

TactileLab: Efficient Multimodal Tactile Simulation for Shear-Sensitive Sim2Real Dexterous Robotic Manipulation

Yijiong Lin, Nathan F. Lepora

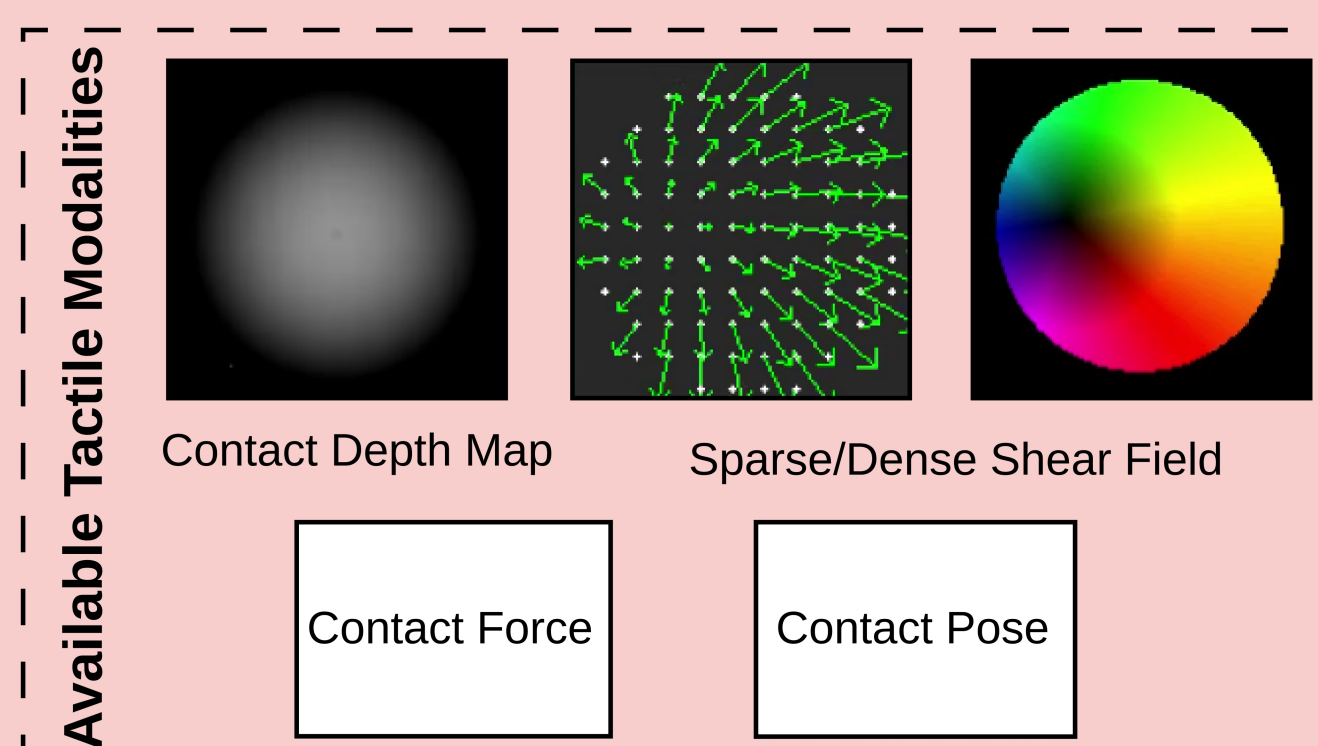

Department of Engineering Mathematics and Bristol Robotics Laboratory, University of Bristol, U.K.

