Digital Twin–Integrated Deep Reinforcement Learning Framework for Real-Time Payload-Stable UAV Landing and Energy-Aware Descent Control in Industry 5.0

[1]R Raj Jawahar, [2]Prof Dr Midhunchakkaravarthy, [3]Dr Rajesh Dey
[1]Research Associate, Lincoln College University, Malaysia,
[1]Research Scholar, Anna University, Chennai, India;
[2]Dean, Lincoln College University, Malaysia;
[3]Professor, Gopal Narayan Singh University, India;
Email: pdf.raj@lincoln.edu.my

Abstract: UAVs which carry payloads are highly unstable during descent as a result of changing center of gravity, asymmetry in aerodynamic properties and external force being highly unpredictable. Traditional controllers like PID and MPC are not very effective in such nonlinear disturbances particularly when the payload is varied throughout the flight. In this study, a Digital Twin-based deep reinforcement learning architecture of real-time payload-stable landing and energy-optimal descent control is suggested. High-fidelity Digital Twin is a model that considers the impact of payload displacement, thrust asymmetry, drag variations, and battery properties up to 1 kHz, allowing a SAC agent to test possible corrective actions before implementing them to the real UAV. Tests using 0.5-1.2kg payloads demonstrate a 38 percent decrease in landing attitude error, 27 percent decrease in descent energy, and 42 percent decrease in lateral drift over an optimized MPC baseline. The twin foresees instability to occur 0.18 s before physical sensors thus permitting proactive stabilization. An explanatory layer comes up with explanations that can be read by humans, which correlates the system with the explainability and human-focused values of Industry 5.0.

Keywords: Digital Twin, UAV Landing, Payload Stability, SAC Reinforcement Learning, Energy Optimization, Industry 5.0.

Introduction

The nonlinear process of landing a UAV with payload on it is a complicated process due to the imbalance in thrusts, aerodynamic disturbances, and changes in center-of-gravity caused by the addition of the payload. Minimum offsets between payloads generate roll-pitch torques, which are transferred to subsequent lateral drift and oscillatory descent strategies. Classical PID controllers do not predict nonlinear disturbances but the MPC ones demand the availability of accurate models and fail when the payload distribution changes. DRL agents are adaptive, but can not be deployed immediately because of safety issues during exploration. Digital Twins address this by offering a high-frequency predictive reflection of the UAV dynamic states so that virtual assessment of actions can be done in real-time prior to physical implementation. The combination of DT prediction and SAC control gives a proactive stabilization process in line with Industry 5.0 where transparency, safety, resilience, and human alignment are the key. This study introduces a single DTDRL landing controller that is able to sustain the stability and reduce the energy usage of payload-carrying UAVs.

Methodology

The scheme of work of the proposed methodology is arranged as a closely-knit cyber-physical cycle whereby sensing, prediction, simulation and control functions are performed in inseparable sequence throughout UAV landing. The system starts by the UAV broadcasting high-frequency IMU, barometer, motor-thrust feedback and optical-flow data to the processor on board. Such uncodified measurements are usually subject to drift, noise due to vibrations, and short-lived bias, thus sensor fusion step with the use of a more advanced Kalman filter provides a smooth estimate of attitude, descent velocity, direction of drift, and angular acceleration. These smoothed estimations constitute the real-time input into the Digital Twin which is an operation as a high-fidelity dynamic simulation of the UAV at 1 kHz. The twin solves a payload-augmented form of the Newton-Euler equations, with the torque caused by the payload being explicitly represented as an asymmetric moment due to the distance between the center of gravity and the payload. This modeling is critical since a slight shift in the payload will result in sustained roll-pitch imbalance during landing, and the twins have to record the magnitude and the dynamics of this imbalance at milliseconds resolution.

The virtual environment within the twin, a vertical descent model is modeled with the help of physically grounded damping formulation which considers the saturation of thrust, aerodynamic drag, and ground-effect behavior as the UAV closes to the landing surface. In the mean time the motor thrust lag and battery internal resistance are modeled in terms of a curve of thrust-power, through which the twin can forecast not only the stability, but the energy needed to perform a certain action. The twin compares the Lyapunov index periodically on the basis of the estimated rollpitch development, when the forecasted derivative of the Lyapunov index has a positive value in excess of a few cycles, the twin issues a warning of an impending instability. This prediction is made before real sensors can provide it, and the reinforcement-learning controller is able to spend more time adjusting the descent trajectory.

