Supplementary Material for "Spatially Generalizable Mobile Manipulation via Adaptive Experience Selection and Dynamic Imagination"

Ping Zhong[1,2], **Liangbai Liu**[1], **Bolei Chen**[1*], **Tao Wu**[1], **Jiazhi Xia**[1,2], **Chaoxu Mu**[3], **Jianxin Wang**[1]

[1]School of Computer Science and Engineering, Central South University

[2]Xiangjiang Laboratory

[3]School of Artificial Intelligence, Anhui University

{ping.zhong, liangbailiu, boleichen, 8102221321, xiajiazhi}@csu.edu.cn, cxmu@tju.edu.cn, jxwang@mail.csu.edu.cn

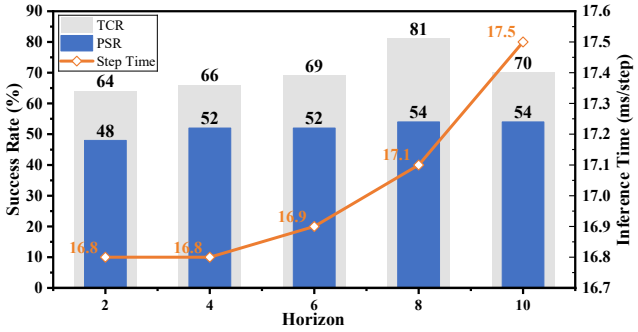

Figure 1: Success rate metrics (TCR and PSR) and single-step inference times when using different horizon parameters H.

1 Studies on H and Computational Efficiency

Considering that RSSM-based MPFP inevitably introduces additional computational cost, we investigate performance boosts induced by different dynamic imagination horizons H. Accordingly, we also report the computational efficiency. Fig. 1 illustrates the trend of the success rate metrics (TCR and PSR) and the single-step inference time with respect to different horizons H. By weighing MM performance and computational cost, we find that $H = 8$ is optimal. Setting $H = 8$ allows our MM policy to make decisions at a frequency of ∼58 Hz. The single-step inference times are measured on a computer with an NVIDIA GeForcw RTX 3090 graphics card and an AMD® Ryzen 9 9900x CPU with 12 physical cores and 24 logical cores.

2 Sim-to-Real Experiments

We deploy our method to a real robotic platform consisting of a Husky A200[1] mobile base and a UR5 manipulator[2] to verify its practicality. The robot's upper computer is a laptop with an Intel Core i9-13900HX CPU and GeForce RTX 4060 GPU. The robot uses a RealSense D435 RGB-D camera and YOLO V7 to identify target objects. We use ROS[3] to organize the hardware and software resources of the robotic system. Considering the security, we set the maximum speed and decision frequency of the mobile base to 0.3 m/s and 1 Hz, respectively. Our real-world robotics experiment considers the water delivery task of grabbing a water bottle and delivering it to a specified goal position, the process of which is shown in Fig. 2. Notably, our MM policy allows to accomplish this task efficiently in a vehicle-arm synergistic manner, instead of performing navigation and stationary manipulation separately. The sim-to-real experiment demonstrates the practicality of our approach. Please see our attached video for a more intuitive MM process.

*Corresponding Authors.

[1]https://wiki.ros.org/Robots/Husky

[2]https://github.com/ros-industrial/universal_robot

[3]https://wiki.ros.org/

3 More Experimental Configurations and Visualizations

The home-scene-oriented, warehouse-oriented, and dynamic-scene-oriented experimental configurations are shown in Fig. 3, Fig. 4, and Fig. 5, respectively. In addition, these figures also illustrate the trend of TCR metrics with the number of training episodes in different experimental configurations, respectively. These curves reflect the superior sample efficiency and MM performance of our method compared to strong baselines including N^2M^2 [Honerkamp *et al.*, 2023], Dreamer V3-base MM policy [Hafner *et al.*, 2025], and TD-MPC2-based MM policy [Hansen *et al.*, 2024].