{yijiong.lin, n.lepora}@bristol.ac.uk

Overview and Contributions

TactileLab: GPU-Accelerated Tactile Simulation

Various Contact-Rich Tasks with Different Robots

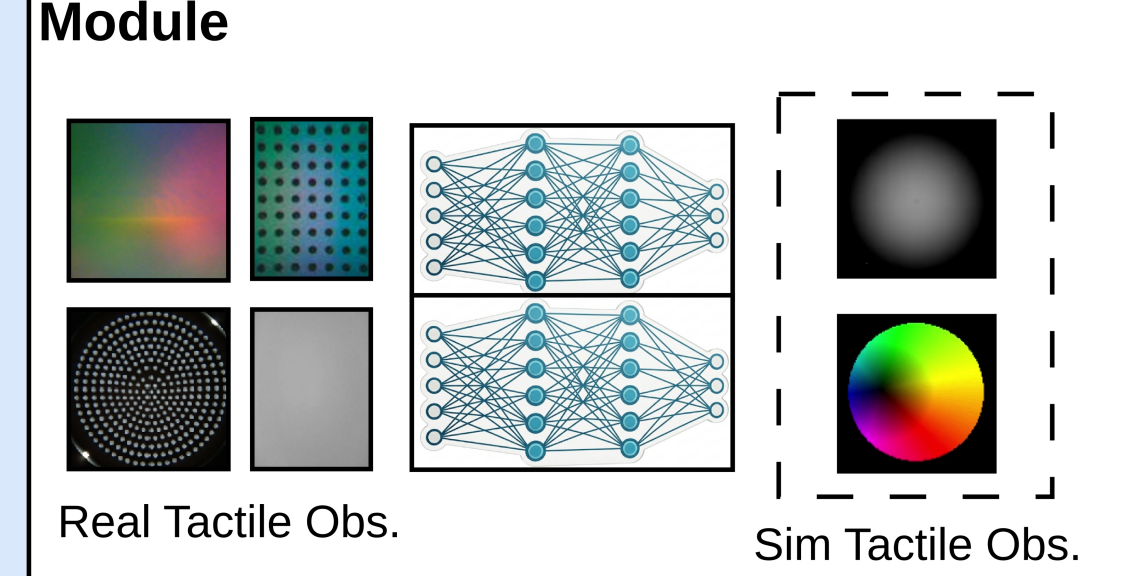


Real-to-Sim Tactile Transfer


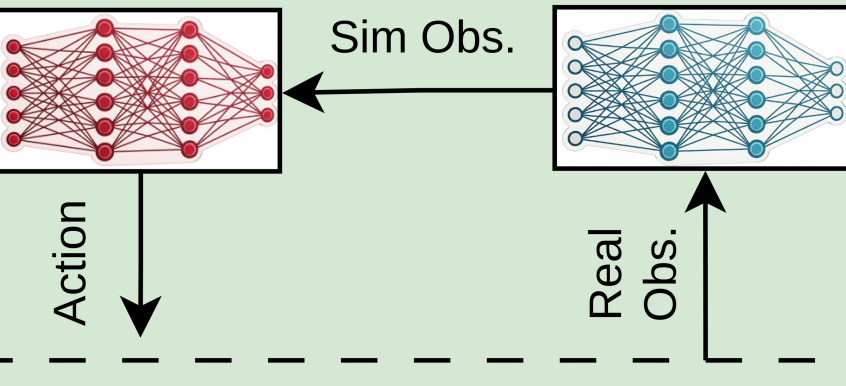
Real-World Tactile Sensors



