Fully AI-based UAV controllers can rival hybrid DRL–PID systems in performance, especially under adversarial conditions.

An End-to-End Reinforcement Learning Controller for Inclined UAV Landing

Background

- ❑ Conventional controllers such as PID and linear MPC are effective for flat-surface UAV landings but struggle with the nonlinear, underactuated dynamics present in inclined scenarios.
- ❑ To our knowledge, this is the first work to explore an end-to-end DRL controller for UAV landing on sloped platforms under environmental disturbances.

Methods

1 Drone Controller Architecture

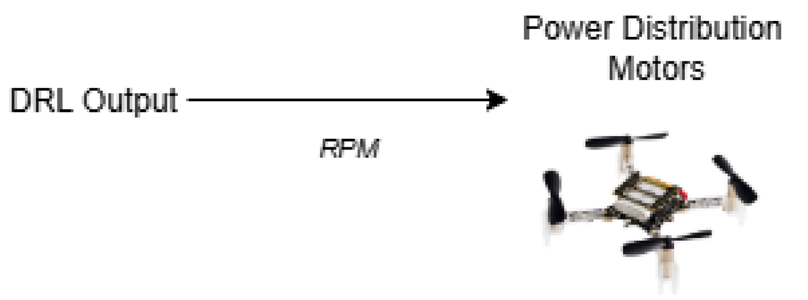

2 DRL Framework

$$\mathbf{o}_t = [\mathbf{p}_t, \boldsymbol{\phi}_t, \mathbf{v}_t, \boldsymbol{\omega}_t] \in \mathbb{R}^{12}$$

$$\mathbf{a}_t = [P_0, P_1, P_2, P_3]_t \in [0,1]^4$$

$$r_t = \begin{cases} 1 & \text{if } s_t \in S_{\text{success}} \\ -8 & \text{if } s_t \in S_{\text{crash}} \\ -1 & \text{if } s_t \in S_{\text{trunc}} \cup S_{\text{bounds}} \\ -0.1 & \text{otherwise} \end{cases}$$

3 Curriculum Learning

$$\delta_d(k) = \max\left(\delta_d^{\text{final}}, \delta_d^{\text{start}} - \alpha_d(k - k_0)\right)$$
$$\delta_v(k) = \max\left(\delta_v^{\text{final}}, \delta_v^{\text{start}} - \alpha_v(k - k_0)\right)$$
$$\delta_a(k) = \min\left(\delta_a^{\text{final}}, \delta_a^{\text{start}} + \alpha_a(k - k_0)\right)$$

4 Disturbance Modelling

$$G_i = k_G k_F \left(\frac{r_P}{4h_i}\right)^2 P_i^2$$

$$\mathbf{D} = -\mathbf{k}_D \left(\sum_{i=0}^{3} \frac{2\pi P_i}{60}\right) \dot{\mathbf{x}}$$

$$\mathbf{F}_{\text{wind}}(t) = \begin{cases} \mathbf{F}_{\text{applied}}(t), & \text{if } W_e = 1 \text{ and } W_t = 1 \\ 0, & \text{otherwise} \end{cases}$$