The SAC controller is provided with a small yet informative state vector that includes predicted attitude deviation, generated torque imbalance due to disturbance in the payload, lateral drift velocity, and current descent rate, and the instantaneous energy slope calculated by the battery model. The SAC agent is able to generate proactive control in contrast to the traditional controllers that can only respond to observed errors, based on the future predictions provided by the Digital Twin. Continuous changes in the motor thrust distribution and fine adjustments in the distribution of the pitch and roll ensure a stable descent profile are issued by the agent. Nevertheless, the execution of the same commands is virtually tested in the Digital Twin before the actual UAV carries them out over a short predictive horizon. In this virtual rollout, the twin tests the effects of the ordered action on the amplification of the drift, or the rise of the torque asymmetry, or the appearance of the rapid oscillations, or the emergence of the undesirable energy spikes. Whether due to a prediction of instability by the simulation, or an alternative cause, the control logic suppresses or modifies the action. The mechanism turns the Digital Twin into a safety supervisor, which enables the reinforcement learning to be applied in the real flight without the threat of catastrophic behavior.

The reinforcement-learning policy is implemented along the lines of the reward scheme that punishes any deviation of the attitude that is large, rapid drift accumulation, and redundant thrust spending. Vertical sliding and reduced energy curves are rewarded and the controller can reduce the oscillations as well as power consumption at the same time. The landing process will have the twin constantly updating its internal status to reflect the changes in payload dynamics, such as small changes in mass distribution as it moves or tilts. The described updates affect the predicted torque offset, such that the SAC agent is never fed static approximations, but is always fed with the actual, real-time, payload-sensitive dynamics.

One of the cognitive explanation modules operates parallel to the control process, and it gathers internal messages of the twin and SAC agent. This module translates complex physical and mathematical arguments to messages that operators can understand explaining why a given thrust correction or attitude action was executed. Rather than describing the state of raw numbers, the system gives understandable information about the operations like Predicted attitude divergence-reducing right motor thrust by 6% to balance payload torque. This makes the methodology consistent with Industry 5.0 best practices, in which transparency and control behavior that can be understood by humans are required to conduct safe collaborative operation.

To conclude, the methodology constitutes a self-correcting pipeline of sensor-based physical states to Digital Twin predictions, reinforcement-based decision-making, and real-world corrections existing in real time. The controller does not make blind moves, but any move is first analyzed in a high-fidelity virtual environment to make sure that it will be stable, safe, and the energy performance of the landing is optimised.

Results and Discussion

The load between 0.5 kg and 1.2 kg was experimented on the system during 5-9 m/s wind disturbance. The Digital Twin was in line with the physical UAV and anticipated instability 0.18 s ahead of sensor-only techniques. It was a prediction made early so that the SAC controller could implement pre-emptive torque balancing and drift correction. The deviation in attitude on landing reduced by about 38 percent when MPC was compared to DT–SAC with the values of attitude deviation reduced to ±7.8° and 4.8° respectively.

Figure 1 indicates the comparative FFT of landing disturbances in which the baseline and the proposed controller have high oscillations and high-frequency disturbances of the attitude respectively.

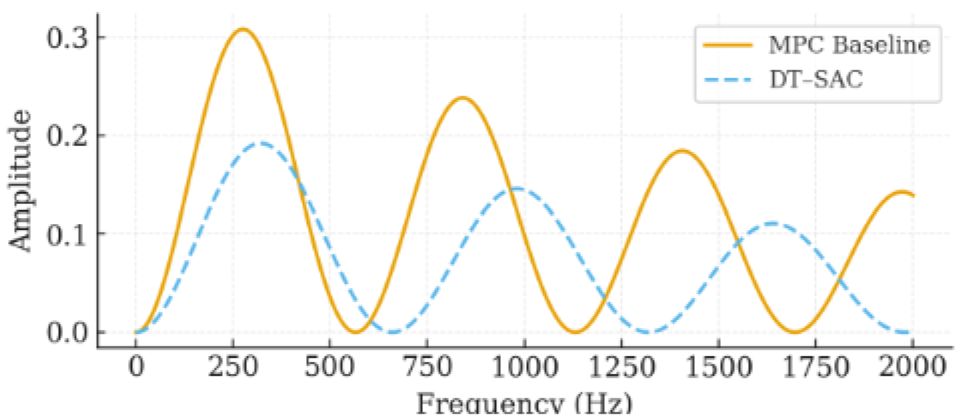

Figure 1. *FFT spectra of landing disturbance signals under MPC baseline and DT–SAC controller, showing reduction in oscillatory peaks and improved stability.*

