

Learning Swing-up Maneuvers for a Suspended Aerial Manipulation Platform in a Hierarchical Control Framework

Hemjyoti Das * Minh Nhat Vu ** Christian Ott **

* *Automation and Control Institute (ACIN), TU Wien, Gusshausstraße 27-29, 1040, Vienna, Austria (email: hemjyoti.das@tuwien.ac.at, minh.vu@tuwien.ac.at, christian.ott@tuwien.ac.at)*

** *Austrian Institute of Technology (AIT) GmbH, 1210, Vienna, Austria*

*** *Institute of Robotics and Mechatronics, German Aerospace Center (DLR), Oberpfaffenhofen, Muenchener Strasse 20, 82234, Wessling, Germany*

Abstract: In this work, we present a novel approach to augment a model-based control method with a reinforcement learning (RL) agent and demonstrate a swing-up maneuver with a suspended aerial manipulation platform. These platforms are targeted towards a wide range of applications on construction sites involving cranes, with swing-up maneuvers allowing it to perch at a given location, inaccessible with purely the thrust force of the platform. Our proposed approach is based on a hierarchical control framework, which allows different tasks to be executed according to their assigned priorities. An RL agent is then subsequently utilized to adjust the reference set-point of the lower-priority tasks to perform the swing-up maneuver, which is confined in the nullspace of the higher-priority tasks, such as maintaining a specific orientation and position of the end-effector. Our approach is validated using extensive numerical simulation studies.

Copyright © 2025 The Authors. This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

Keywords: Aerial manipulation, swing-up control, hierarchical control, reinforcement learning

1. INTRODUCTION

Aerial robotic manipulation has advanced significantly in recent years for complex physical-interaction tasks such as inspection and maintenance of pipelines, electricity lines, and remotely located construction sites (Ruggiero et al. (2018); Ollero et al. (2021); Ryll et al. (2017)). These systems, usually comprising of a multirotor platform equipped with a manipulator (Ollero et al. (2021)), typically demand a high amount of energy for the stabilization of the complete system, which further leads to a reduced flight time and payload capacity (Paredes et al. (2017); Villa et al. (2020)). Additionally, the large size of the propellers necessary to compensate for the mass of the heavy manipulator can cause turbulence in the surroundings (Kondak et al. (2014)).

To tackle these challenges, a suspended aerial manipulator (SAM) platform is proposed in Sarkisov et al. (2019), which consists of an omnidirectional platform suspended from a carrier system using rigid cables. The SAM platform is equipped with a 7 Degrees of Freedom (DoF) manipulator. It has been utilized to demonstrate compliant physical interaction with an unknown environment in a hierarchical control framework (Coelho et al. (2021); Sarkisov et al. (2023); Gabellieri et al. (2020)). In Yigit et al. (2021); Kong et al. (2022), similar suspended aerial manipulator platforms have been proposed to demonstrate

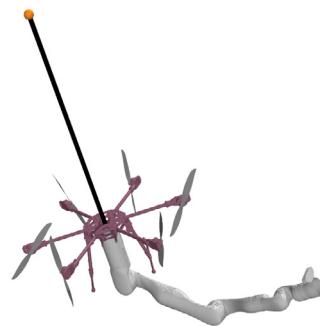


Fig. 1. Suspended aerial manipulation platform.

accurate end-effector trajectory tracking and teleoperation tasks with the environment.

This work extends the application of such a suspended aerial platform for performing swing-up maneuvers by utilizing the motion of only the manipulator. The problem of generating swing-up motion for different mechanical systems has been widely studied in the literature. For instance, during the 1960s, research into stabilizing single inverted pendulums gained significant attention, as demonstrated in works like Cannon (1967) and Ogata and Yang (1970). These studies primarily relied on specially designed mechanical systems that generate the necessary forces and torques for swing-up and stabilization tasks.

However, implementing such tasks on industrial robots presents additional challenges, including constraints on joint positions, velocities, and torques. Winkler and Suchý (2009) implemented a combined feedforward and balancing control strategy to achieve the swing-up and stabilization of a single inverted pendulum on a KUKA KR3 robot. More recently, advancements in the control of robotic arms with seven degrees of freedom (DoFs) have utilized Bayesian optimization (Marco et al. (2016)) and model-based policy search (Doerr et al. (2017)) to automate the tuning of control strategies for balancing tasks. The stabilization of more complex pendulums, such as spherical pendulums, is first successfully demonstrated on the “DLR LWR II” robot by Schreiber et al. (2001). Building on this foundation, recent studies have begun exploring the swing-up of spherical pendulums using industrial robots, with promising experimental results reported in Vu et al. (2021).

