

DexSkin: High-Coverage Conformable Robotic Skin for Learning Contact-Rich Manipulation

Author Names Omitted for Anonymous Review.

Abstract: Human skin offers a rich tactile sensing stream over large, contoured areas, but replicating this in robotics remains challenging. We present DexSkin, a soft, conformable capacitive skin that provides localized, calibratable sensing and adapts to varied geometries. Integrated on gripper fingers with near-full coverage, DexSkin enables learning-based manipulation, transfers across sensor instances via calibration, and supports real-world online RL. *This work has been published at CoRL 2025 (main track).*

Keywords: tactile sensing, contact-rich manipulation



Figure 1: **DexSkin sensors enable contact-rich robotic manipulation policies.** We demonstrate the applications of DexSkin in in-hand pen reorientation, box packaging, and berry transporting tasks. In each frame, we visualize tactile readings from DexSkin as heatmaps, where brighter readings correspond to larger forces. DexSkin’s high-coverage sensing capabilities enable it to provide contact signals for a variety of robotic manipulation tasks. At the same time, it possesses desirable properties for learning-based manipulation systems.

1 Introduction

Tactile feedback is essential for dexterous manipulation, yet most robotic sensors remain rigid, low-coverage, and poorly adaptable to complex surfaces. We introduce DexSkin, a soft capacitive skin that conforms to end-effector geometry, offers high-density localized sensing, and can be calibrated across hardware instances. Each taxel is individually addressable, allowing localization of contacts across surfaces, and the sensor can be calibrated for consistent readings across hardware instances.

In this work, we describe DexSkin’s design and integration with a robotic fingertip, then evaluate its utility for robot learning. We test whether DexSkin’s coverage expands learnable manipulation tasks, assess calibration for model transfer, and demonstrate real-world reinforcement learning on a delicate object picking task. The results highlight DexSkin’s practicality for contact-rich, data-driven manipulation.

2 Related Work

Tactile sensing in robotics. Robotic tactile sensors generally fall into three categories: vision-based, magnetic, and electrically addressable. Vision-based sensors such as GelSight [1, 2], GelSlim [3], DIGIT [4], and DenseTact [5, 6] capture elastomer deformations with cameras, but are bulky and difficult to miniaturize [7]. Magnetic sensors such as uSkin [8], ReSkin [9], and AnySkin [10]

measure flux changes, but provide sparse signals and require data-heavy calibration. Electrically addressable sensors include resistive [11, 12], impedance (e.g., BioTac [13]), and capacitive [14, 15, 16, 17]. Capacitive designs offer high sensitivity and low power consumption but often rely on costly microfabrication. Our approach also employs a capacitive mechanism, but enables conformable coverage and low-cost, rapid prototyping with individually addressable taxels.

Learning robotic manipulation with tactile sensors. Tactile sensing has been integrated into learning robotic manipulation via reinforcement learning [18, 19, 20], predictive models [21, 22, 23, 4], and imitation learning [24, 25, 26, 27]. Since simulating contacts is difficult [28, 29, 30, 31], many works rely on real-world data. We follow this paradigm, focusing on tasks requiring broad coverage and simultaneous contacts.

3 The DexSkin Framework

We introduce the main components of the DexSkin framework. While applicable to many morphologies, we focus on fingertip-shaped gripper jaws for a parallel gripper, shown in Figure 2.

Low-cost, high-performance, conformable DexSkin sensors. Electrically addressable tactile sensors remain rare in robotics [32, 26] due to sensitivity, cost, and integration challenges. DexSkin is a stretchable capacitive sheet fabricated for under \$10 per pair (at 1,000 units) using accessible tools. It provides nearly full-finger coverage, high sensitivity, and individually addressable taxels that can resolve multiple simultaneous contacts.

Tailorable fabrication. Fabrication begins with computer-aided electrode pattern design, which determines taxel layout and size. The electrode layers are deposited onto a thin elastomer substrate and sealed into a flexible sheet. This single continuous piece can be conformed around both hemispherical domes and cylindrical finger surfaces, yielding dense tactile coverage (Figure 2).