Multi-Tactile-Modal Real-to-Sim Transfer Module

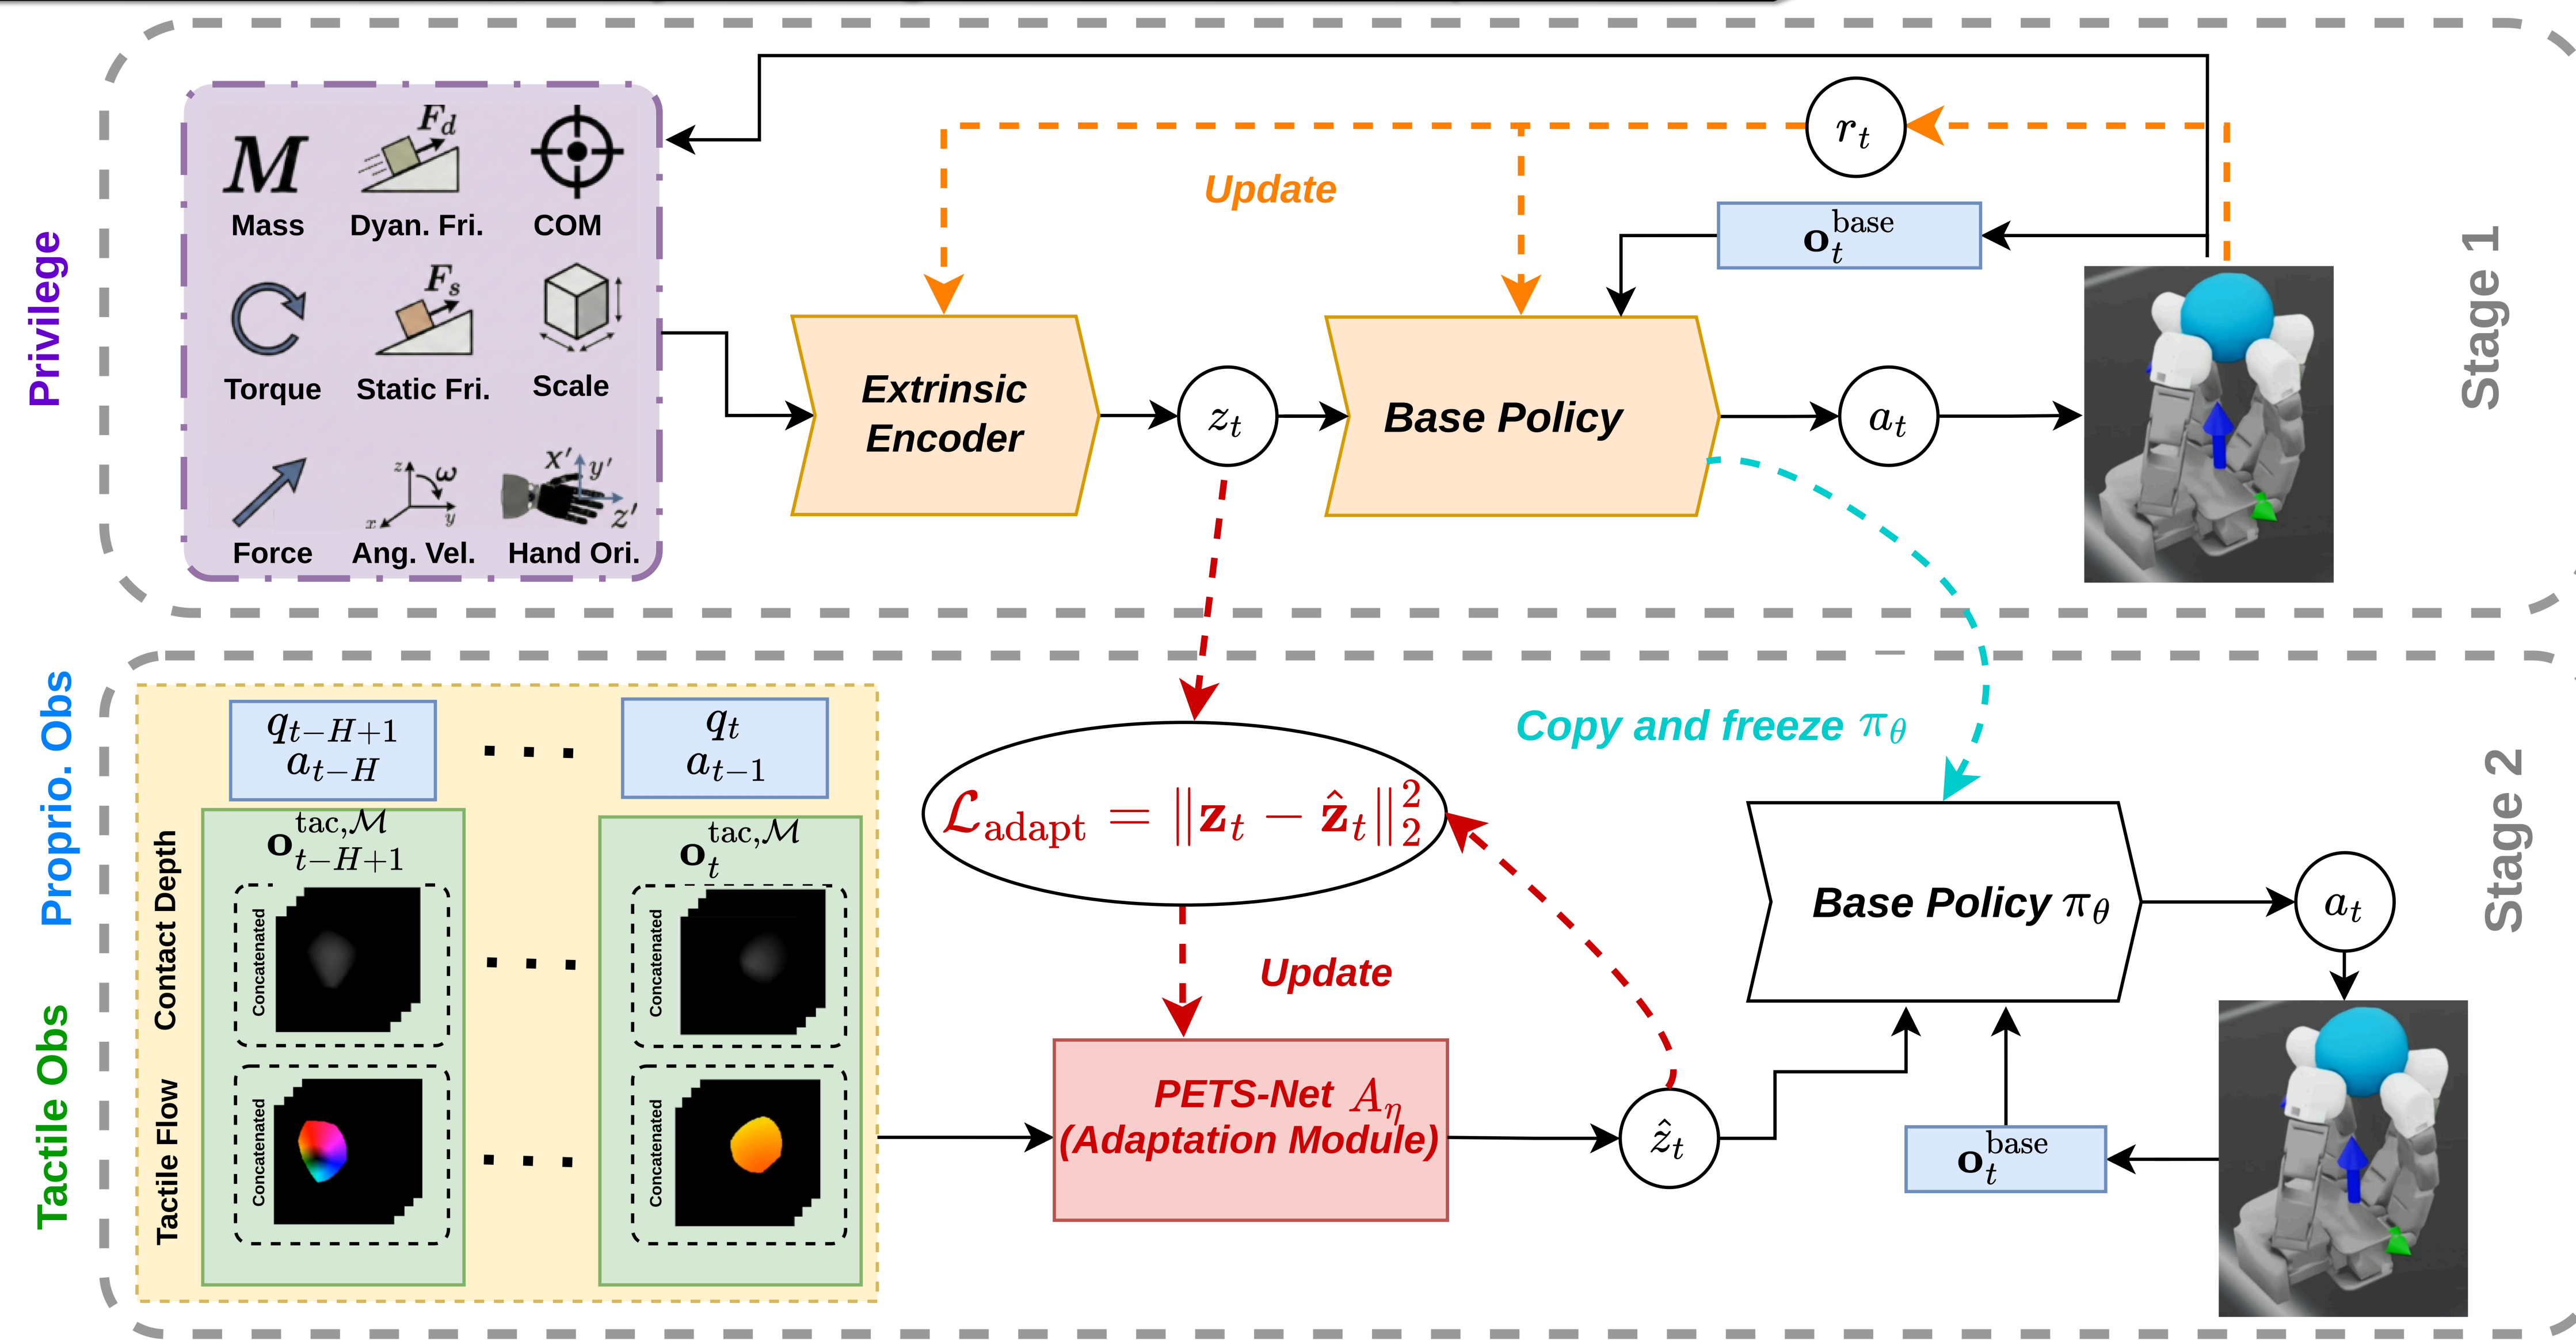


Sim-to-Real Tactile RL Policy Deployment



1. A GPU-based and scalable multimodal tactile simulation framework
2. A new simulated tactile modality for shear-relevant contact interaction: Optical-based Tactile Flow
3. A diverse suite of tactile robotic manipulation environments
4. A sample-efficient policy architecture for proprioceptive-history and tactile-observation fusion
5. An extended real-to-sim tactile transfer method for multimodal tactile observations
6. Real-world deployment of TactileLab-trained policies

