TactileLab: Efficient Shear-Sensitive Tactile Simulation for Sim-to-Real Multi-Finger Dexterous Manipulation

Anonymous Authors (Paper under double-blind review)

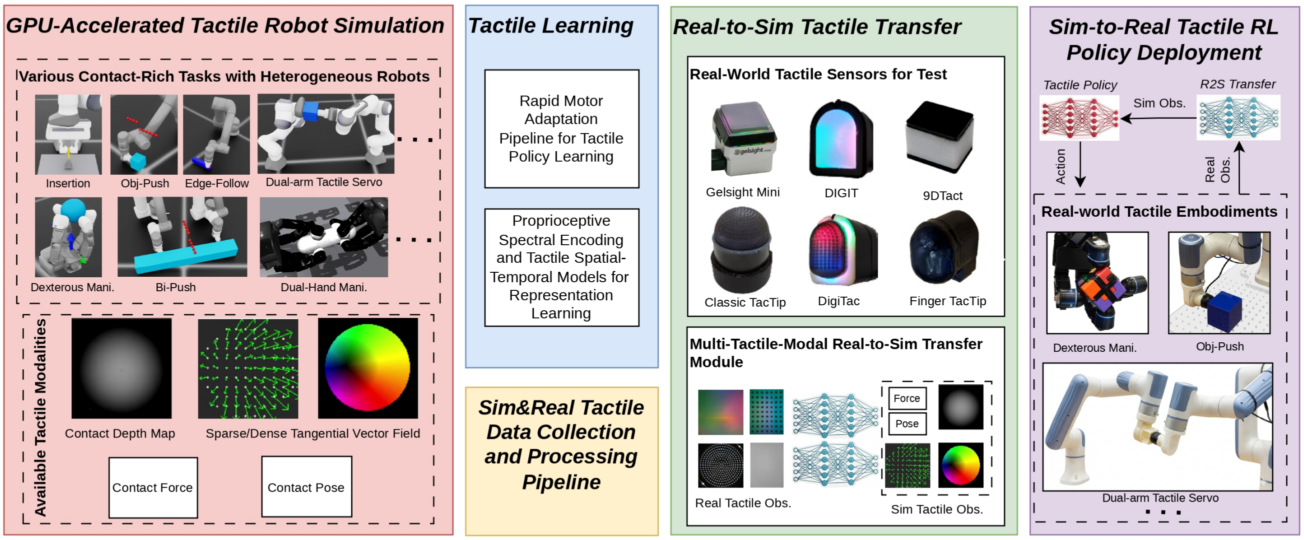

Figure 1. TactileLab is an end-to-end framework for efficient shear-sensitive tactile robot learning. We developed a GPU-accelerated tactile simulation based on IsaacLab for diverse contact-rich tasks and heterogeneous robots, tactile learning modules for policy and representation learning, a multi-modal real-to-sim tactile transfer pipeline for multiple real tactile sensors, and sim-to-real tactile RL policy deployment on real robotic systems. This unified design supports scalable training, efficient simulated tactile data collection, and transfer from simulation to real-world tactile manipulation.

I. INTRODUCTION

Tactile sensing is especially important for dexterous manipulation with multi-finger hands, where contact is distributed across multiple fingertips and evolves continuously during object reorientation and stabilization. In these settings, control depends not only on normal contact geometry but also on tangential interaction, slip tendency, and disturbance recovery. However, scalable tactile simulation methods still struggle to represent such shear-sensitive interaction efficiently.

Most existing tactile simulators fall into two categories. Deformable simulators can model compliant contact with high fidelity [1], but are often too slow for large-scale reinforcement learning. Rigid-body simulators are fast and scalable, but usually provide only depth-like or force-like signals [2]–[6]. As a result, they are well suited for normal-contact reasoning but less effective for tasks in which shear, slip, or tangential motion matters. This limitation is important for manipulation problems such as in-contact trajectory following, insertion, and in-hand object manipulation in the presence of external perturbations. In these tasks, a depth-only observation cannot fully describe how contact evolves over time. Different contact states may produce similar instantaneous depth signals while leading to very different future outcomes. This causes partial observability and makes policy learning more difficult. Although a recent state-of-the-art tactile simulator (TacSL) can model shear force fields [7], it relies on prior object shape knowledge, which limits generalization to unseen objects and reduces its applicability to real-world deployment.

To address this gap, we introduce *TactileLab*, an efficient tactile simulation and learning framework that augments rigid-body tactile sensing with a dense tangential displacement field. Rather than explicitly simulating elastomer deformation, the framework represents contact using two complementary components: normal-contact depth and a shear-sensitive in-plane tangential-displacement-vector field. This produces a compact and transferable tactile observation that is efficient to compute and suitable for large-scale reinforcement learning.