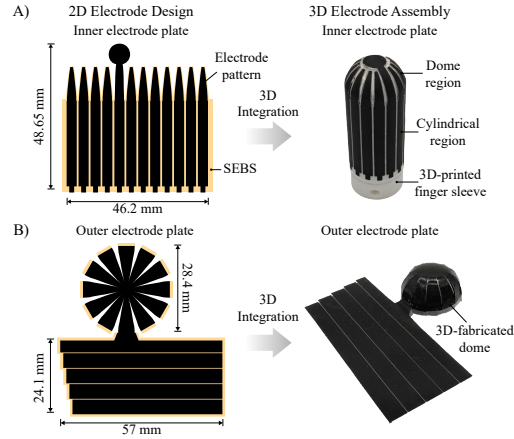


Figure 2: **Designs of electrode patterns.** The 2D pattern design of the inner plate (top left) and outer plate (bottom left) electrodes on the soft SEBS substrate and its appearance after its vertical traces have been conformed onto the dome and cylindrical circumference of the finger sleeve (top right and bottom right). Note that SEBS is transparent and can be hard to distinguish in the post-assembly photo (right).

We are committed to open-sourcing the detailed fabrication instructions and materials for DexSkin.

4 Evaluating DexSkin for Learning Robotic Manipulation

We evaluate DexSkin for learning robotic manipulation from the following perspectives:

1. Can DexSkin’s coverage and tailorability enable robots to learn a range of manipulation tasks?
2. Can calibrating DexSkin allow learned policies to be transferred across sensor instances?
3. Is DexSkin suitable for applications to learning robot behaviors *online*?

4.1 Learning Manipulation with Expanded Coverage and Tailorability

We test whether DexSkin’s coverage and resolution improve learned manipulation through two tasks. Policies are trained from 50 teleoperated demonstrations using GELLO [33] and diffusion policies [34, 35]. The first task is in-hand pen reorientation, where the robot must rely on tactile and proprioceptive feedback to reorient a pen and remain robust to human perturbations. The second is box packaging, where the robot must secure a container lid with an elastic band that may be intact




Figure 3: Examples of box packaging task rollouts. (*Top*) With an initially provided perforated band, the DexSkin policy correctly detects and discards the perforated band based on the tactile reading (col. 1), retrieves the backup (col. 2), and wraps it around the box (cols. 3-5). We visualize DexSkin sensor readings as if the viewer is facing the robot, looking at the gripper. (*Bottom*) With an initially provided unperforated band, the policy without tactile information unnecessarily replaces the band (cols. 1-2) and then fails to package the box as the band slips off the gripper fingers (cols. 3-5).

Model	Pen reorientation		Box: Non-Perforated Band		Box: Perforated Band	
	No perturb	Perturb	Select	Wrap	Select	Wrap
No tactile	19/20	0/20	0/20	0/20	19/20	6/20
DIGIT [4]	19/20	0/20	20/20*	20/20*	0/20	0/20
DexSkin (ablated)	12/20	0/20	19/20	14/20	1/20	0/20
DexSkin (ours)	19/20	19/20	18/20	17/20	19/20	15/20

Table 1: Pen reorientation and box packaging success rates. Policies with full DexSkin information achieve the most consistent performance across tasks requiring high sensing coverage. For box packaging, we report success in two cumulative stages for evaluation only: (1) determining the correct elastic band to use and (2) physically wrapping the band around the box. *Placing the band around DIGIT’s bulkier geometry elongates the band more, making wrapping it around the box easier.

or perforated, requiring it to detect band properties through dorsal fingertip sensing and adapt its behavior accordingly.