This work considers such swing-up maneuvers for a suspended aerial manipulation platform. However, our work aims to accomplish other tasks with the platform while simultaneously performing the swing-up maneuver, for which we utilize a hierarchical control framework (Dietrich and Ott (2019)). Thus, one of the contributions of this work is to formulate a task-hierarchy approach, where we ensure that the end-effector always maintains a specific orientation and position along a particular axis with a given yaw orientation of the platform. This is motivated by the choice of the mounted camera system on both the platform and end-effector of the manipulator, which should always observe the given target to grasp once the swing-up motion is completed. Concurrently, the swing-up control is performed as a second-priority task in a dynamically decoupled manner. A reinforcement learning agent is then proposed to adjust the reference coordinates for the lower-priority tasks of the designed hierarchical whole-body control framework and thus successfully learn the swing-up maneuver in the nullspace of the higher-priority tasks, which is the main contribution of this paper. The effectiveness of our proposed approaches is validated using extensive simulation studies.

The rest of the paper is organized as follows. In Section 2, we present the system dynamics of the suspended aerial platform with the attached manipulator. Section 3 introduces the task definition and hierarchical control formulation. The reinforcement learning policy and task references necessary for the swing-up maneuver are presented in Section 4. Finally, the simulation results are shown in Section 5, while Section 6 concludes the paper.

2. SYSTEM DYNAMICS AND OVERVIEW

In this work, we performed the swing-up maneuvers using the planar-thrust suspended aerial platform previously developed by our group in Das et al. (2024). The oscillations of the suspended platform are modeled as a single spherical pendulum, similar to the work by Sarkisov et al. (2020), with the cable used for the suspension assumed to be rigid and taut. As seen from Fig. 2, we denote the generalized coordinates of the suspended platform as q_1, q_2 and q_3 which are the rotations of the suspension cable about its z , y and x axis, respectively. The platform can apply a planar

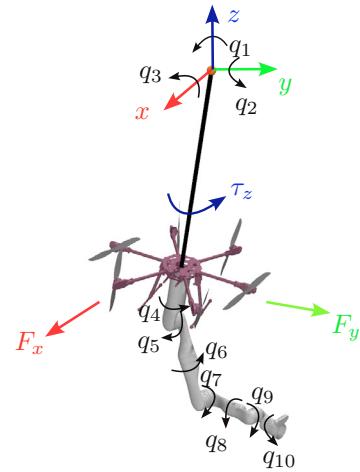


Fig. 2. Schematic representation of the cable-suspended aerial manipulation platform.

wrench \mathbf{u}_p about its axis (Das et al. (2024)), which consists of the translational forces along its x and y axis, denoted as F_x and F_y , respectively, and the torque τ_z exerted about its z axis. The planar wrench is related to the exerted joint torque τ_p at the platform joints using Jacobian $\mathbf{J}_u \in \mathbb{R}^{3 \times 3}$ as

$$\tau_p = \mathbf{J}_u^T \mathbf{u}_p \quad (1)$$

In this work, we considered the 7-DoF Kinnova Gen3 robotic manipulator attached to the suspended platform. Our framework is highly versatile and not restricted to a specific manipulator, allowing compatibility with a wide range of manipulators attached to the suspended platform. The generalized coordinates of the manipulator is denoted as $\mathbf{q}_m \in \mathbb{R}^7$, which consists of the joints $q_4, q_5, q_6, q_7, q_8, q_9$ and q_{10} as shown in Fig. 2. The joint torque exerted by the manipulator is denoted by the vector $\tau_m \in \mathbb{R}^7$. The complete suspended aerial manipulation system is thus modeled as a 10 DoF system, with its system dynamics summarized as

$$\mathbf{M}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} + \mathbf{g}(\mathbf{q}) = \boldsymbol{\tau}, \quad (2)$$

where $\mathbf{M}(\mathbf{q}) \in \mathbb{R}^{10 \times 10}$ is the inertia matrix, $\mathbf{C}(\mathbf{q}, \dot{\mathbf{q}}) \in \mathbb{R}^{10 \times 10}$, $\mathbf{g}(\mathbf{q}) \in \mathbb{R}^{10}$ is the gravity torque exerted on each joint. $\boldsymbol{\tau}$ is the generalized joint torque vector summarized as

$$\boldsymbol{\tau} = [\tau_p \ \tau_m]^T \quad (3)$$

3. CONTROLLER DESIGN

In this section, we design a whole-body controller in a hierarchical framework. Such control approaches (Dietrich and Ott (2019)) allow us to execute multiple tasks while assigning a certain priority to each task, such that the lower-priority tasks do not interfere with the execution of the higher-priority tasks. This strict hierarchy between different tasks is achieved by defining dynamically consistent null-space projectors, which leads to decoupled task-space velocities. In this work, we define the highest priority task $\mathbf{x}_1 \in \mathbb{R}^5$ as the stabilization of the end-effector orientation, end-effector position along the axis perpendicular to the swing-up axis, and the platform yaw angle. The second priority task $\mathbf{x}_2 \in \mathbb{R}^5$ concerns the motion of the specific