References

[Hafner *et al.*, 2025] Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. Mastering diverse control tasks through world models. *Nature*, pages 1–7, 2025.

[Hansen *et al.*, 2024] Nicklas Hansen, Hao Su, and Xiaolong Wang. Td-mpc2: Scalable, robust world models for continuous control. In *International Conference on Learning Representations (ICLR)*, 2024.

[Honerkamp *et al.*, 2023] Daniel Honerkamp, Tim Welschehold, and Abhinav Valada. N^2m^2: Learning navigation for arbitrary mobile manipulation motions in unseen and dynamic environments. *IEEE Transactions on Robotics*, 2023.

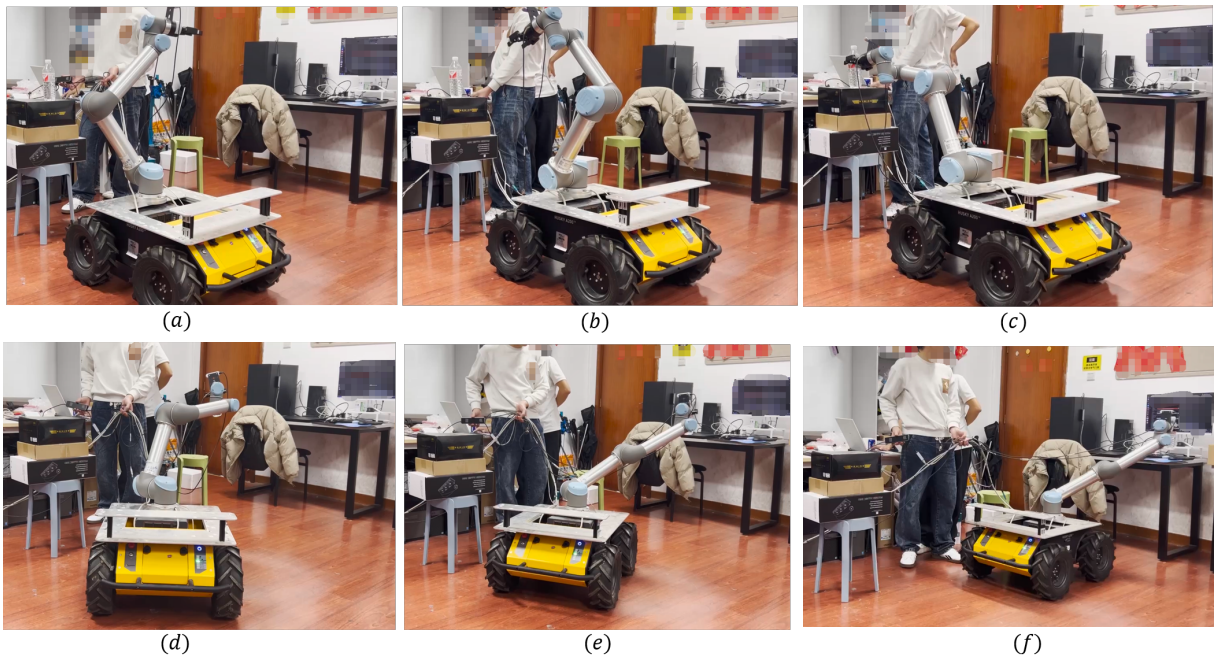

Figure 2: (a) Identify and locate the goal position. (b) Detect and locate the water bottle using the YOLO V7 algorithm. (c) Move the robotic arm to grab the water bottle. (d) Begin to deliver the water bottle in a vehicle-arm synergistic manner. (e) The robotic arm and mobile base move together for efficiency, not just moving the base. (f) The water bottle is delivered to the specified goal position. **Please see our attached video for a more intuitive MM process.**