Comparative evaluation. To test our hypothesis, we compare the following settings:

- *DexSkin (ours)*: Uses the full 120 taxels of DexSkin readings, proprioception, and wrist camera RGB images. Note wrist camera images are provided for the box task only.
- *No tactile*: A baseline identical to *DexSkin (ours)* but excluding all tactile information.
- Ablations include spatially pooled readings for the pen task (*DexSkin (spatial pooling)*), mimicking low-resolution sensors (eg. load cells), and inner-column readings for the box task (*DexSkin (inner cols. only)*), approximating flat sensors. These test the importance of DexSkin’s resolution and coverage.
- *DIGIT*: Uses a pair of DIGIT [4] sensors as gripper fingers and as input to the policy, comparing our system to existing and commercially available sensors.

Pen task results. We test two settings: the training setup and one with human perturbation that rotates the pen back to horizontal. Results (Table 1) show most policies succeed without perturbation, but only the full DexSkin policy maintains success under disturbance. The no-tactile baseline repeats the same motion regardless, spatial pooling often fails to estimate pose even unperturbed, and DIGIT misses many contacts outside its limited sensing region.

Box packaging results. As shown in Table 1 and Figure 3, only the full DexSkin policy consistently succeeds with both intact and perforated bands. Other policies adopt fixed strategies, misusing bands that are visually indistinguishable, while DexSkin correctly selects based on perceived tension. Full DexSkin input also improves wrapping success, completing 86% of rollouts versus 32% (No tactile) and 70% (inner cols. only).

4.2 Calibration and Model Transfer Across Sensor Instances

Tactile sensors are difficult to manufacture consistently, as optical, magnetic, and resistive designs often produce signal variations across units [10, 36]. Because learned models are highly sensitive to such shifts, replacing sensors can render prior training unusable. We introduce a calibration procedure for DexSkin to mitigate these issues and enable policy transfer across hardware instances.

Calibration procedure. We develop two calibration protocols. First, DexSkin is enclosed in a 3D-printed airtight chamber with an Ecoflex 00-50 membrane, and internal pressure is ramped from 0–6 psi to impose uniform stress; fitting each taxel’s response to an exponential curve provides both a forward mapping ($\Delta C/C_0 \rightarrow \text{pressure}$) and an inverse mapping to align new data with legacy outputs. Second, DexSkin is mounted on a motorized stage with a force gauge, recording three loading–unloading cycles per taxel to map outputs to normal force.

Policy transfer experiment. We revisit the perturbed pen reorientation task using a policy trained on one sensor pair (*Source*) and transferred to swapped fingers (*Target*), simulating sensor replacement. Results (Table 2) show that DexSkin transfers reasonably well without calibration and improves further with it, while DIGIT policies fail entirely when gel windows are swapped due to appearance changes.

Sensor configuration	Stage 1	Stage 2
Source sensors	20/20	20/20
Target sensors (no calib.)	17/20	12/20
Target sensors (calib.)	18/20	16/20
Source sensors (DIGIT)	20/20	0/20
Target sensors (DIGIT)	0/20	0/20

Table 2: Pen reorientation policy performance when transferred across sensor hardware. We report successes across two stages: (1) successfully reorienting the pen the first time and (2) detecting and fixing human perturbation.

4.3 Real-World Online Robot Learning with DexSkin

Beyond imitation learning, tactile sensing can provide reward signals for RL, but many sensors lack durability and clear reward definitions. We evaluate DexSkin on a delicate blueberry grasping task (where a single excessive force causes failure), augmenting a non-tactile imitation policy with DexSkin and training a residual policy via soft actor-critic [37, 38], with rewards penalizing excessive forces, large actions, and failed grasps.

Figure 4 shows that the residual policy refines the base actions to produce gentler grasps. Without tactile input, the base policy crushes berries, while the DexSkin-enhanced policy handles them successfully. This demonstrates that DexSkin can adapt policies trained without tactile data and provide natural reward signals without extra classifiers or calibration.

5 Conclusion

We present DexSkin, a soft tactile sensor that conforms to varied geometries for localized, high-coverage sensing. Integrated on a gripper, DexSkin improves imitation learning for contact-rich tasks, enables calibration-based transfer across hardware, and provides reward feedback in real-world reinforcement learning, advancing toward practical tactile sensors with human-skin-like coverage and sensitivity.