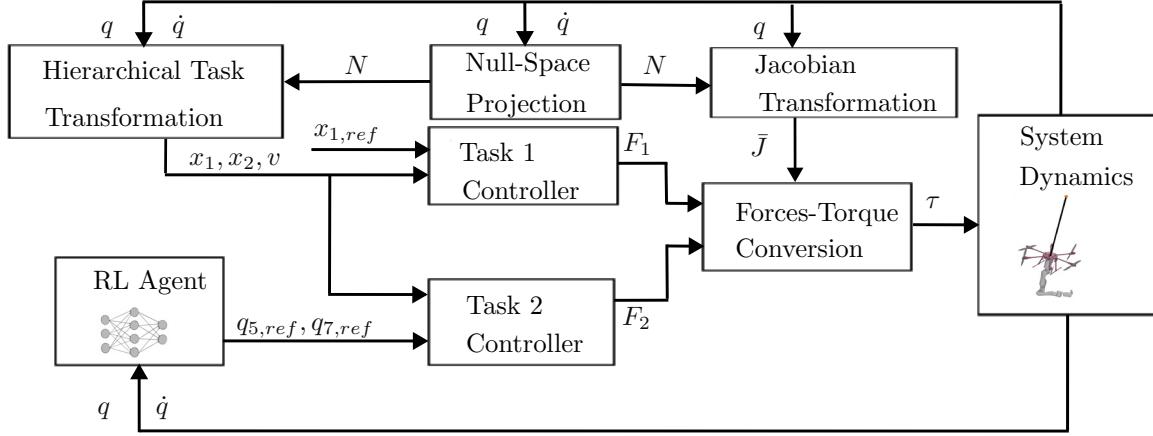


Fig. 3. Overview of our proposed control framework.

joints, which enable the swing-up of the complete system. \mathbf{x}_2 also contains the platform state q_2 along the swing-up axis, as will be explained in the subsequent section. The task-spaces \mathbf{x}_1 and \mathbf{x}_2 are thus summarized as follows,

$$\begin{aligned}\mathbf{x}_1 &= [y_e \ \phi \ \theta \ \psi \ q_1]^T \\ \mathbf{x}_2 &= [q_2 \ q_4 \ q_5 \ q_6 \ q_7]^T,\end{aligned}\quad (4)$$

where y_e is the end-effector's position along its y axis. ϕ , θ and ψ are the Euler angles of the end-effector defined about its x , y and z axis, respectively. Note that the goal of this work is to perform the swing-up maneuver along the x - z plane, as shown in Fig. 2 while keeping the end-effector along a given orientation and a given y -axis position, in order to observe the given target to be grasped after swinging up. Additionally, the first-priority task includes the yaw orientation q_1 for similar reasons and thus should always point in a specific direction to observe the target. The second priority task \mathbf{x}_2 concerns the motion of certain joints, which will aid in the swing-up maneuver. Note that in this work, we aim to swing up the complete system by primarily utilizing the motion of the manipulator while using minimal propulsion force from the platforms, to ensure energy efficiency. Based on these definitions, we summarize the task-space velocities as

$$\begin{aligned}\dot{\mathbf{x}}_1 &= \mathbf{J}_1 \dot{\mathbf{q}} \\ \dot{\mathbf{x}}_2 &= \mathbf{J}_2 \dot{\mathbf{q}},\end{aligned}\quad (5)$$

where the task-space Jacobians $\mathbf{J}_1 \in \mathbb{R}^{5 \times 10}$ and $\mathbf{J}_2 \in \mathbb{R}^{5 \times 10}$ are obtained kinematically and have full row-rank. Next, to decouple the task-space velocities and to define a strict hierarchy amongst the different tasks, we define the null-space projection matrix (Dietrich and Ott (2019)) for the second task as

$$\mathbf{N} = \mathbf{I} - \mathbf{J}_1^T \mathbf{J}_1^{M+,T}, \quad (6)$$

where $\mathbf{I} \in \mathbb{R}^{10 \times 10}$ is the identity matrix, while the operator ${}^{M+}$ denotes a dynamically consistent pseudoinversion (Dietrich and Ott (2019)). Next, we define the decoupled Jacobian matrices $\bar{\mathbf{J}}$ as

$$\bar{\mathbf{J}} = \begin{bmatrix} \bar{\mathbf{J}}_1 \\ \bar{\mathbf{J}}_2 \end{bmatrix} = \begin{bmatrix} \mathbf{J}_1 \\ \mathbf{J}_2 \mathbf{N}^T \end{bmatrix}. \quad (7)$$

