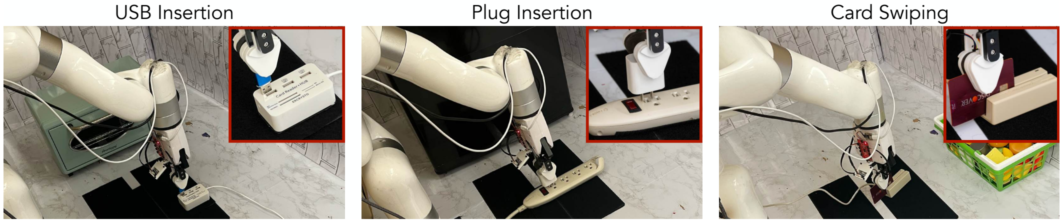

Figure 3: We evaluate the replaceability of AnySkin on a set of contact-rich, precision tasks.

The BAKU [27] architecture is used as the policy architecture. BAKU tokenizes each input using a modality-specific encoder: image inputs from cameras and DIGITs use ResNet-18 [28] encoders, while AnySkin and ReSkin inputs use MLP encoders. The policy uses a deterministic action head and action chunking [29] with exponential temporal averaging.

4.3.1 Evaluating cross-instance generalizability of trained policies

To investigate replaceability, we evaluate behavior cloning policies trained using a single instance of AnySkin on a new instance. We use different training and test skins for *each* of the presented tasks to avoid over-indexing on specific skin instances. Table 2 compares how well the policy performs with both the original and swapped skins for each of the precise, contact-rich tasks outlined above. Additionally, we train and evaluate a policy that only takes camera images as input to serve as a control experiment and verify that the policies rely on tactile data. We find that across tasks, performance drops by an average of just 15.6% and visuotactile policies with swapped skins continue to do significantly better than purely visual policies. This result demonstrates the strength and uniqueness of AnySkin as a tactile sensor for contact-rich manipulation.

Table 2: Success rates (out of 10) for policies when swapping out skins (averaged over 3 seeds)

Task	Cameras only	Cameras + Skin	
		Original skin	Swapped skin
Cross-instance generalization			
Plug Insertion	1.7 ± 0.6	6.7 ± 1.5	5.3 ± 2.5
Card Swiping	2.0 ± 1.0	7.0 ± 1.7	6.3 ± 0.6
USB Insertion	1.7 ± 1.2	5.7 ± 1.5	3.0 ± 1.0
Comparison across sensors – Plug Insertion			
AnySkin	1.7 ± 0.6	$\mathbf{6.7 \pm 1.5}$	$\mathbf{5.3 \pm 2.5}$
ReSkin	1.7 ± 1.2	6.0 ± 1.7	1.7 ± 1.2
DIGIT	1.7 ± 1.5	2.3 ± 0.6	1.3 ± 0.6

Additionally, we find that visuotactile policies trained with ReSkin and AnySkin have similar performance on solving the plug insertion task. However, when the skin is replaced, the performance of the ReSkin policy falls 43% to the same level as the camera-only policy, while the performance of AnySkin policies only drops by 13%. The difference is even more stark with DIGIT sensors where visuotactile policies perform significantly worse than those trained with ReSkin/AnySkin, and swapping skins results in poorer performance than vision-only policies. This demonstrates the uniqueness of AnySkin for learning generalizable visuotactile policies for precise manipulation.

5 Conclusions

In this paper, we present AnySkin, a new magnetic tactile sensor. AnySkin is versatile, self-adhering and improves on the repeatability of sensor response across different instances of the skin. Furthermore, to the best of our knowledge, AnySkin is the first sensor to demonstrate zero-shot generalization of visuotactile policies to new instances of the tactile skin. This opens the door to exciting scaling efforts like training large tactile models and real world deployment of models trained in the laboratory. Future work could also investigate simple calibration schemes or conditional learning frameworks to completely close the gap between training and swapped skin instances.

ACKNOWLEDGMENTS

Special thanks for Krishna Bodduluri, Mike Lambeta, and team from Meta AI Research for providing the DIGIT sensors for comparison. NYU authors are supported by grants from Honda, Hyundai, NSF award 2339096 and ONR awards N00014-21-1-2758 and N00014-22-1-2773. LP is supported by the Packard Fellowship.

References

[1] JR Jenner and JA1197256 Stephens. Cutaneous reflex responses and their central nervous pathways studied in man. *The Journal of physiology*, 333(1):405–419, 1982.

[2] Roland S Johansson. Sensory control of dexterous manipulation in humans. In *Hand and brain*, pages 381–414. Elsevier, 1996.

[3] Zipeng Fu, Ashish Kumar, Ananye Agarwal, Haozhi Qi, Jitendra Malik, and Deepak Pathak. Coupling vision and proprioception for navigation of legged robots. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17273–17283, June 2022.

[4] Abhishek Padalkar, Acorn Pooley, Ajinkya Jain, Alex Bewley, Alex Herzog, Alex Irpan, Alexander Khazatsky, Anant Rai, Anikait Singh, Anthony Brohan, et al. Open x-embodiment: Robotic learning datasets and rt-x models. *arXiv preprint arXiv:2310.08864*, 2023.