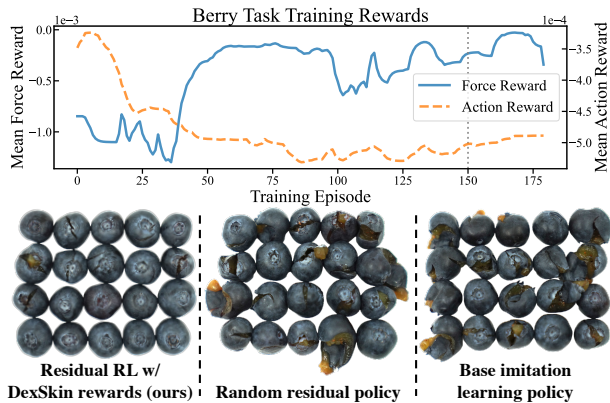


Figure 4: (Top) Residual policy learning curves showing reduced DexSkin-output forces above threshold (*force reward*) and balanced action costs (*action reward*). Training used a faux berry until episode 150, then real berries for 30 episodes. (Bottom) Berries after grasp and transport; the finetuned residual policy causes much less damage.

References

- [1] M. K. Johnson, F. Cole, A. Raj, and E. H. Adelson. Microgeometry capture using an elastomeric sensor. *ACM Transactions on Graphics (TOG)*, 30(4):1–8, 2011.
- [2] R. Li, R. Platt, W. Yuan, A. ten Pas, N. Roscup, M. A. Srinivasan, and E. Adelson. Localization and manipulation of small parts using gelsight tactile sensing. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2014.
- [3] E. Donlon, S. Dong, M. Liu, J. Li, E. Adelson, and A. Rodriguez. GelSlim: A High-Resolution, Compact, Robust, and Calibrated Tactile-sensing Finger. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018.
- [4] M. Lambeta, P.-W. Chou, S. Tian, B. Yang, B. Maloon, V. R. Most, D. Stroud, R. Santos, A. Byagowi, G. Kammerer, D. Jayaraman, and R. Calandra. DIGIT: A Novel Design for a Low-Cost Compact High-Resolution Tactile Sensor With Application to In-Hand Manipulation. *IEEE Robotics and Automation Letters*, 2020.
- [5] W. K. Do, B. Jurewicz, and M. Kennedy. Densetact 2.0: Optical tactile sensor for shape and force reconstruction. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 12549–12555. IEEE, 2023.
- [6] W. K. Do, A. K. Dhawan, M. Kitzmann, and M. Kennedy. Densetact-mini: An optical tactile sensor for grasping multi-scale objects from flat surfaces. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6928–6934. IEEE, 2024.
- [7] J. Di, Z. Dugonjic, W. Fu, T. Wu, R. Mercado, K. Sawyer, V. R. Most, G. Kammerer, S. Speidel, R. E. Fan, et al. Using fiber optic bundles to miniaturize vision-based tactile sensors. *arXiv preprint arXiv:2403.05500*, 2024.
- [8] T. P. Tomo, M. Regoli, A. Schmitz, L. Natale, H. Kristanto, S. Somlor, L. Jamone, G. Metta, and S. 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. doi:[10.1109/LRA.2018.2812915](https://doi.org/10.1109/LRA.2018.2812915).
- [9] R. Bhirangi, T. Hellebrekers, C. Majidi, and A. Gupta. Reskin: versatile, replaceable, lasting tactile skins. In *5th Annual Conference on Robot Learning*, 2021.
- [10] R. Bhirangi, V. Pattabiraman, E. Erciyes, Y. Cao, T. Hellebrekers, and L. Pinto. Anyskin: Plug-and-play skin sensing for robotic touch. *arXiv preprint arXiv:2409.08276*, 2024.
- [11] S. Sundaram, P. Kellnhofer, Y. Li, J.-Y. Zhu, A. Torralba, and W. Matusik. Learning the signatures of the human grasp using a scalable tactile glove. *Nature*, 569(7758), 2019.
- [12] Y. Luo, Y. Li, P. Sharma, W. Shou, K. Wu, M. Foshey, B. Li, T. Palacios, A. Torralba, and W. Matusik. Learning human–environment interactions using conformal tactile textiles. *Nature Electronics*, 4(3):193–201, 2021.
- [13] R. S. J. Nicholas Wettels, Veronica J. Santos and G. E. Loeb. Biomimetic tactile sensor array. *Advanced Robotics*, 22(8):829–849, 2008. doi:[10.1163/156855308X314533](https://doi.org/10.1163/156855308X314533).
- [14] J. Yang, S. Luo, X. Zhou, J. Li, J. Fu, W. Yang, and D. Wei. Flexible, tunable, and ultrasensitive capacitive pressure sensor with microconformal graphene electrodes. *ACS Applied Materials & Interfaces*, 11(16):14997–15006, 2019. doi:[10.1021/acsami.9b02049](https://doi.org/10.1021/acsami.9b02049).
- [15] N. Bai, L. Wang, Q. Wang, J. Deng, Y. Wang, P. Lu, J. Huang, G. Li, Y. Zhang, J. Yang, K. Xie, X. Zhao, and C. Guo. Graded intrafillable architecture-based iontronic pressure sensor with ultra-broad-range high sensitivity. *Nature Communications*, 11:209, 01 2020. doi:[10.1038/s41467-019-14054-9](https://doi.org/10.1038/s41467-019-14054-9).