The decoupled task-space velocities are then defined as

$$\mathbf{v} = \begin{bmatrix} \mathbf{v}_1 \\ \mathbf{v}_2 \end{bmatrix} = \begin{bmatrix} \bar{\mathbf{J}}_1 \\ \bar{\mathbf{J}}_2 \end{bmatrix} \dot{\mathbf{q}} = \bar{\mathbf{J}} \dot{\mathbf{q}}. \quad (8)$$

By utilizing these task-space coordinates, we conduct the coordinate transformation of (2) to obtain the hierarchical decoupled equations of motion as

$$\boldsymbol{\Lambda}(\mathbf{q}) \dot{\mathbf{v}} + \boldsymbol{\mu}(\mathbf{q}, \dot{\mathbf{q}}) \mathbf{v} = \bar{\mathbf{J}}^{-T} (\boldsymbol{\tau} - \mathbf{g}), \quad (9)$$

where $\boldsymbol{\Lambda} = (\bar{\mathbf{J}} \mathbf{M}^{-1} \bar{\mathbf{J}}^T)^{-1}$ is the decoupled task-space inertia and $\boldsymbol{\mu} = \boldsymbol{\Lambda} (\bar{\mathbf{J}} \mathbf{M}^{-1} \mathbf{C} \bar{\mathbf{J}}^{-1} - \dot{\mathbf{J}} \bar{\mathbf{J}}^{-1})$ is the transformed Coriolis term in the task space. Next, we define the joint torque $\boldsymbol{\tau}$ as

$$\boldsymbol{\tau} = \mathbf{g} + \boldsymbol{\tau}_u + \bar{\mathbf{J}}^T \mathbf{F}, \quad (10)$$

where the torque $\boldsymbol{\tau}_u$ compensates for the cross-coupling Coriolis components (Dietrich and Ott (2019)), whereas the task-space force $\mathbf{F} = [\mathbf{F}_1 \ \mathbf{F}_2]^T$ consists of the forces necessary to accomplish task 1 and task 2, defined by \mathbf{F}_1 and \mathbf{F}_2 , respectively, and both are chosen as PD control laws, similar to Das et al. (2023). The control law (10) thus ensures the execution of the different tasks hierarchically. The control law also ensures that the swing-up maneuver will be executed primarily due to the motion of the manipulator by choosing an appropriate task force \mathbf{F}_2 , as will be discussed subsequently.

4. LEARNING POLICY AND TASK REFERENCE

In this section, the hierarchical control framework is integrated with a reinforcement learning (RL) agent to achieve desired reference trajectories for the suspended aerial manipulation system. The Soft Actor-Critic (SAC) algorithm (Haarnoja et al. (2018)) is selected for its ability to handle continuous action spaces, which are inherent to the control of robotic manipulators. Furthermore, SAC balances exploration and exploitation, ensuring effective learning and robust policy performance.

Actor-critic methods combine value-based and policy-based reinforcement learning principles, addressing some of the limitations of standalone approaches. As an Actor-Critic method, SAC utilizes a policy gradient to optimize a parameterized policy. Unlike REINFORCE, see Arulkumaran et al. (2017), which updates policies based on episodic returns, SAC employs a Critic to estimate

the action-value function, enhancing sample efficiency. The policy gradient update in SAC incorporates the advantage function; see Haarnoja et al. (2018):

$$\nabla_{\theta} J(\pi_{\theta}) = \mathbb{E} \left[\nabla_{\theta} \log \pi_{\theta}(\mathbf{u}_k | \mathbf{x}_k) A^{\pi}(\mathbf{x}_k, \mathbf{u}_k) \right], \quad (11)$$

where $A^{\pi}(\mathbf{x}_k, \mathbf{u}_k) = Q^{\pi}(\mathbf{x}_k, \mathbf{u}_k) - V^{\pi}(\mathbf{x}_k)$ represents the advantage function, with \mathbf{x}_k and \mathbf{u}_k denoting the observation and the action space vector at the i^{th} iteration, respectively. This formulation ensures stable training dynamics and improved performance in continuous action spaces.