5 Pybullet Physics engine simulation environment for Model Training

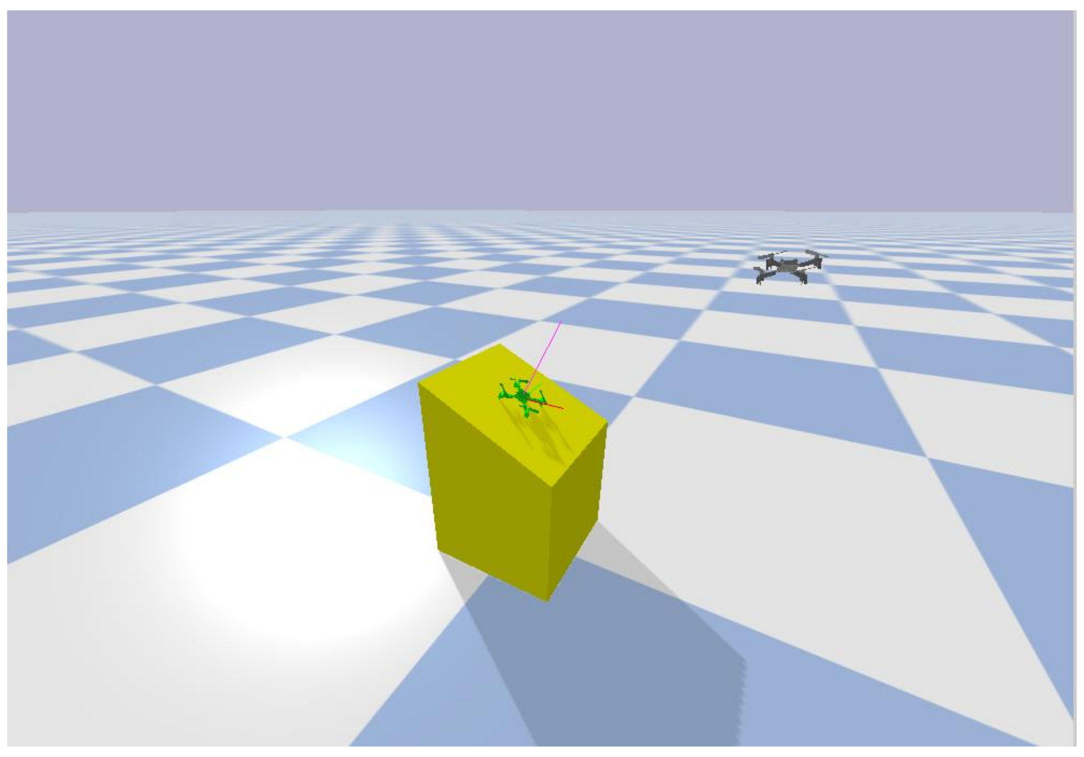

Result 1: Landing success rates of E2E controller vs hybrid controller with and without disturbance

Table 1: LSR under no disturbance

Slope Angle (°)	Hybrid (%)	End-to-End PPO (%)
0	95	95
15	95	80
30	90	55

Table 2: LSR with probabilistic disturbances activation

Slope Angle (°)	Hybrid (%)	End-to-End PPO (%)
0	95	95
15	100	100
30	95	90

Result 2: Trajectory performance indicative of stable and robust learned policy.

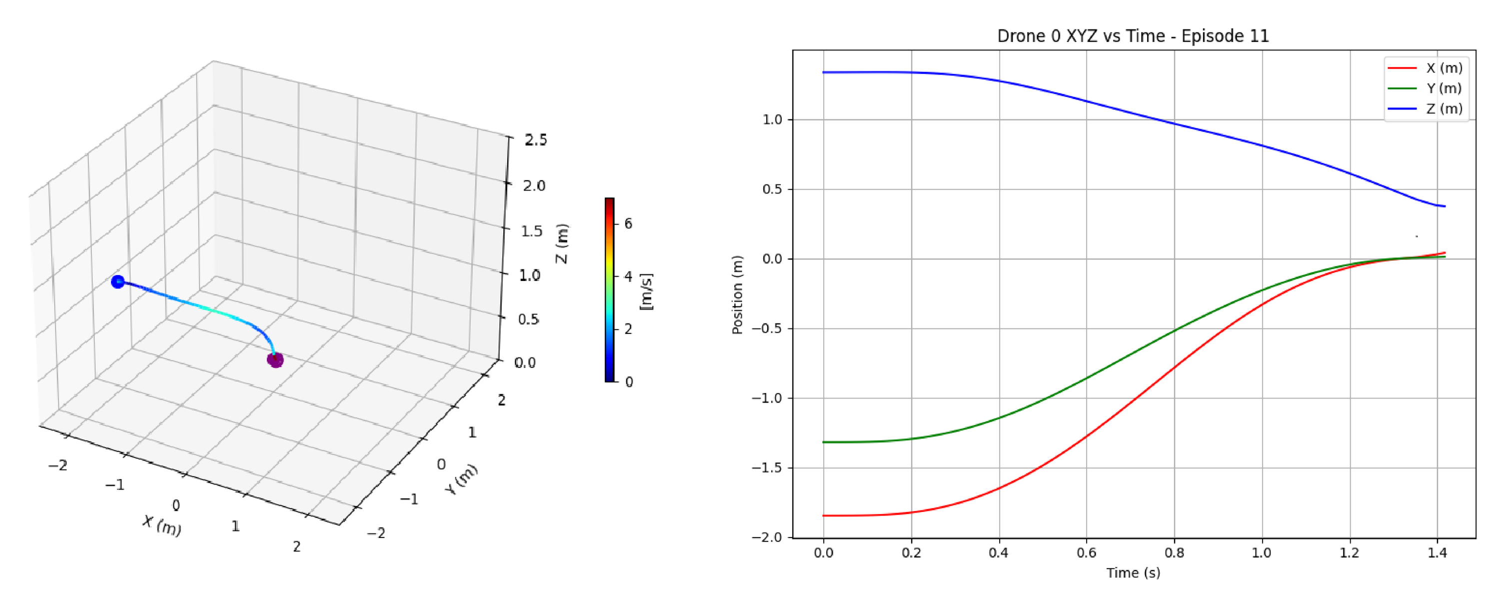

References

[1] J. Amendola, L. R. Cenkeramaddi, and A. Jha, "Drone Landing and Reinforcement Learning: State-of-Art, Challenges and Opportunities," IEEE Open Journal of Intelligent Transportation Systems, vol. 5, pp. 520–539, 2024.

[2] E. Kaufmann, L. Bauersfeld, and D. Scaramuzza, "A Benchmark Comparison of Learned Control Policies for Agile Quadrotor Flight," Feb. 2022

[3] J. E. Kooi and R. Babuška, "Inclined Quadrotor Landing using Deep Reinforcement Learning," 2021.

[4] F. Rein, "The Landing of a Quadcopter on Inclined Surfaces using Reinforcement Learning," 2020.

Nyakotey Nathaniel
nnyakotey@gmail.com