[5] Cheng Chi, Siyuan Feng, Yilun Du, Zhenjia Xu, Eric Cousineau, Benjamin Burchfiel, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. *arXiv preprint arXiv:2303.04137*, 2023.

[6] Homanga Bharadhwaj, Jay Vakil, Mohit Sharma, Abhinav Gupta, Shubham Tulsiani, and Vikash Kumar. Roboagent: Generalization and efficiency in robot manipulation via semantic augmentations and action chunking, 2023.

[7] Raunaq Bhirangi, Tess Hellebrekers, Carmel Majidi, and Abhinav Gupta. Reskin: versatile, replaceable, lasting tactile skins. In *5th Annual Conference on Robot Learning*, 2021.

[8] Oliver Glauser, Daniele Panozzo, Otmar Hilliges, and Olga Sorkine-Hornung. Deformation capture via soft and stretchable sensor arrays. *ACM Transactions on Graphics (TOG)*, 38(2):1–16, 2019.

[9] Te-Yen Wu, Lu Tan, Yuji Zhang, Teddy Seyed, and Xing-Dong Yang. Capacitivo: Contact-based object recognition on interactive fabrics using capacitive sensing. In *Proceedings of the 33rd annual acm symposium on user interface software and technology*, pages 649–661, 2020.

[10] Daniel Xu, Andreas Tairych, and Iain Alexander Anderson. Stretch not flex: programmable rubber keyboard. *Smart Materials and Structures*, 25, 2015.

[11] 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.

[12] Subramanian Sundaram, Petr Kellnhofer, Yunzhu Li, Jun-Yan Zhu, Antonio Torralba, and Wojciech Matusik. Learning the signatures of the human grasp using a scalable tactile glove. *Nature*, 569(7758):698–702, 2019.

[13] Tapomayukh Bhattacharjee, Advait Jain, Sarvagya Vaish, Marc D Killpack, and Charles C Kemp. Tactile sensing over articulated joints with stretchable sensors. In *2013 World Haptics Conference (WHC)*, pages 103–108. IEEE, 2013.

[14] Stefano Stassi, Valentina Cauda, Giancarlo Canavese, and Candido Fabrizio Pirri. Flexible tactile sensing based on piezoresistive composites: A review. *Sensors*, 14(3):5296–5332, 2014.

[15] Ian H Taylor, Siyuan Dong, and Alberto Rodriguez. Gelslim 3.0: High-resolution measurement of shape, force and slip in a compact tactile-sensing finger. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 10781–10787. IEEE, 2022.

[16] Wenzhen Yuan, Siyuan Dong, and Edward H Adelson. Gelsight: High-resolution robot tactile sensors for estimating geometry and force. *Sensors*, 17(12):2762, 2017.

[17] Mike Lambeta, Po-Wei Chou, Stephen Tian, Brian Yang, Benjamin Maloon, Victoria Rose Most, Dave Stroud, Raymond Santos, Ahmad Byagowi, Gregg Kammerer, et al. Digit: A novel design for a low-cost compact high-resolution tactile sensor with application to in-hand manipulation. *IEEE Robotics and Automation Letters*, 5(3):3838–3845, 2020.

[18] Julia Di, Zdravko Dugonjic, Will Fu, Tingfan Wu, Romeo Mercado, Kevin Sawyer, Victoria Rose Most, Gregg Kammerer, Stefanie Speidel, Richard E. Fan, Geoffrey Sonn, Mark R. Cutkosky, Mike Lambeta, and Roberto Calandra. Using fiber optic bundles to miniaturize vision-based tactile sensors, 2024.

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

[20] 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.

[21] Youcan Yan, Zhe Hu, Zhengbao Yang, Wenzhen Yuan, Chaoyang Song, Jia Pan, and Yajing Shen. Soft magnetic skin for super-resolution tactile sensing with force self-decoupling. *Science Robotics*, 6, 2021.

[22] Vani H Sundaram, Raunaq Bhirangi, Mark E Rentschler, Abhinav Gupta, and Tess Hellebrekers. Dragonclaw: A low-cost pneumatic gripper with integrated magnetic sensing. In *2023 IEEE International Conference on Soft Robotics (RoboSoft)*, pages 1–8. IEEE, 2023.

[23] Raunaq Bhirangi, Abigail DeFranco, Jacob Adkins, Carmel Majidi, Abhinav Gupta, Tess Hellebrekers, and Vikash Kumar. All the feels: A dexterous hand with large-area tactile sensing. *IEEE Robotics and Automation Letters*, 2023.

[24] Raunaq Bhirangi, Chenyu Wang, Venkatesh Pattabiraman, Carmel Majidi, Abhinav Gupta, Tess Hellebrekers, and Lerrel Pinto. Hierarchical state space models for continuous sequence-to-sequence modeling. In *Forty-first International Conference on Machine Learning*, 2024.

[25] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.

[26] Aadhithya Iyer, Zhuoran Peng, Yinlong Dai, Irmak Guzey, Siddhant Haldar, Soumith Chintala, and Lerrel Pinto. Open teach: A versatile teleoperation system for robotic manipulation. *arXiv preprint arXiv:2403.07870*, 2024.

[27] Siddhant Haldar, Zhuoran Peng, and Lerrel Pinto. Baku: An efficient transformer for multi-task policy learning. *arXiv preprint arXiv:2406.07539*, 2024.

[28] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

[29] 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.