While TactileLab supports a broad range of contact-rich manipulation tasks, this paper particularly emphasizes multi-finger dexterous in-hand manipulation, which is central to the workshop theme and provides a demanding testbed for tactile sensing, large-scale simulation, and sim-to-real transfer.

The contributions of this work are:

1. We develop *TactileLab*, a unified tactile learning framework in IsaacLab for GPU-parallelized multi-modal tactile simulation, large-scale RL training, Real2Sim tactile transfer, and Sim2Real tactile policy deployment across multiple robotic embodiments.

2. We propose a new shear-sensitive tactile representation based on contact depth and tangential displacement for efficient learning of contact dynamics beyond depth-only sensing.

3. We present preliminary evidence that this richer tactile representation improves challenging multi-finger in-hand dexterous manipulation under randomized initial hand orientation and external force/torque perturbations.

II. METHODS

A. TactileLab Framework

TactileLab is an end-to-end framework for tactile robot learning built on IsaacLab [8]. It unifies efficient tactile simulation, contact-rich task design, tactile reinforcement learning, real-to-sim tactile transfer, and sim-to-real deployment within a single GPU-parallelized pipeline. The framework is designed to support heterogeneous robotic embodiments, including single-arm, dual-arm, gripper-based, and multi-finger platforms, while maintaining a common interface for tactile observation, policy learning, and deployment.

At the system level, TactileLab contains four components. First, it provides a highly efficient tactile simulation stack with multiple tactile modalities. Second, it offers a unified library of tactile robotic tasks spanning manipulation regimes from non-prehensile exploration to dexterous in-hand control. Third, it includes learning tools for tactile reinforcement learning, policy distillation, and tactile data collection. Fourth, it supports real-to-sim tactile transfer and real-world deployment, enabling policies trained in simulation to run with real tactile observations. Together, these components make TactileLab a scalable platform for studying contact-rich manipulation under a shared simulation-to-deployment workflow.

B. Shear-Sensitive Tactile Representation

The key representation in TactileLab augments contact depth with a dense tangential displacement field defined on the tactile surface. Contact depth captures normal interaction, while the tangential field captures local in-plane motion at each spatial location. Together, they form a compact shear-sensitive tactile observation that encodes both the geometry of contact and its temporal evolution.

This design is motivated by the limitation of depth-only tactile sensing in shear-dependent tasks. A depth map can indicate where contact occurs, but it does not explicitly describe how the contact patch moves, drifts, or deforms along the tangent plane. As a result, different contact states may produce similar instantaneous depth observations while corresponding to different future dynamics. By incorporating tangential displacement, the observation becomes more informative for tasks that require tracking relative motion, regulating shear, or reacting to incipient instability.

Importantly, this representation is defined as an efficient abstraction rather than a full deformable-material simulation. Instead of explicitly modeling elastomer mechanics, the framework represents shear interaction through tangential movement cues that remain compatible with fast rigid-body simulation. This preserves scalability for reinforcement learning while providing a transferable tactile interface for real-world sensing.

C. Multi-Finger Dexterous In-Hand Manipulation

We place particular emphasis on multi-finger in-hand dexterous manipulation, since it provides one of the most demanding settings for tactile robot learning. In contrast to simpler contact-rich tasks, in-hand manipulation requires the policy to reason over distributed, multi-point, and continuously evolving contacts while maintaining stable object control. This makes it a natural testbed for evaluating whether richer tactile sensing can improve dexterous manipulation performance.

Table I
Preliminary results on multi-finger in-hand rotation. Higher is better.

Method	Return ↑	Ep. Len. ↑
HORA (Prop.) [9]	48.3	421.4
AnyRotate (F+P) [10]	78.2	509.6
Ours (Depth+TM)	**103.5**	**552.7**

We conducted preliminary experiments on in-hand object rotation using a rich tactile sensing representation that combines contact depth maps with tangential motion fields. To further increase task difficulty, we randomly initialize the hand orientation at the beginning of each episode and apply random external force and torque perturbations to the manipulated object during execution. These modifications create a substantially more challenging setting by increasing contact uncertainty, disturbance sensitivity, and recovery difficulty.

Table I compares three observation settings: proprioception only (Prop.) used by HORA [9], force and contact pose (F+P) used by AnyRotate [10], and our richer tactile representation based on contact depth and tangential motion (Depth+TM). As shown in the table, the proposed method outperforms both HORA and AnyRotate in terms of mean return and mean episode length. Our method achieves a mean return of **103.5** and a mean episode length of **552.7**, compared with 48.3 and 421.4 for HORA, and 78.2 and 509.6 for AnyRotate. This corresponds to a **114.3%** improvement in mean return and a **31.2%** improvement in mean episode length over HORA, and a **32.4%** improvement in mean return and a **8.5%** improvement in mean episode length over AnyRotate. Although these results are preliminary, they provide encouraging evidence that rich shear-sensitive tactile observations can improve multi-finger dexterous manipulation in challenging disturbance-rich conditions.