To train the RL agent, we consider the observation space as the joints \mathbf{q} and its velocities $\dot{\mathbf{q}}$. The action space consists of the two elbow joints q_5 and q_7 . We choose the reward function as the exponential error of the platform joint q_2 along the swing-up axis. As mentioned previously, the primary task \mathbf{x}_1 does not contribute to the swing-up that is accomplished using the secondary task \mathbf{x}_2 . Additionally, due to the involvement of certain joints q_4 and q_6 perpendicular to the swing-up axis, a zero reference is commanded to them. In this work, we demonstrate the swing-up motion by using minimal thrust force from the platform. Therefore, the task force \mathbf{F}_2 corresponding to the joint q_2 is used to counteract its corresponding gravity compensation g_2 in (10). For similar reasons, the component $\tau_{u,2}$ of the cross-coupling compensation torque τ_u corresponding to q_2 is also considered as zero. The elbow joint q_9 is not considered in the second task due to its comparatively low inertia and possible conflict with the primary task.

The proposed reinforcement learning framework thus utilizes an observation space $[\mathbf{q} \ \dot{\mathbf{q}}]^T$ and an action space $[q_{5,\text{ref}} \ q_{7,\text{ref}}]^T$, where $q_{5,\text{ref}}$ and $q_{7,\text{ref}}$ represent the reference commands for the corresponding joints. The agent is rewarded according to the exponential function $e^{-5|q_2-1|}$, thus favoring actions that bring q_2 close to 1 rad. To ensure proper task execution, the references $\mathbf{x}_{1,\text{ref}} = [0 \ 0 \ 0 \ -\frac{\pi}{2} \ 0]^T$ and $\mathbf{x}_{2,\text{ref}}[2:5] = [0 \ q_{5,\text{ref}} \ 0 \ q_{7,\text{ref}}]^T$, where [2:5] indicates elements 2 through 5 of $x_{2,\text{ref}}$ are defined for the respective tasks, with the thrust force of the platform commanded to cancel out the gravity compensation g_2 and the cross-coupling Coriolis compensation $\tau_{u,2}$ through the task force $\mathbf{F}_2[1]$ corresponding to q_2 . This setup ensures stable platform behavior and effective regulation of the target joints through the learned policy.

5. RESULTS

This section presents the numerical simulation results and analyzes the swing-up maneuver with our proposed approach. We considered the mass of the suspended aerial platform to be 4 kg, with its principal moment of inertia along its x , y , and z axis being 0.0646, 0.0646 and 0.0682 kg m², respectively, similar to our previous work in Das et al. (2024). The platform is attached using a rigid cable of length 1 m, with its mass being 0.3 kg. The mass of the 7-DoF manipulator is considered as 8.2 kg, with the maximum reach of the end-effector being 0.9 m. We performed our analysis in the Mujoco Simulation environment (Todorov et al. (2012)), with a sampling time of 1000 Hz. We utilized the Stable-baseline3 RL library

(Raffin et al. (2021)), which is based on PyTorch, to train the RL agent. Both the actor and critic networks contain four hidden layers, each containing 256 neurons. We trained the networks for a total of around 1300 episodes, with each episode lasting 15 s. The training is accomplished in a custom-made environment based on OpenAI Gym. The total training time is roughly 100 min, with 30 environments being used in parallel to speed up the training time.

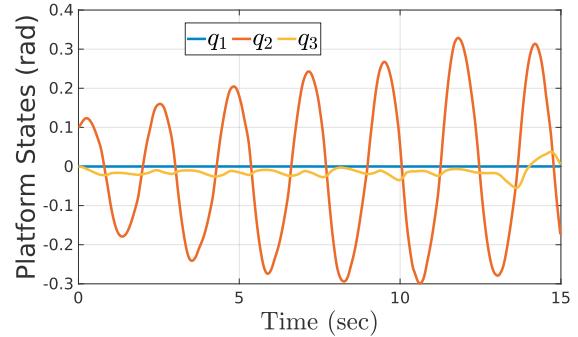


Fig. 4. Swing-up behaviour of the platform.

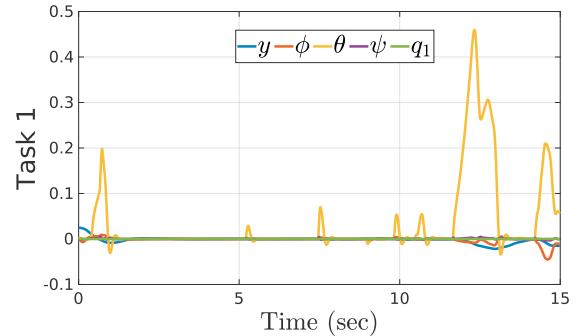


Fig. 5. Behaviour of task \mathbf{x}_1 , with y expressed in m, while ϕ, θ, ψ and q_1 are expressed in rad.