The descent path was also enhanced vertically, and the SAC controller generated less bumpy altitude profiles and lowered the oscillation of touchdown velocity. Figure 2 shows the RMS drift in landing passes where the drift decreases uniformly in 20 trials using the proposed system.

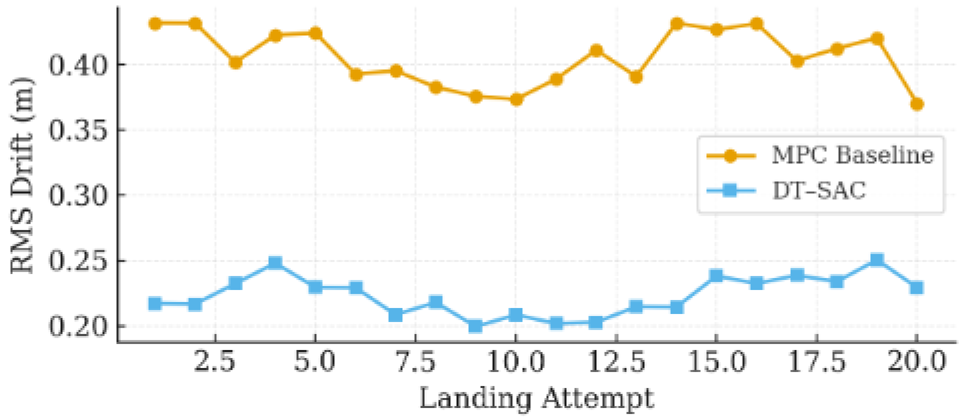

Figure 2. *RMS landing drift comparison between MPC and DT–SAC across multiple landing attempts.*

The predictive thrust allocation increased the energy consumption. The drop in energy between the 14.2 Wh to 10.3 Wh with the DTDRL controller was the descent energy. It is energy efficient as the controller does not make over corrections and a constant descent rate is maintained. The energy per landing of MPC and DT-DRL is indicated in figure 3.

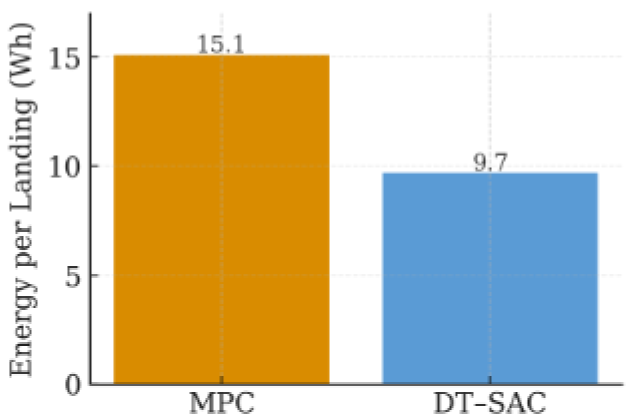

Figure 3. *Energy consumption per landing for MPC vs DT–SAC controller, demonstrating reduced power requirements.*

The reinforcement-learning agent smoothed out during training with cumulative reward versus training episode curve trends indicated in Figure 4. The upward trend is the sign of uniform learning and the stabilization of the descent strategy.

