

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

Raunaq Bhirangi Venkatesh Pattabiraman Enes Erciyes Yifeng Cao Tess Hellebrekers Lerrel Pinto




Fig. 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. AnySkin senses contact through distortions in magnetic field generated by magnetized iron particles in the sensing surface. The flexible surface is physically separated from its electronics, which allows for easy replaceability when damaged.

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 simplistic design of ReSkin, and decoupling the sensing electronics from the sensing interface, AnySkin simplifies integration making it 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.

INTRODUCTION

Touch sensing is widely recognized as a crucial modality for biological movement and control [9]. Unlike vision, sound, or proprioception, touch provides sensing at the point of contact, allowing agents to perceive and reason about forces and pressure. 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 sensing in conjunction with proprioception to train manipulation skills [4, 1], often ignoring touch. If touch is indeed 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 [2], 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).

Our main findings from experiments are summarized below:

- 1) AnySkin can readily be used on a variety of robots including xArm, Franka, and Leap hand.
- 2) AnySkin is compatible with ML techniques for slip detection and visuo-tactile policy learning for precise tasks such as inserting USBs.
- 3) AnySkin takes an average of 12 seconds to replace and can be reused after replacement.
- 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 the 43% drop in performance with ReSkin [2] sensors.

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




Fig. 3: Fabrication of AnySkin

Fig. 4: (a) AnySkin is made by mixing Smooth-On DragonSkin 10 Slow and MQFP-15-7(25 μ m) 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) Skins made with MQP-15-7(-80 mesh) and MQFP-15-7(25 μ m) particles. Note the concentration of particles at the surface of the former due to the larger particle size.

I. ANYSKIN: FABRICATION

AnySkin builds on ReSkin [2], a tactile skin composed of a soft magnetized skin coupled with magnetometer-based sensing circuitry. By detecting distortions in magnetic fields, ReSkin measures skin deformations caused by normal and shear forces [7, 2]. Its adaptability enables integration across various applications, from robotic hands [3] to arm sleeves and even dog shoes. AnySkin uses the same 5-magnetometer circuitry as ReSkin, while introducing key design and fabrication changes to the skin to improve durability, repeatability, and replaceability.

- Magnetizing skins post-curing using a pulse magnetizer.
- Introducing physical separation between magnetic elastomer and magnetometer circuit.
- Utilizing finer magnetic particles to achieve a more uniform particle distribution.
- Implementing a self-aligning design for reduced variability in the positioning of elastomers and circuitry.

While some of these additions have been used in isolation in prior work [3, 11], there has been little discussion on their effect on sensor response.

Applying a magnetic field during elastomer curing increases variability in the signal response. Before curing, magnetic particles are free to move through the liquid elastomer under the effect of the magnetic field. As a result, the distribution of particles is influenced by the temporal evolution of the applied magnetic field, i.e. how you move the magnets into place, which can be difficult to control when fabricating. To circumvent these disadvantages, we propose using a pulse magnetizer to magnetize the skins post-curing in line with [3], as shown in Fig. 3. The pulse magnetizer can apply a large enough magnetic field to magnetize the dipoles in the magnetic elastomer. Curing outside the influence of magnetic

fields allows for a more uniform distribution of magnetic particles through the bulk of the sensor, thereby improving magnetic field consistency.

II. EXPERIMENTS AND RESULTS

In this section, we perform extensive experiments to demonstrate the capabilities of AnySkin as a tactile sensor, and within the context of policy learning. These experiments are designed to answer the following questions:

- How do the fabrication changes outlined in Section I influence signal characteristics?
- Can AnySkin sensors be used to detect slip?
- How does AnySkin’s ease of replaceability compare with other sensors like DIGIT and ReSkin?
- How does replacing AnySkin affect the performance of learned policies, and compare with other sensors like ReSkin and DIGIT?

A. Comparison between ReSkin and AnySkin signal

To quantitatively demonstrate the effect of each of the fabrication changes listed in Section I towards improving the consistency of AnySkin, we present the following set of experiments analyzing the raw signal from the four different skins shown in Table I, tracking the progression from ReSkin to AnySkin:

1) *Effect of pulse magnetizer on signal strength:* To understand the effect of the pulse magnetizer on signal strength, we take five instances of each skin type and measure the raw signal corresponding to each instance. We average the absolute values across the three axes of the five magnetometers, and report the results in Table I. We see a significant increase in the raw magnetic field for both sets of pulse magnetized skins. This increase allows us to add a physical separation between sensing skin and the sensory electronics, which improves

TABLE I: AnySkin’s signal strength is comparable to ReSkin with lower variability across instances, and physical separation from the magnetometer electronics. Statistics computed over 5 samples of each type (PM: Pulse magnetizer, FP: finer particles, SA: self-aligning).