D. Broader Contact-Rich Task Suite

Beyond the multi-finger in-hand manipulation setting emphasized in this paper, TactileLab supports a broad range of contact-rich tactile manipulation tasks spanning multiple robotic embodiments and interaction regimes. The task suite is organized as a unified progression from simpler contact-rich behaviours to more dynamic and dexterous manipulation. It includes established single-arm tactile tasks such as object pushing, edge following, and surface following [2], [3]; bi-manual tactile tasks such as bi-pushing, bi-reorienting, and bi-gathering [11]; shear-sensitive tactile tracking tasks such as dual-arm object tracking [12]; and precision assembly tasks such as peg placement, peg insertion, and bolt-on-nut alignment [7]. Together, these tasks allow TactileLab to cover non-prehensile interaction, dual-arm coordination, shear-sensitive tactile servoing, and precision contact alignment within a single framework.

An important contribution of TactileLab is that these tasks are implemented within a GPU-parallelized simulation and training pipeline. In contrast to earlier tactile-learning setups that were typically developed task-by-task with more limited parallelism, our framework supports large numbers of environments running concurrently, substantially accelerating reinforcement learning for contact-rich tasks. The same parallelized design also makes simulated tactile data collection far more efficient, enabling scalable generation of training data for policy learning, tactile representation learning, and real-to-sim transfer. As a result, TactileLab contributes not only a broad task suite, but also a significantly more efficient infrastructure for tactile robotics research.

This broader task suite is important for two reasons. First, it shows that the same tactile abstraction can support a wide range of embodiments and contact regimes, rather than being tied only to dexterous hand manipulation. Second, it helps identify where shear-sensitive tactile observations matter most. In relatively lower-dimensional settings such as edge following or object pushing, tangential cues improve interaction awareness and robustness. In more demanding settings such as bimanual reorientation and precision insertion, they become increasingly important because control depends on continuously evolving relative motion at the contact interface.

E. Real2Sim Tactile Transfer and Sim2Real RL Policy Deployment

To bridge the gap between simulated and real tactile observations, we proposed a new real-to-sim tactile transfer module that maps real images into the simulated multi-tactile-modal representation used during policy training. This generated observation space reduces the mismatch between sim and real, allowing policies learned in simulation to operate directly on real robotic systems. Since the proposed tactile representation is defined at the level of contact force, contact pose, contact depth map and tangential movement field, it is naturally suited to transfer from compliant vision-based tactile sensors that exhibit observable surface displacement under shear. This makes it possible to use simulated policies as the basis for real-world tactile control without requiring full high-fidelity soft-body image simulation during training. Combined with GPU-parallelized simulation and large-scale tactile data generation, this transfer pipeline enables an efficient end-to-end workflow from simulated training to real-world deployment.

III. Conclusion

TactileLab presents a unified framework for efficient shear-sensitive tactile robot learning. By combining GPU-parallelized tactile simulation, a shear-aware tactile representation, and real-to-sim tactile transfer, the framework supports scalable learning for contact-rich manipulation from simulation to real deployment.

We particularly highlight preliminary results on multi-finger dexterous in-hand manipulation, where rich tactile sensing based on contact depth and tangential motion improves performance over strong existing baselines in a more challenging setting with randomized hand initialization and external force/torque perturbations. These results suggest that efficient shear-sensitive tactile simulation can be a useful foundation for advancing sensing, skill learning, and sim-to-real transfer in dexterous multi-finger robotic hands.

References

[1] Xuyang Zhang, Jiaqi Jiang, Zhuo Chen, Yongqiang Zhao, Tianqi Yang, Daniel Fernandes Gomes, Jianan Wang, and Shan Luo. Simtac: A physics-based simulator for vision-based tactile sensing with biomorphic structures. *Cyborg and Bionic Systems*, 7:0510, 2026.

[2] Alex Church, John Lloyd, Raia Hadsell, and Nathan F Lepora. Tactile sim-to-real policy transfer via real-to-sim image translation. In *Conference on Robot Learning*, pages 1645–1654, 2022.

[3] Yijiong Lin, John Lloyd, Alex Church, and Nathan F Lepora. Tactile gym 2.0: Sim-to-real deep reinforcement learning for comparing low-cost high-resolution robot touch. *IEEE Robotics and Automation Letters*, 7(4):10754–10761, 2022.