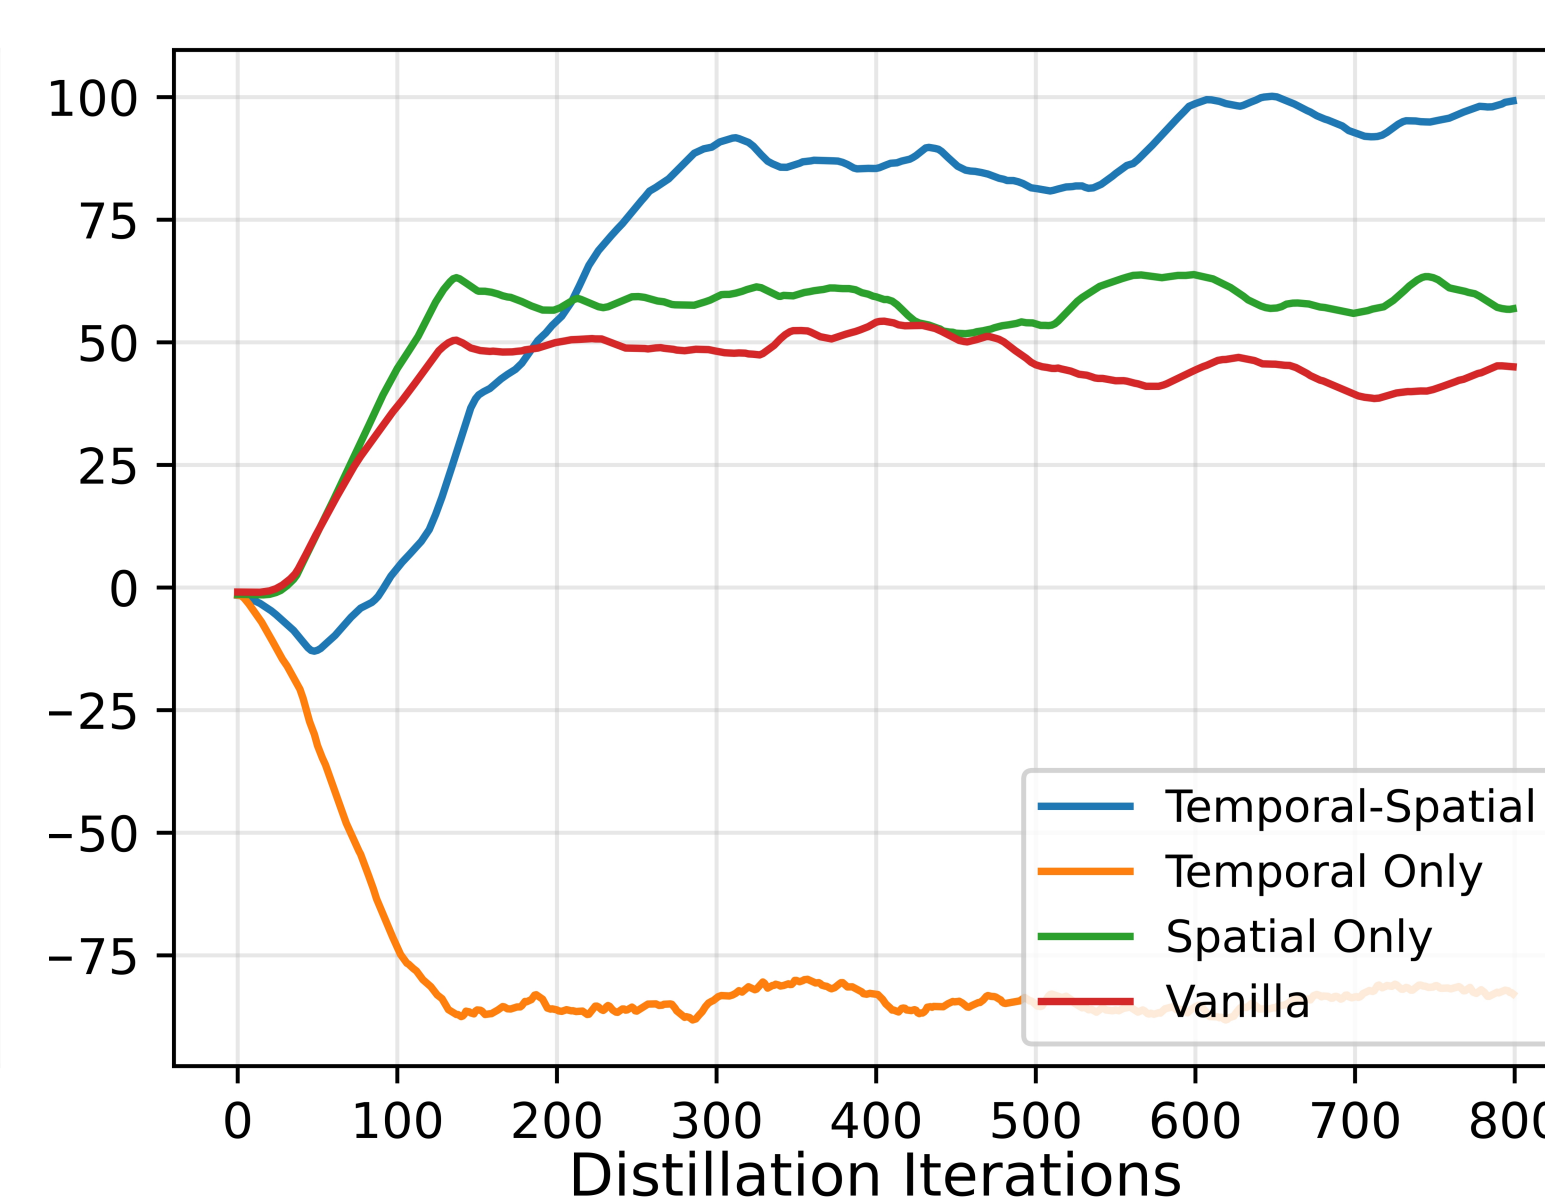
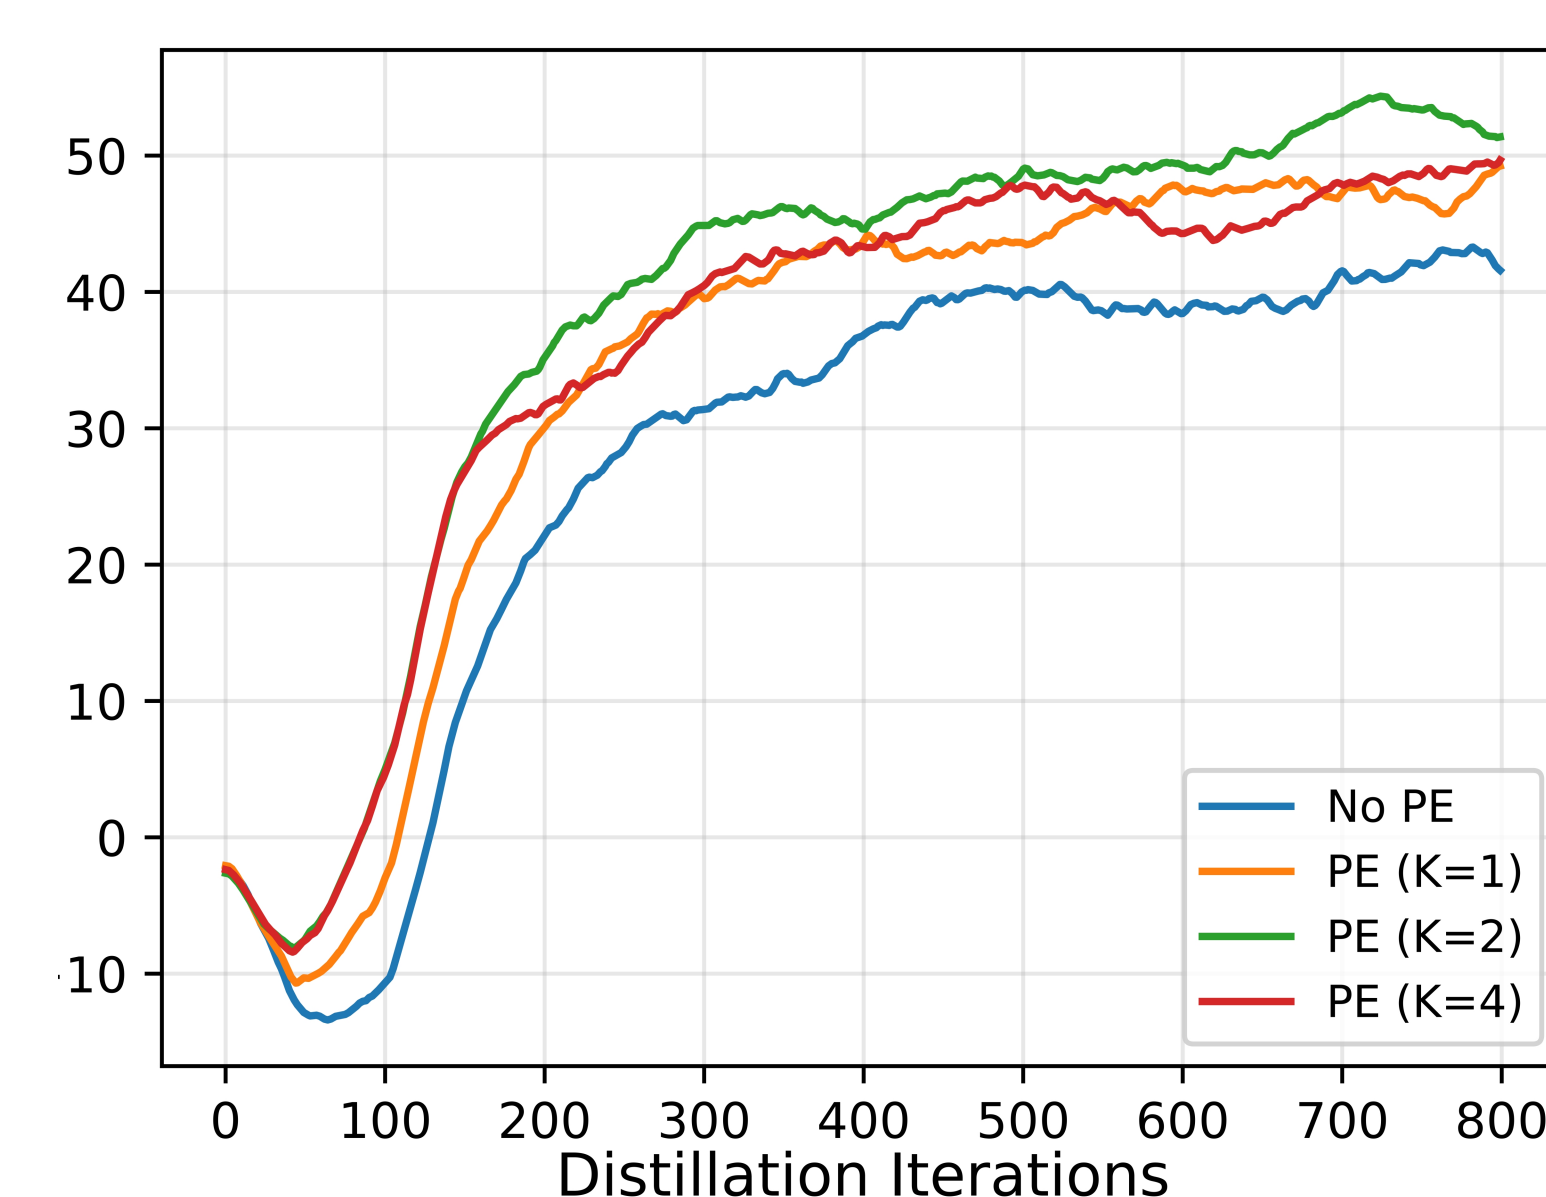