- [16] Z. Luo, J. Chen, Z. Zhu, L. Li, Y. Su, W. Tang, O. M. Omisore, L. Wang, and H. Li. High-resolution and high-sensitivity flexible capacitive pressure sensors enhanced by a transferable electrode array and a micropillar–pvdf film. *ACS Applied Materials & Interfaces*, 13(6):7635–7649, 2021. doi:[10.1021/acsami.0c23042](https://doi.org/10.1021/acsami.0c23042). PMID: 33539065.
- [17] Y.-C. Huang, Y. Liu, M. Chao, H.-C. Cheng, Q. He, H. Wu, C. Wang, C.-Y. Lin, Y. Huang, and X. Duan. Sensitive pressure sensors based on conductive microstructured air-gap gates and two-dimensional semiconductor transistors. *Nature Electronics*, 3:59–69, 01 2020. doi:[10.1038/s41928-019-0356-5](https://doi.org/10.1038/s41928-019-0356-5).
- [18] M. A. Lee, Y. Zhu, K. Srinivasan, P. Shah, S. Savarese, L. Fei-Fei, A. Garg, and J. Bohg. Making sense of vision and touch: Self-supervised learning of multimodal representations for contact-rich tasks. In *International Conference on Robotics and Automation (ICRA)*, 2019.
- [19] S. Dong, D. K. Jha, D. Romeres, S. Kim, D. Nikovski, and A. Rodriguez. Tactile-rl for insertion: Generalization to objects of unknown geometry. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2021.
- [20] S. Kim and A. Rodriguez. Active extrinsic contact sensing: Application to general peg-in-hole insertion. In *International Conference on Robotics and Automation (ICRA)*, 2022.
- [21] R. Calandra, A. Owens, M. Upadhyaya, W. Yuan, J. Lin, E. H. Adelson, and S. Levine. The feeling of success: Does touch sensing help predict grasp outcomes? In *Conference on Robot Learning*, 2017.
- [22] R. Calandra, A. Owens, D. Jayaraman, J. Lin, W. Yuan, J. Malik, E. H. Adelson, and S. Levine. More Than a Feeling: Learning to Grasp and Regrasp Using Vision and Touch. *IEEE Robotics and Automation Letters*, 2018.
- [23] S. Tian, F. Ebert, D. Jayaraman, M. Mudigonda, C. Finn, R. Calandra, and S. Levine. Manipulation by feel: Touch-based control with deep predictive models. In *International Conference on Robotics and Automation (ICRA)*, 2019.
- [24] H. Li, Y. Zhang, J. Zhu, S. Wang, M. A. Lee, H. Xu, E. Adelson, L. Fei-Fei, R. Gao, and J. Wu. See, Hear, and Feel: Smart Sensory Fusion for Robotic Manipulation. In *Conference on Robot Learning*, 2023.
- [25] K. Yu, Y. Han, Q. Wang, V. Saxena, D. Xu, and Y. Zhao. Mimictouch: Leveraging multi-modal human tactile demonstrations for contact-rich manipulation. In *8th Annual Conference on Robot Learning*, 2024.
- [26] B. Huang, Y. Wang, X. Yang, Y. Luo, and Y. Li. 3d-vitac: Learning fine-grained manipulation with visuo-tactile sensing. *arXiv preprint arXiv:2410.24091*, 2024.
- [27] T. Lin, Y. Zhang, Q. Li, H. Qi, B. Yi, S. Levine, and J. Malik. Learning visuotactile skills with two multifingered hands. *arXiv:2404.16823*, 2024.
- [28] H. Qi, B. Yi, S. Suresh, M. Lambeta, Y. Ma, R. Calandra, and J. Malik. General In-hand Object Rotation with Vision and Touch. In *Conference on Robot Learning*, 2023.
- [29] S. Suresh, H. Qi, T. Wu, T. Fan, L. Pineda, M. Lambeta, J. Malik, M. Kalakrishnan, R. Calandra, M. Kaess, J. Ortiz, and M. Mukadam. Neural feels with neural fields: Visuo-tactile perception for in-hand manipulation. *arXiv:2312.1346*, 2023.
- [30] Y. Yuan, H. Che, Y. Qin, B. Huang, Z.-H. Yin, K.-W. Lee, Y. Wu, S.-C. Lim, and X. Wang. Robot Synesthesia: In-Hand Manipulation with Visuotactile Sensing. *arXiv:2312.01853*, 2023.
- [31] J. Xu, Y. Jia, D. Yang, P. Meng, X. Zhu, Z. Guo, S. Song, and M. Ciocarlie. Tactile-based object retrieval from granular media. *arXiv preprint arXiv:2402.04536*, 2023.