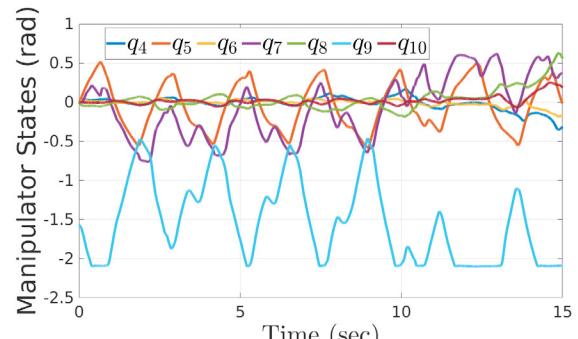


Fig. 6. Behaviour of the manipulator.

In this work, we command the task \mathbf{x}_2 to execute the swing-up maneuver in the nullspace of the primary task \mathbf{x}_1 . Additionally, to aid in the swing-up maneuver, we considered the initial angle $q_{2,0}$ for the joint q_2 during both the training and evaluation phases as 0.1 rad, which can be achieved by the propulsion force of the platform. The sequence of the swing-up experiment is shown in Fig. 7. From Fig. 4, we observe that the platform can reach the

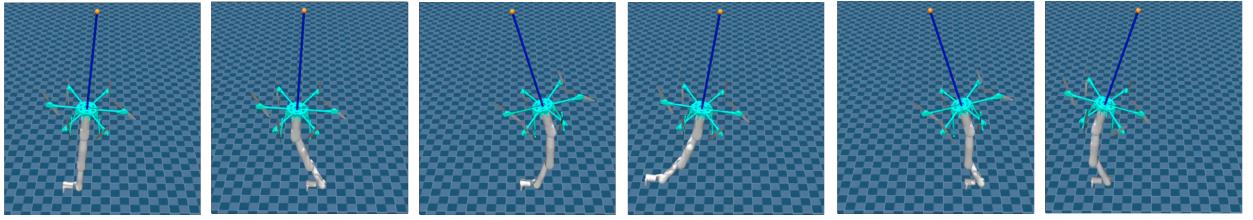


Fig. 7. Sequence of the swing-up experiments at different times $t = 0$ s, 5.3 s, 8.1 s, 9.2 s, 12.9 s, and 14.2 s, displayed from left to right.

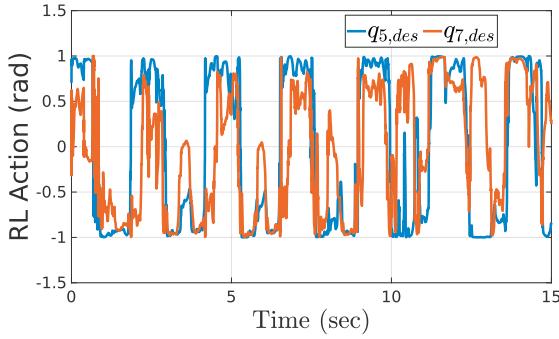


Fig. 8. Commanded angles by the RL agent for the joints q_5 and q_7 .

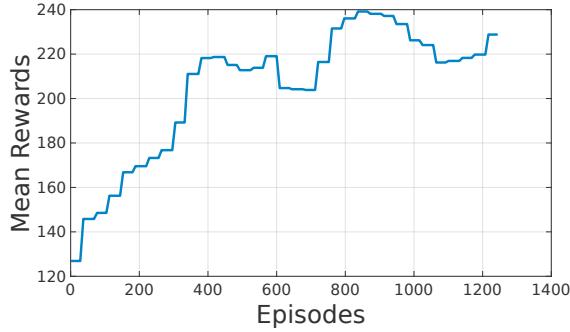
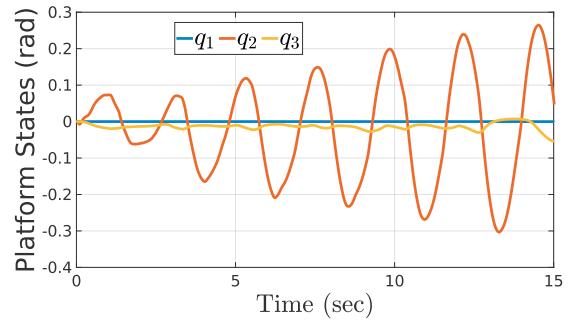
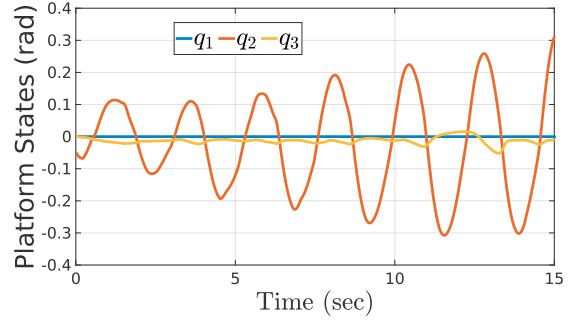