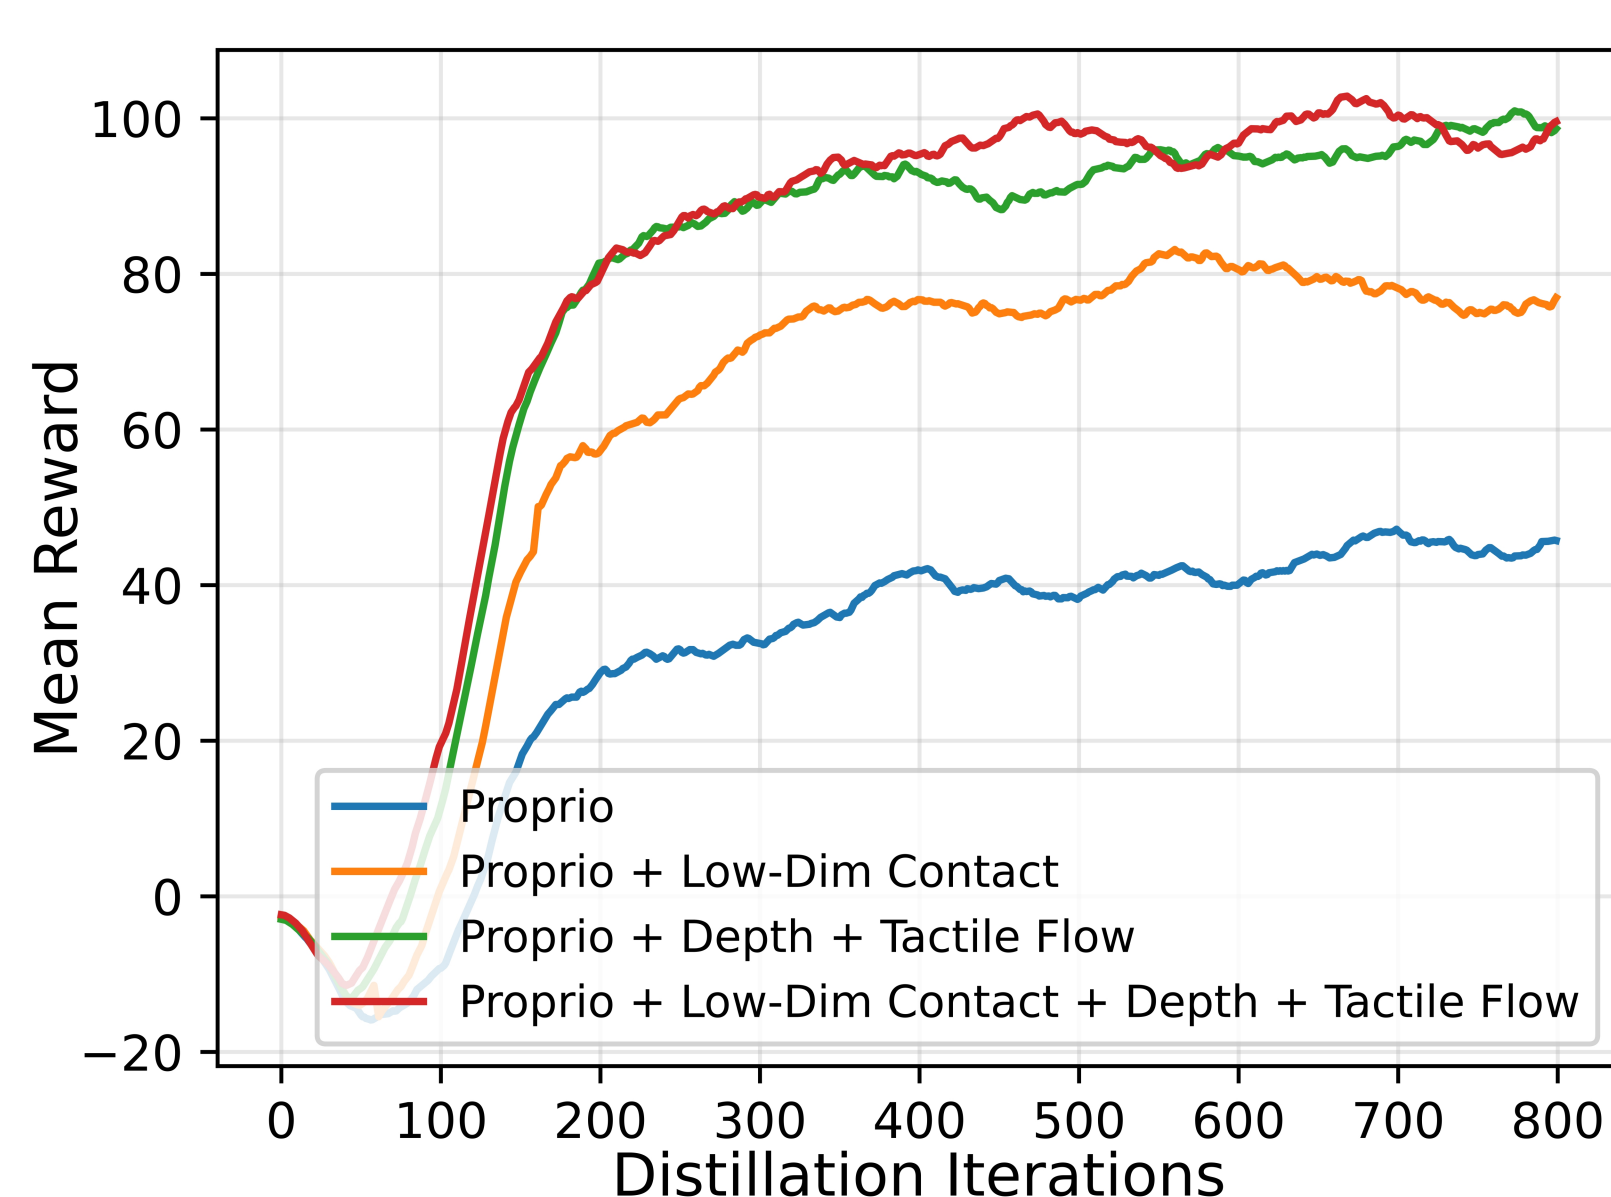
Multimodal Tactile Policy Learning for Dexterous Manipulation