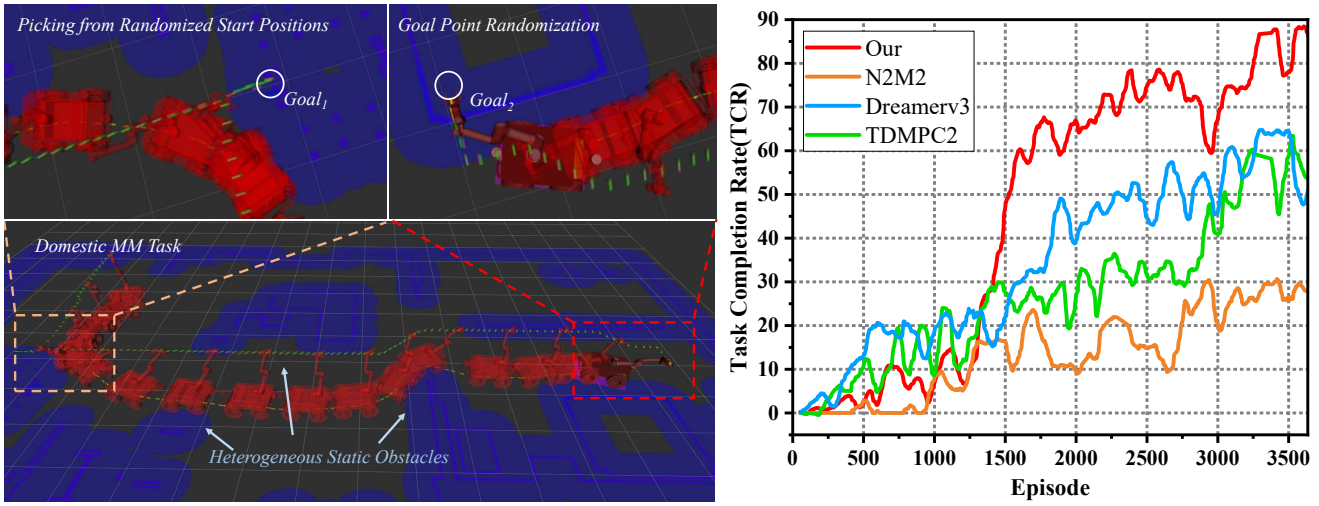

Figure 3: (**Left**) An illustration of the home-scene-oriented MM task. In each episode, the positions of the robot, the picking point (Goal_1), and the placing point (Goal_2) are randomly initialized. The layout of static obstacles in the environment is invariant. This task requires the robot to pick up an object at Goal_1 and then place the object at Goal_2. (**Right**) TCR metrics change with the number of training episodes in the scene-oriented experimental configuration.

