Learning Compliant Manipulation in Space

Andrej Orsula[1] Matthieu Geist[2] Miguel Olivares-Mendez[1] Carol Martinez[1]

***Abstract*— As humanity ventures deeper into space, the demand for autonomous robotic systems capable of performing complex manipulation tasks is becoming increasingly critical. In this work, we investigate the integration of Operational Space Control with model-based Reinforcement Learning to achieve adaptive compliant manipulation in challenging space environments. Our experimental results across two demanding simulated space scenarios demonstrate that learning to modulate compliance significantly improves convergence and movement smoothness compared to standard differential Inverse Kinematics control. These findings underscore the potential of learned software-defined compliance for robust manipulation under the unpredictable conditions of space. A critical next step involves validating these strategies via sim-to-real transfer on physical robotic platforms to establish practical viability.**

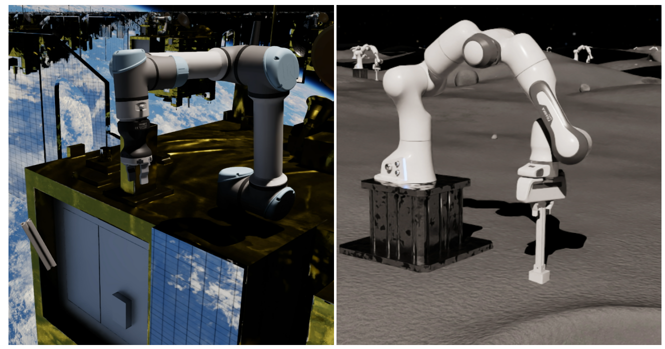

Fig. 1. Our model-based compliant manipulation is trained and evaluated in diverse procedural scenarios with hundreds of parallel environments.

I. Introduction

Robotic manipulation is a key enabler for the success of future space missions by supporting critical operations such as orbital structure assembly and planetary habitat construction. However, the extreme conditions of space coupled with the limited human oversight require robust and adaptive control strategies. This presents a fundamental challenge for developing systems that are resilient enough to withstand harsh conditions while maintaining the flexibility needed to adapt to unpredictable and dynamic environments. Although traditional approaches are effective in structured settings, they often fail in unstructured environments that are inherent to space exploration. To address these challenges, learning-based approaches have gained traction in recent years due to their data-driven nature that enables robots to learn from experience and adapt to new situations. Yet, their application for complex manipulation tasks in space remains limited.

At the same time, rigid robotic systems are often ill-suited for the unpredictable dynamics of space environments, where unmodeled forces and contact interactions can lead to loss of control or damage to both the robot and its environment. This is particularly relevant in scenarios such as on-orbit servicing, where robots must interact with free-floating tools or debris, and planetary exploration, where regolith or dust can result in unmodeled contacts. The need for compliant manipulation strategies is further exacerbated by the limited ability for human intervention in space, as astronauts may not be able to provide real-time feedback or assistance during complex tasks. This highlights the importance of developing robots that can autonomously adapt to their environment and handle unexpected situations.

Learning-based approaches such as Reinforcement Learning (RL) are general in nature, which makes them suitable for a variety of tasks in dynamic environments. While it is possible to learn compliant manipulation strategies through direct control of joint commands, this approach is often impractical in real-world applications due to the complexity of the involved dynamics. Leveraging the knowledge of kinematics through high-level task space commands can significantly simplify the learning process for complex tasks and result in a robot-agnostic policy. However, methods based on Inverse Kinematics (IK) are often brittle and limit the ability of the robot to control its compliance; this is particularly relevant in space environments where the resulting lack of compliance makes the system susceptible to instability or damage from unmodeled forces and contact interactions.