[4] Daniel Fernandes Gomes, Paolo Paoletti, and Shan Luo. Generation of gelsight tactile images for sim2real learning. *IEEE Robotics and Automation Letters*, 6(2):4177–4184, 2021.

[5] Shaoxiong Wang, Mike Lambeta, Po-Wei Chou, and Roberto Calandra. Tacto: A fast, flexible, and open-source simulator for high-resolution vision-based tactile sensors. *IEEE Robotics and Automation Letters*, 7(2):3930–3937, 2022.

[6] Quan Khanh Luu, Nhan Huu Nguyen, et al. Simulation, learning, and application of vision-based tactile sensing at large scale. *IEEE Transactions on Robotics*, 39(3):2003–2019, 2023.

[7] Iretiayo Akinola, Jie Xu, Jan Carius, Dieter Fox, and Yashraj Narang. Tacsl: A library for visuotactile sensor simulation and learning. *IEEE Transactions on Robotics*, 2025.

[8] Mayank Mittal, Pascal Roth, James Tigue, Antoine Richard, Octi Zhang, Peter Du, Antonio Serrano-Muñoz, Xinjie Yao, René Zurbrügg, Nikita Rudin, Lukasz Wawrzyniak, Milad Rakhsha, Alain Denzler, Eric Heiden, Ales Borovicka, Ossama Ahmed, Iretiayo Akinola, Abrar Anwar, Mark T. Carlson, Ji Yuan Feng, Animesh Garg, Renato Gasoto, Lionel Gulich, Yijie Guo, M. Gussert, Alex Hansen, Mihir Kulkarni, Chenran Li, Wei Liu, Viktor Makoviychuk, Grzegorz Malczyk, Hammad Mazhar, Masoud Moghani, Adithyavairavan Murali, Michael Noseworthy, Alexander Poddubny, Nathan Ratliff, Welf Rehberg, Clemens Schwarke, Ritvik Singh, James Latham Smith, Bingjie Tang, Ruchik Thaker, Matthew Trepte, Karl Van Wyk, Fangzhou Yu, Alex Millane, Vikram Ramasamy, Remo Steiner, Sangeeta Subramanian, Clemens Volk, CY Chen, Neel Jawale, Ashwin Varghese Kuruttukulam, Michael A. Lin, Ajay Mandlekar, Karsten Patzwaldt, John Welsh, Huihua Zhao, Fatima Anes, Jean-Francois Lafleche, Nicolas Moënne-Loccoz, Soowan Park, Rob Stepinski, Dirk Van Gelder, Chris Amevor, Jan Carius, Jumyung Chang, Anka He Chen, Pablo de Heras Ciechomski, Gilles Daviet, Mohammad Mohajerani, Julia von Muralt, Viktor Reutskyy, Michael Sauter, Simon Schirm, Eric L. Shi, Pierre Terdiman, Kenny Vilella, Tobias Widmer, Gordon Yeoman, Tiffany Chen, Sergey Grizan, Cathy Li, Lotus Li, Connor Smith, Rafael Wiltz, Kostas Alexis, Yan Chang, David Chu, Linxi "Jim" Fan, Farbod Farshidian, Ankur Handa, Spencer Huang, Marco Hutter, Yashraj Narang, Soha Pouya, Shiwei Sheng, Yuke Zhu, Miles Macklin, Adam Moravanszky, Philipp Reist, Yunrong Guo, David Hoeller, and Gavriel State. Isaac lab: A gpu-accelerated simulation framework for multi-modal robot learning. *arXiv preprint arXiv:2511.04831*, 2025.

[9] Haozhi Qi, Ashish Kumar, Roberto Calandra, Yi Ma, and Jitendra Malik. In-hand object rotation via rapid motor adaptation. In *Conference on Robot Learning*, pages 1722–1732. PMLR, 2023.

[10] Max Yang, Chenghua Lu, Alex Church, Yijiong Lin, Chris Ford, Haoran Li, Efi Psomopoulou, David AW Barton, and Nathan F Lepora. Anyrotate: Gravity-invariant in-hand object rotation with sim-to-real touch. *arXiv preprint arXiv:2405.07391*, 2024.

[11] Yijiong Lin, Alex Church, Max Yang, Haoran Li, John Lloyd, Dandan Zhang, and Nathan F. Lepora. Bi-touch: Bimanual tactile manipulation with sim-to-real deep reinforcement learning. *IEEE Robotics and Automation Letters*, 8(9):5472–5479, 2023.

[12] John Lloyd and Nathan F Lepora. Pose-and-shear-based tactile servoing. *The International Journal of Robotics Research*, 43(7):1024–1055, 2024.