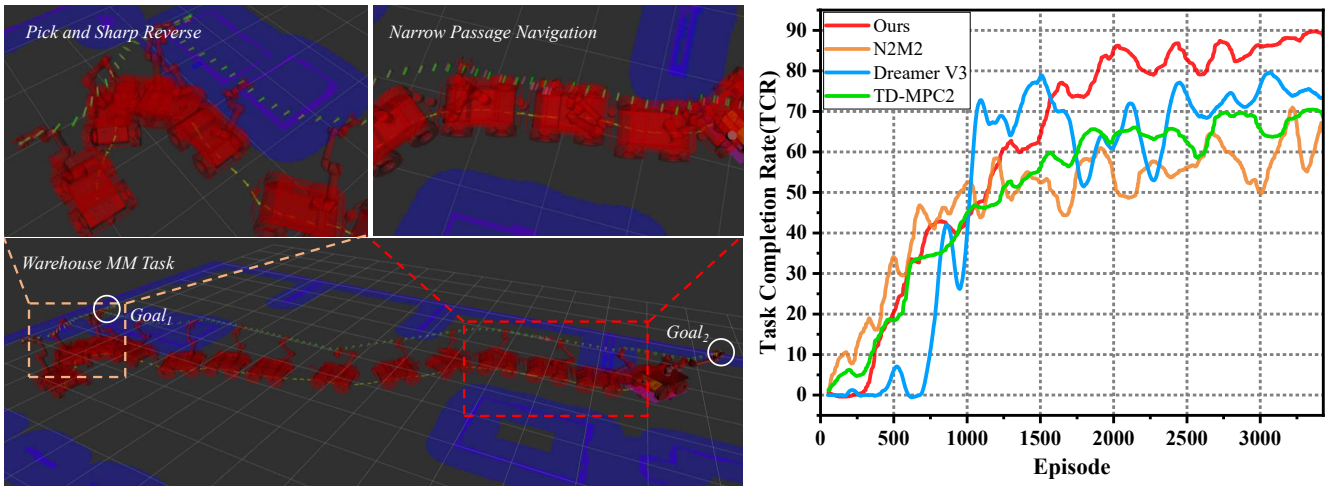

Figure 4: (**Left**) An illustration of the warehouse-oriented MM task. In each episode, the positions of the robot, the picking point (Goal_1), and the placing point (Goal_2) are randomly initialized. The layout of static obstacles in the environment is invariant. This task requires the robot to pick up an object at Goal_1 and then place the object at Goal_2. (**Right**) TCR metrics change with the number of training episodes in the warehouse-oriented experimental configuration.

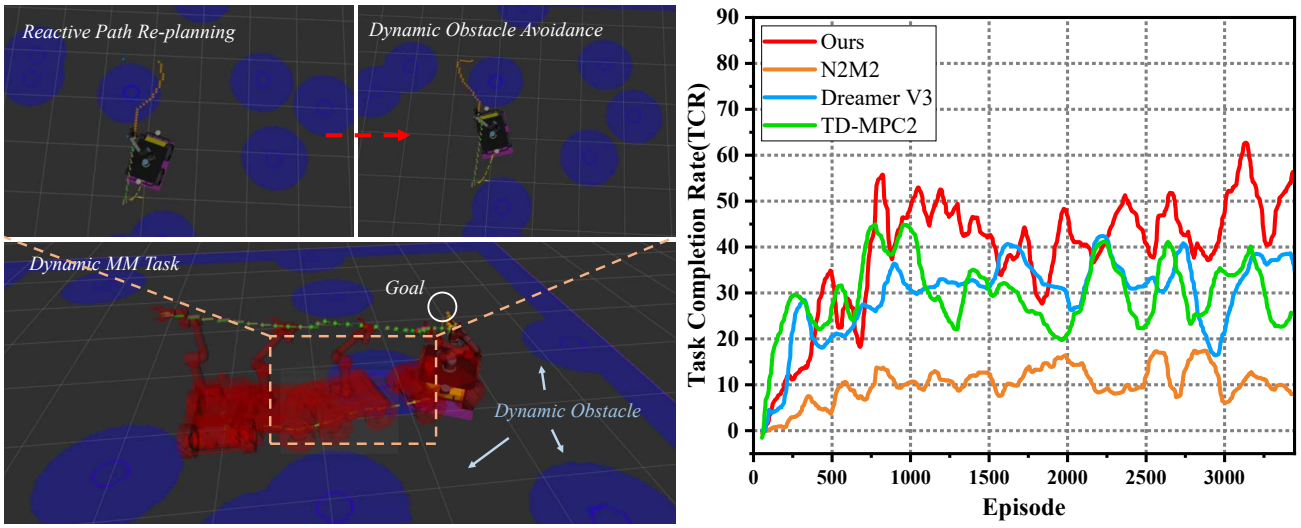

Figure 5: (**Left**) An illustration of the dynamic-scene-oriented MM task. In each episode, the positions of the robot, the picking point (Goal_1), and the placing point (Goal_2) are randomly initialized. A number of free-moving cubic objects are randomly initialized in the scene. This task requires the robot to pick up an object at Goal_1 and then place the object at Goal_2 while avoiding dynamic obstacles. (**Right**) TCR metrics change with the number of training episodes in the dynamic-scene-oriented experimental configuration.