Fig. 9. Mean episodic rewards during training of the RL agent.

swing-up angle of up to 0.32 rad in a time-period of 15 s. The joint angles commanded by the RL agent is limited between -1 rad and 1 rad to avoid any aggressive swing-up maneuver (Fig. 8). The RL agent commands action to the secondary task in the nullspace of the primary task, and therefore do not influence it due to the involved dynamical decoupling. The mean episodic rewards during the training of the RL agent are shown in Fig. 9. The behavior of the task \mathbf{x}_1 is plotted in Fig. 5. We observe that the tracking error is not zero throughout the swing-up maneuver, which can possibly be reduced by increasing the actuator bandwidth and its saturation limits. The behavior of the manipulator is plotted in figure 6, where we observe the motion of the joints q_5 and q_7 utilized for the swing-up maneuver. We also observe that the joint q_9 is primarily used for maintaining the end-effector orientation along the y - z plane. Furthermore, in order to demonstrate the robustness of our trained network, we also demonstrate the swing-up maneuver of the platform from different initial configurations, as shown in Fig. 10.



(a) Swing-up behaviour of the platform with $q_{2,0} = 0$ rad.



(b) Swing-up behaviour of the platform with $q_{2,0} = -0.05$ rad.

Fig. 10. Swing-up behaviour of the platform with different initial conditions.

6. CONCLUSION

In this work, we integrated a reinforcement learning agent to adjust the secondary task references in the null space of the higher-priority task. Our framework is based on a hierarchical control framework, which ensures that the primary tasks related to both the orientation and the y axis position of the end-effector, and also the platform's yaw angle, are not interfered with by the execution of the secondary tasks related to the swing-up maneuver. We utilized an off-policy RL agent based on the state-of-the-art SAC algorithm to obtain the reference setpoints of the two elbow joints of the manipulator and we successfully demonstrated the swing-up maneuver of the suspended aerial manipulation system.

We will implement the RL agent on the physical system as future work. This would further necessitate domain randomization to ensure the robustness of the learned policy to unknown system dynamics and sensor noise. We would also like to analyze the use of the thrusters of the platform to further aid the swing-up maneuver by integrating it into the RL agent.

REFERENCES

Arulkumaran, K., Deisenroth, M.P., Brundage, M., and Bharath, A.A. (2017). Deep reinforcement learning: A brief survey. *IEEE Signal Processing Magazine*, 34(6), 26–38.

Cannon, R.H. (1967). *Dynamics of physical systems*. McGraw-Hill.

Coelho, A., Sarkisov, Y.S., Lee, J., Balachandran, R., Franchi, A., Kondak, K., and Ott, C. (2021). Hierarchical control of redundant aerial manipulators with enhanced field of view. In *Proc. IEEE Int. Conf. on Unm. Air. Syst.*, 994–1002.

Das, H., Sæbø, B.K., Pettersen, K.Y., and Ott, C. (2023). Hardware-in-the-loop simulation of vehicle-manipulator systems for physical interaction tasks. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 4390–4396. IEEE.

Das, H., Vu, M.N., Egle, T., and Ott, C. (2024). Observer-based controller design for oscillation damping of a novel suspended underactuated aerial platform. In *2024 IEEE Int. Conf. Robot. Automat. (ICRA)*, 8415–8421. doi: 10.1109/ICRA57147.2024.10610305.

Dietrich, A. and Ott, C. (2019). Hierarchical impedance-based tracking control of kinematically redundant robots. *IEEE Transactions on Robotics*, 36(1), 204–221.

Doerr, A., Nguyen-Tuong, D., Marco, A., Schaal, S., and Trimpe, S. (2017). Model-based policy search for automatic tuning of multivariate PID controllers. *Proceedings of the IEEE International Conference on Robotics and Automation*, 5295–5301.