- [32] S. H. Huang, M. Zambelli, J. Kay, M. F. Martins, Y. Tassa, P. M. Pilarski, and R. Hadsell. Learning gentle object manipulation with curiosity-driven deep reinforcement learning. *arXiv preprint arXiv:1903.08542*, 2019.
- [33] P. Wu, Y. Shentu, Z. Yi, X. Lin, and P. Abbeel. Gello: A general, low-cost, and intuitive teleoperation framework for robot manipulators. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 12156–12163. IEEE, 2024.
- [34] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song. Diffusion policy: Visuomotor policy learning via action diffusion. In *Proceedings of Robotics: Science and Systems (RSS)*, 2023.
- [35] C. Chi, Z. Xu, S. Feng, E. Cousineau, Y. Du, B. Burchfiel, R. Tedrake, and S. Song. Diffusion policy: Visuomotor policy learning via action diffusion. *The International Journal of Robotics Research*, 2024.
- [36] S. Stassi, V. Cauda, G. Canavese, and C. F. Pirri. Flexible tactile sensing based on piezoresistive composites: A review. *Sensors*, 14(3):5296–5332, 2014. ISSN 1424-8220. doi:10.3390/s140305296.
- [37] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pages 1861–1870. Pmlr, 2018.
- [38] A. Raffin, A. Hill, A. Gleave, A. Kanervisto, M. Ernestus, and N. Dormann. Stable-baselines3: Reliable reinforcement learning implementations. *Journal of Machine Learning Research*, 22(268):1–8, 2021. URL <http://jmlr.org/papers/v22/20-1364.html>.