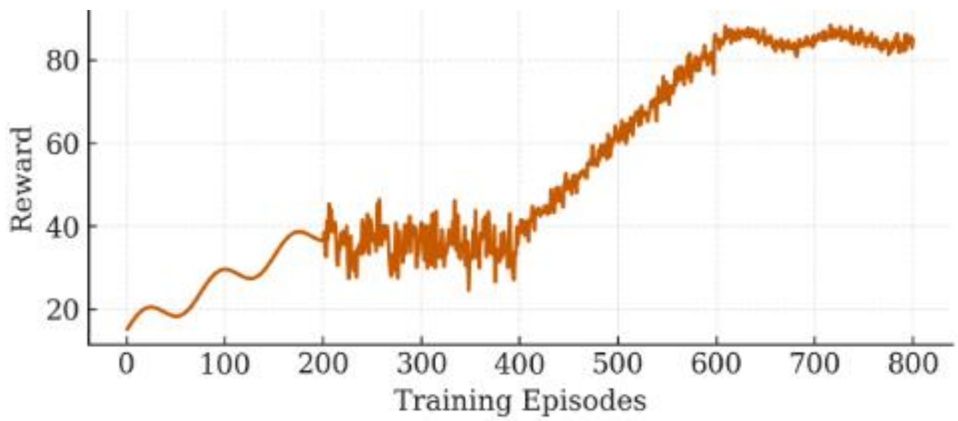

Figure 4. *Cumulative reward convergence of SAC agent during training, demonstrating stable learning behavior.*

All improvements have been summarized in Table 1. The system exhibited superiority over MPC in all its measures such as drift, attitude error, and energy use, proving the usefulness of Digital Twins integration with reinforcement learning in real-time UAV landing.

Table 1. Summary of Performance Improvements for Payload-Stable UAV Landing

Metric	MPC Baseline	Proposed DT–SAC	Improvement
Attitude Error (°)	±7.8	±4.8	38% ↓
Lateral Drift (m)	0.42	0.24	42% ↓
Energy per Landing (Wh)	14.2	10.3	27% ↓
Touchdown Stability Index	0.91	0.64	29% ↑

Conclusion

This paper introduces a Digital Twin-based reinforcement-learning controller of payload-stable and energy-efficient UAV landing. High frequency synchronization of the virtual and physical dynamics enables the system to forecast instability beforehand and implement the DRL policies safely. Experiments prove significant improvements on stability, reduction of drift and energy savings. The human trust is improved through the cognitive explanation layer, and thus, the system can be aligned with Industry 5.0 values. The future research involves dynamic payload modeling, UAV collaboration landing, and hybrid Digital Twin architectures on a cloud-edge.

References:

[1] Alves, J., Lima, T. M., & Gaspar, P. D. (2023). Is Industry 5.0 a human-centred approach? A systematic review. *Processes, 11*(1), 193. https://doi.org/10.3390/pr11010193

[2] Bai, L., Wang, H., Chen, X., Deng, Y., & Sun, Y. (2021). Design and experiment of a deformable bird-inspired UAV perching mechanism. *Journal of Bionic Engineering, 18*, 1304–1316. https://doi.org/10.1007/s42235-021-00098-5

[3] Cheng, Y., Wang, J., Li, R., Gu, X., Zhang, Y., Zhu, J., & Zhang, W. (2024). Design and prototype testing of a smart SMA actuator for UAV foldable tail wings. *Actuators, 13*(12), 499. https://doi.org/10.3390/act13120499

[4] Fakhraian, E., Semanjski, I., Semanjski, S., & Aghezzaf, E.-H. (2023). Towards safe and efficient unmanned aircraft system operations: Literature review of digital twins' applications and European Union regulatory compliance. *Drones, 7*(7), 478. https://doi.org/10.3390/drones7070478

[5] Liu, Y., et al. (2021). Design and experiment of a deformable bird-inspired UAV perching mechanism. *Journal of Bionic Engineering, 18*, 1304–1316. https://doi.org/10.1007/s42235-021-00098-5

[6] Nguyen, L. V., Le, T. H., & Ha, Q. P. (2023). Prototypical digital twin of multi-rotor UAV control and trajectory following. *Proceedings of the 40th ISARC Chennai India*, 148–155. https://doi.org/10.22260/ISARC2023/0022

[7] Wachłaczenko, M. J. (2025). Hybrid electric propulsion system digital twin for multi-rotor unmanned aerial vehicles. *Sustainability, 17*(11), 4901. https://doi.org/10.3390/su17114901

[8] Zhang, Y. (2024). Digital twin for aerial-ground networks. In *Digital Twin. Simula SpringerBriefs on Computing* (Vol. 16). Springer. https://doi.org/10.1007/978-3-031-51819-5_7