The framework uses a two-stage teacher–student learning pipeline. In Stage 1, a privileged teacher policy is trained in simulation using privilege information that is not available on the real robot. The extrinsic encoder maps these privileged variables to a latent representation, which is used by the base policy to learn contact-rich manipulation behaviours.

In Stage 2, the base policy is copied and frozen. The proposed **PETS-Net** replaces the privileged encoder and predicts the latent representation from real-world-available observations, including proprioceptive history and tactile observations such as contact depth and tactile flow. PETS-Net is trained to match the teacher latent representation, enabling the student policy to use dense multimodal tactile feedback without requiring privileged simulation states during deployment.

Preliminary Experimental Results



- Left: **High-dimensional tactile observations** improve student distillation. Depth and tactile flow outperform proprioception-only [2] and low-dimensional contact [3] baselines, while combining low-dimensional contact with depth and flow gives the best performance.
- Middle: **Positional Encoding (PE)** module improves learning from proprioceptive/action histories. Increasing the frequency bands generally improves distillation efficiency compared with no positional encoding.
- Right: The **Temporal-Spatial (TS)** architecture is essential for high-dimensional tactile inputs. Temporal-only processing of stacked images fails to distill the teacher effectively, while combining spatial tactile encoding with temporal aggregation gives the strongest result.

References

- [1] Iretoiyo Akinola, Jie Xu, Jan Carius, Dieter Fox, and Yashraj Narang. Tacsl: A library for visuotactile sensor simulation and learning. *IEEE Transactions on Robotics*, 2025.
- [2] Haozhi Qi, Ashish Kumar, Roberto Calandra, Yi Ma, and Jitendra Malik. In-hand object rotation via rapid motor adaptation. *In Proceedings of The 6th Conference on Robot Learning*, 2023.
- [3] M. Yang, C. Lu, A. Church, Y. Lin, C. Ford, H. Li, E. Psomopoulou, D. AW Barton, and N. F. Lepora. Anyrotate: Gravity-invariant in-hand object rotation with sim-to-real touch. *In Proceedings of The 7th Conference on Robot Learning*, 2024.