Gabellieri, C., Sarkisov, Y.S., Coelho, A., Pallottino, L., Kondak, K., and Kim, M.J. (2020). Compliance control of a cable-suspended aerial manipulator using hierarchical control framework. In *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 7196–7202.

Haarnoja, T., Zhou, A., Abbeel, P., and Levine, S. (2018). Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, 1861–1870. PMLR.

Kondak, K., Huber, F., Schwarzbach, M., Laiacker, M., Sommer, D., Bejar, M., and Ollero, A. (2014). Aerial manipulation robot composed of an autonomous helicopter and a 7 degrees of freedom industrial manipulator. In *Proc. IEEE Int. Conf. Robot. Automat.*, 2107–2112.

Kong, F., Monteleone, S., Grioli, G., Catalano, M.G., and Bicchi, A. (2022). A robotic aerial platform with functionally anthropomorphic arms designed for physical interaction. In *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 7581–7588.

Marco, A., Hennig, P., Bohg, J., Schaal, S., and Trimpe, S. (2016). Automatic LQR tuning based on gaussian process global optimization. *Proceedings of the IEEE International Conference on Robotics and Automation*, 270–277.

Ogata, K. and Yang, Y. (1970). *Modern control engineering*. Prentice Hall.

Ollero, A., Tognon, M., Suarez, A., Lee, D., and Franchi, A. (2021). Past, present, and future of aerial robotic manipulators. *IEEE Trans. Robot.*, 38(1), 626–645.

Paredes, J.A., Saito, C., Abarca, M., and Cuellar, F. (2017). Study of effects of high-altitude environments on multicopter and fixed-wing uavs' energy consumption and flight time. In *2017 13th IEEE Conference on Automation Science and Engineering (CASE)*, 1645–1650. IEEE.

Raffin, A., Hill, A., Gleave, A., Kanervisto, A., Ernestus, M., and Dormann, N. (2021). Stable-baselines3: Reliable reinforcement learning implementations. *Journal of Machine Learning Research*, 22(268), 1–8.

Ruggiero, F., Lippiello, V., and Ollero, A. (2018). Aerial manipulation: A literature review. *IEEE Robot. Automat. Lett.*, 3(3), 1957–1964.

Ryll, M., Muscio, G., Pierri, F., Cataldi, E., Antonelli, G., Caccavale, F., and Franchi, A. (2017). 6d physical interaction with a fully actuated aerial robot. In *2017 Proc. IEEE Int. Conf. Robot. Automat. (ICRA)*, 5190–5195. IEEE.

Sarkisov, Y.S., Coelho, A., Santos, M.G., Kim, M.J., Tsetserukou, D., Ott, C., and Kondak, K. (2023). Hierarchical whole-body control of the cable-suspended aerial manipulator endowed with winch-based actuation. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 5366–5372. IEEE.

Sarkisov, Y.S., Kim, M.J., Bicego, D., Tsetserukou, D., Ott, C., Franchi, A., and Kondak, K. (2019). Development of sam: cable-suspended aerial manipulator. In *2019 International Conference on Robotics and Automation (ICRA)*, 5323–5329. IEEE.

Sarkisov, Y.S., Kim, M.J., Coelho, A., Tsetserukou, D., Ott, C., and Kondak, K. (2020). Optimal oscillation damping control of cable-suspended aerial manipulator with a single imu sensor. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 5349–5355. IEEE.

Schreiber, G., Ott, C., and Hirzinger, G. (2001). Interactive redundant robotics: Control of the inverted pendulum with nullspace motion. *Proceedings of the IEEE International Conference on Intelligent Robots and Systems*, 1, 158–164.

Todorov, E., Erez, T., and Tassa, Y. (2012). Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ international conference on intelligent robots and systems*, 5026–5033. IEEE.

Villa, D.K., Brando, A.S., and Sarcinelli-Filho, M. (2020). A survey on load transportation using multirotor uavs. *Journal of Intelligent & Robotic Systems*, 98, 267–296.

Vu, M.N., Hartl-Nesic, C., and Kugi, A. (2021). Fast swing-up trajectory optimization for a spherical pendulum on a 7-dof collaborative robot. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 10114–10120. IEEE.

Winkler, A. and Suchý, J. (2009). Erecting and balancing of the inverted pendulum by an industrial robot. *IFAC Proceedings Volumes*, 42(16), 323–328.

Yigit, A., Perozo, M.A., Cuvillon, L., Durand, S., and Gangloff, J. (2021). Novel omnidirectional aerial manipulator with elastic suspension: Dynamic control and experimental performance assessment. *IEEE Robot. Automat. Lett.*, 6(2), 612–619.