To address these challenges, Operational Space Control (OSC) [1] provides a principled framework for implementing controlled compliance that enables direct end-effector control with adjustable stiffness and damping. In addition to manually tuning the compliance gains, OSC can be extended with learning-based approaches that dynamically adjust the stiffness and/or damping based on the sensory feedback and requirements of the task [2], with the goal of enhancing the adaptability of autonomous systems. While prior research has combined OSC with model-free RL for terrestrial tasks [3], its application to space robotics remains unexplored. Similarly, the integration of OSC with model-based RL has not been studied. To this end, we propose an approach that integrates the principles of OSC with model-based RL to learn compliant manipulation in space environments. By leveraging the strengths of both OSC and RL, we aim to develop a control strategy that can effectively adapt to the challenges of space manipulation while maintaining the required precision and robustness.

[1]Space Robotics Research Group (SpaceR), Interdisciplinary Centre for Security, Reliability and Trust (SnT), University of Luxembourg
Contact: andrej.orsula@uni.lu
[2]Earth Species Project

II. Learning Compliant Manipulation

We explore the integration of OSC with model-based RL to learn compliant manipulation strategies suitable for space robotics. The core concept involves using OSC to provide software-defined compliance via adjustable stiffness and damping parameters. These parameters are dynamically modulated by an RL agent based on task requirements and sensory feedback. For the learning agent, we utilize DreamerV3 [4] that has demonstrated strong performance on complex tasks across diverse domains.

We investigate four distinct control strategies, primarily differing in how compliance is handled and learned:

- **IK** [$\dim(\mathcal{A}) = 6$]: Non-compliant baseline using standard differential IK.
- **OSC-CONST** [$\dim(\mathcal{A}) = 6$]: OSC with fixed compliance gains that provide constant passive compliance.
- **OSC-STIFF** [$\dim(\mathcal{A}) = 12$]: OSC where the agent learns to modulate stiffness (K_p) gains.
- **OSC-VAR** [$\dim(\mathcal{A}) = 18$]: OSC where the agent learns to modulate stiffness (K_p) and damping (K_d) gains.

In all strategies, a core component of the action space $\mathcal{A}$ includes the desired end-effector translational and rotational displacements ($\Delta\mathbf{x} \in \mathbb{R}^6$). For the adaptive strategies, the action space is augmented with **OSC-STIFF** adding 6 dimensions for the stiffness gains K_p, while **OSC-VAR** adds another 6 dimensions for the damping gains K_d. This allows the agent to actively control the robot's impedance based on the learned policy and the perceived state.

III. Results

We evaluate the performance of the four control strategies on two tasks from the Space Robotics Bench [5] shown in Fig. 1, namely the `debris_capture` task, which involves capturing and de-tumbling a free-floating object in orbit, and the `peg_in_hole_assembly` task, which requires grasping and inserting a peg into a hole on a static surface.

Learning Performance: Fig. 2 presents the learning curves for both tasks. The results consistently show that all OSC-based approaches outperform the non-compliant **IK** baseline. Notably, the adaptive strategies **OSC-STIFF** and **OSC-VAR** demonstrate a more stable convergence in the contact-intensive `peg_in_hole_assembly` task, which highlights the benefit of learned compliance.

Motion Smoothness: Beyond task success, we analyzed motion smoothness by evaluating the average joint-space jerkiness during execution. Fig. 3 compares the jerkiness of different strategies on the `peg_in_hole_assembly` task. The results indicate that **OSC-VAR** achieves significantly smoother motion by exhibiting the lowest jerkiness among all tested methods. This reduction in jerkiness suggests more controlled and predictable movements, which is highly desirable for minimizing vibrations and potential instability during delicate manipulation tasks in space environments.

