DyWA: Dynamics-adaptive World Action Model for Generalizable Non-prehensile Manipulation

Jiangran Lyu[1,2], Ziming Li[1,2], Xuesong Shi[2], Chaoyi Xu[2], Yizhou Wang[1,†], He Wang[1,2,†]
[1]Center on Frontiers of Computing Studies, School of Computer Science, Peking University [2]Galbot
Accepted by ICCV 2025. Project Page: https://pku-epic.github.io/DyWA/

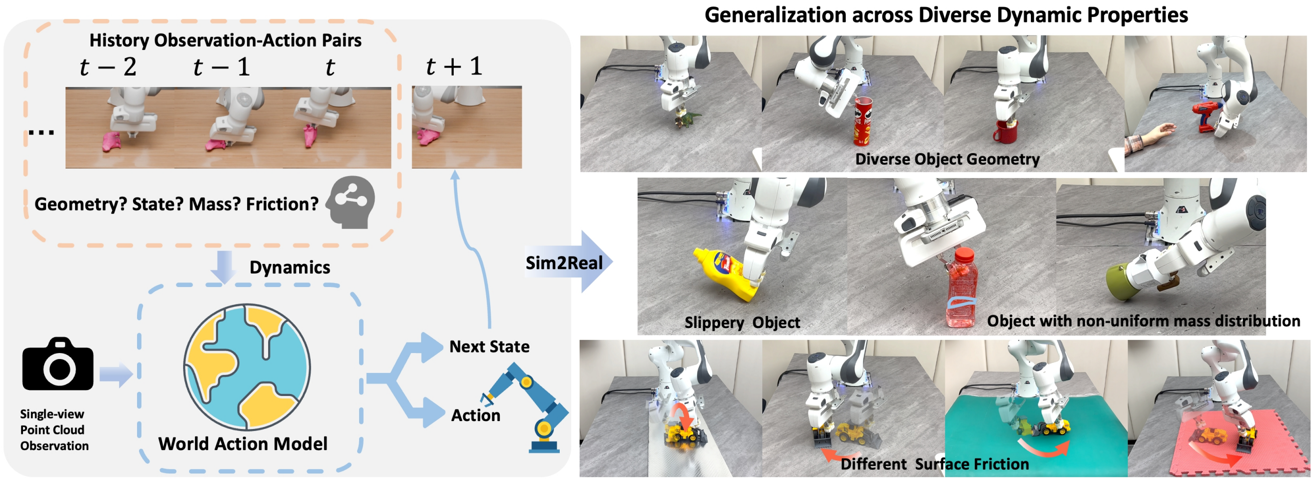

Figure 1. **Illustration of the high-level idea and generalization ability of DyWA.** Given a target object's 6D pose and *single-view* object point cloud, our non-prehensile manipulation policy aims to rearrange the object without grasping. **Left**: Our key insight is to enhance action learning by jointly predicting future states while adapting to dynamics from historical trajectories. **Right**: After being trained in simulation, our policy achieves zero-shot sim-to-real transfer and generalizes across diverse dynamic properties, including variations in object geometry, object physical property (e.g., slipperiness and non-uniform mass distribution), and surface friction.

1. abstract

Non-prehensile manipulation—such as pushing, sliding, and toppling—extends robotic capabilities beyond traditional grasping, enabling task execution under geometric, clutter, or workspace constraints. While planning-based methods [2–5] have shown success, they rely on precise object properties (e.g., mass, friction, CAD models), which are rarely available in the real world.

Recent learning-based approaches [6] shift toward end-to-end policy learning from visual input, demonstrating stronger generalization. For instance, HACMan [7] and CORN [1] exploit vision-based RL or distillation to acquire contact-rich skills. However, these methods remain geometry-centric and shows poor robustness under dynamic variations such as friction or mass changes.

To achieve generalization across dynamic variations, we argue that contact-rich manipulation fundamentally requires world modeling: policies must not only output actions but also internalize how interactions shape future states. Under this lens, existing teacher-student distillation frameworks fall short—while the privileged RL teacher can exploit full dynamics, the student policy trained from partial observations suffers due to (1) single-view partial point cloud observation, (2) Markovian policies collapsing over multiple dynamics, and (3) weak supervision limited to action imitation.

To address this, we introduce Dynamics-adaptive World Action Model (DyWA). DyWA explicitly integrates world modeling into action learning:(i) a dynamics adaptation module infers latent physical properties from observation-action history, (ii) action prediction is reformulated as joint prediction of actions and future states, providing richer supervision, and (iii) Feature-wise Linear Modulation (FiLM) bridges inferred dynamics with policy learning.

We benchmark DyWA against strong baselines on CORN, varying camera and tracking settings. DyWA improves success by 31.5% in simulation, and in real-world tests achieves 68% success across diverse geometries, frictional conditions, and mass distributions (e.g., half-filled bottles). We further demonstrate DyWA's compatibility with VLMs for challenging thin/wide-object scenarios.

References

[1] Yoonyoung Cho, Junhyek Han, Yoontae Cho, and Beomjoon Kim. Corn: Contact-based object representation for nonprehensile manipulation of general unseen objects. In *12th International Conference on Learning Representations, ICLR 2024*. International Conference on Learning Representations, ICLR, 2024. 1

[2] Igor Mordatch, Zoran Popović, and Emanuel Todorov. Contact-invariant optimization for hand manipulation. In *Proceedings of the ACM SIGGRAPH/Eurographics symposium on computer animation*, pages 137–144, 2012. 1

[3] João Moura, Theodoros Stouraitis, and Sethu Vijayakumar. Non-prehensile planar manipulation via trajectory optimization with complementarity constraints. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 970–976. IEEE, 2022.

[4] Michael Posa, Cecilia Cantu, and Russ Tedrake. A direct method for trajectory optimization of rigid bodies through contact. *The International Journal of Robotics Research*, 33 (1):69–81, 2014.

[5] William Yang and Michael Posa. Dynamic on-palm manipulation via controlled sliding. *arXiv preprint arXiv:2405.08731*, 2024. 1

[6] Xiang Zhang, Siddarth Jain, Baichuan Huang, Masayoshi Tomizuka, and Diego Romeres. Learning generalizable pivoting skills. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5865–5871. IEEE, 2023. 1

[7] Wenxuan Zhou, Bowen Jiang, Fan Yang, Chris Paxton, and David Held. Hacman: Learning hybrid actor-critic maps for 6d non-prehensile manipulation. In *Conference on Robot Learning*, pages 241–265. PMLR, 2023. 1