Experiment	ReSkin		+PM		+FP		+SA (AnySkin)	
	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 std. deviation across instances	0.54	0.87	0.34	0.12	0.21	0.15	0.12	0.10
Normalized std. deviation across 1mm misalignments	1.38	1.43	0.25	0.11	0.18	0.07	Self-aligning	




Fig. 5: Training and test locations of the target objects interacted with for plug insertion, card swiping and USB insertion (left to right). The blue region represents the extent of variation in the location of the target object, while the green-orange blocks denote held-out test configurations used for evaluation.

replaceability, as well as repeatability of the signal as discussed below.

2) *Comparison of signal consistency across skins*: To compare signal consistency across the sensing skins, we compute the standard deviation along each axis of the five magnetometers across the five instance of each skin type. To account for the larger signal strengths of the pulse magnetized skins and allow for a fairer comparison, we normalize the computed standard deviations by the mean absolute values along xy and z axes for each skin type. Aggregated statistics for the different skins are presented in Table I.

B. 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. Our model is able to detect slip on unseen objects with 92% accuracy.

C. Replaceability in Policy Learning

The most important consequence of the signal consistency and replaceability of AnySkin outlined so far, is its ability to enable policy generalization across different instances of the skin. In this section, we demonstrate the cross-instance generalizability of AnySkin across three precise manipulation

tasks. We follow this up with a comparison of the cross-instance generalizability of policies trained on DIGIT, ReSkin and AnySkin on the plug insertion task.

1) *Experimental Setup*: For our policy learning experiments, we train behavior cloning models for a set of precise manipulation tasks. Our experimental setup consists of an X-Arm 7DOF robot in a robot cage. A Meta Quest 3 and the accompanying joystick controller are used to teleoperate the robot using Open-Teach [8], an open-source teleoperation framework.

We demonstrate the replaceability of AnySkin on a set of three contact-rich manipulation tasks - Plug Insertion, Card Swiping and USB Insertion (See Fig. 5 for test variations).

2) *Model Architecture and Training*: Our policies are trained using behavior cloning. The BAKU [5] architecture is used as the policy architecture. BAKU tokenizes each input using a modality-specific encoder: image inputs from cameras and DIGITs are encoded using ResNet-18 [6] encoders, while AnySkin and ReSkin inputs are encoded using an MLP encoder. An action token is appended to the set of encoded tokens before passing the sequence through a transformer encoder, and the output corresponding to the action token is used to predict actions. We use action chunking [12] and predict the next 10 actions at every timestep. For every training setting, we train three separate models corresponding to three different seeds, and present aggregated statistics on 10 policy rollouts.

3) *Evaluating cross-instance generalizability*: To investigate the replaceability of AnySkin in the context of policy learning, we evaluate behavior cloning policies trained using a single instance of AnySkin on a new instance. Table II presents a comparison between policy performance with the original

and swapped skins for each of the precise, contact-rich tasks outlined above.

4) *Comparison across sensors*: To better contextualize the significance of this result, we present a replaceability comparison with DIGIT [10] and ReSkin [2] sensors. We collect two additional datasets of 96 demonstration trajectories each for the plug insertion task with these sensors similar to AnySkin. Replaceability is evaluated by swapping the training skin for a new skin during evaluation as outlined in the previous section. Success rates from 10 evaluations across three seeds for each setting are reported in Table II.

TABLE II: Success rates (out of 10) for policies when swapping out tactile skins. All statistics computed over 3 training seeds

Task	Cameras only	Cameras + Skin	
		Original skin	Swapped skin
<i>Cross-instance generalization</i>			
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
<i>Comparison across sensors – Plug Insertion</i>			
AnySkin	1.7 ± 0.6	6.7 ± 1.5	5.3 ± 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

Based on these results, we find that visuotactile policies trained with ReSkin and AnySkin have similar performance on solving the plug insertion task, while DIGIT policies are unable to capture the minute interactions. However, when the sensor instance 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%. This transferability is evidence of AnySkin’s superior signal consistency, and is a significant boost to scaling efforts like training large tactile models as well as real world deployment of models trained in the laboratory.

REFERENCES

- [1] 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. URL <https://arxiv.org/abs/2309.01918>.
- [2] Raunaq Bhirangi, Tess Hellebrekers, Carmel Majidi, and Abhinav Gupta. Reskin: versatile, replaceable, lasting tactile skins. In *5th Annual Conference on Robot Learning*, 2021.
- [3] 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.
- [4] 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.
- [5] Siddhant Haldar, Zhuoran Peng, and Lerrel Pinto. Baku: An efficient transformer for multi-task policy learning. *arXiv preprint arXiv:2406.07539*, 2024.
- [6] 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.
- [7] Tess Hellebrekers, Oliver Kroemer, and Carmel Majidi. Soft magnetic skin for continuous deformation sensing. *Advanced Intelligent Systems*, 1(4):1900025, 2019.
- [8] Aadithya 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.
- [9] Roland S Johansson. Sensory control of dexterous manipulation in humans. In *Hand and brain*, pages 381–414. Elsevier, 1996.
- [10] 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.
- [11] 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.
- [12] 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.