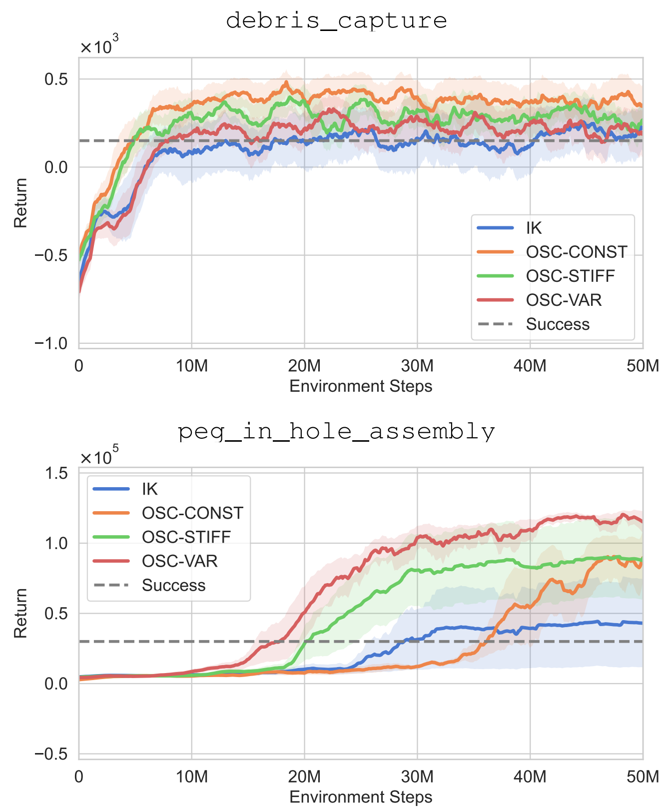

Fig. 2. Learning curves for the four strategies averaged over three seeds.

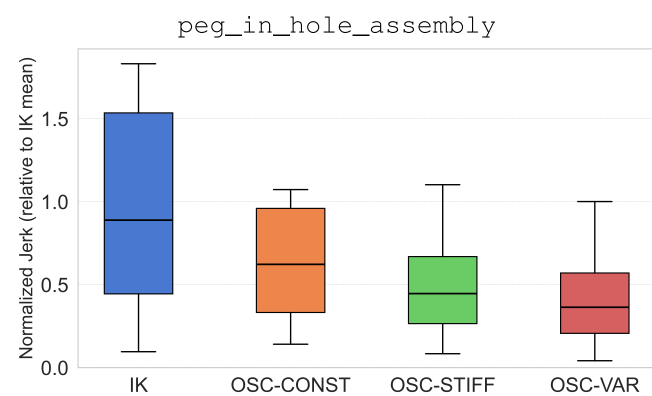

Fig. 3. Comparison of average joint-space jerkiness for the four strategies.

IV. Conclusion and Future Directions

This work demonstrates that integrating learned Operational Space Control with model-based Reinforcement Learning significantly enhances compliant manipulation for space robotics tasks. Our findings show that dynamically modulating compliance gains results in better convergence, training stability, and motion smoothness when compared to standard differential Inverse Kinematics or fixed compliance strategies. While these simulation results are promising, future work must focus on bridging the sim-to-real gap. Validating these learned compliance strategies on physical hardware is the critical next step toward establishing their practical viability for enabling more autonomous and robust robotic operations in future space missions.

REFERENCES

[1] O. Khatib, “A unified approach for motion and force control of robot manipulators: The operational space formulation,” *IEEE Journal on Robotics and Automation*, vol. 3, no. 1, pp. 43–53, 1987.

[2] J. Peters and S. Schaal, “Learning operational space control,” *Proceedings of Robotics: Science and Systems (RSS), Philadelphia, PA*, 08 2006.

[3] J. Wong, V. Makoviychuk, A. Anandkumar, and Y. Zhu, “OSCAR: Data-Driven Operational Space Control for Adaptive and Robust Robot Manipulation,” in *2022 International Conference on Robotics and Automation*, 2022, pp. 10 519–10 526.

[4] D. Hafner, J. Pasukonis, J. Ba, and T. Lillicrap, “Mastering diverse control tasks through world models,” *Nature*, vol. 640, pp. 647–665, 2025.

[5] A. Orsula, A. Richard, M. Geist, M. Olivares-Mendez, and C. Martinez, “Towards Benchmarking Robotic Manipulation in Space,” in *Conference on Robot Learning (CoRL) Workshop on Mastering Robot Manipulation in a World of Abundant Data*, 2024.