PMP–NMPC for Quadrotor Landing Under Learned Downwash Disturbances

Pratik Mukherjee[1]

***Abstract*— This extended abstract presents a Pontryagin-augmented Nonlinear Model Predictive Control (PMP–NMPC) framework for stable quadrotor landing under strong, state-dependent downwash disturbances. The key idea is to embed Pontryagin's Minimum Principle (PMP) directly within the NMPC structure, yielding analytic optimality conditions and an indirect shooting formulation. A spectrally-normalized deep neural network learns the residual downwash dynamics and injects them into the optimal control law, enabling disturbance-aware optimal control. A local finite-time practical stability result is derived using a Lipschitz-based contraction mapping formulation over motor-level inputs. Simulation studies demonstrate stable descent and improved disturbance rejection.**

I. Introduction

Autonomous aerial vehicles often operate in environments requiring precise *close-proximity flight*, where strong, state-dependent *downwash* effects pose major challenges for stability and safety. During descent, propeller-induced airflow generates nonlinear aerodynamic disturbances that degrade tracking performance and can destabilize standard NMPC controllers, particularly when these residual dynamics are difficult to model accurately.

To address this, we develop a unified *Pontryagin-based Nonlinear Model Predictive Control (PMP–NMPC)* framework that embeds the optimality conditions from Pontryagin's Minimum Principle (PMP) directly within the predictive control architecture. This indirect structure provides analytic gradients, enhances sensitivity to modeling errors, and enables principled adaptation to unmodeled dynamics. A key theoretical contribution is a rigorous *stability analysis* showing that the resulting PMP–NMPC controller remains locally finite-time practically stable despite residual disturbances.

To capture the nonlinear and time-varying nature of downwash, a spectrally-normalized Deep Neural Network (DNN) is integrated into the PMP–NMPC loop. The network learns residual downwash forces offline and injects these estimates into the computed optimal control law, enabling disturbance-aware optimal control without sacrificing stability or constraint satisfaction. This architecture (Fig. 1) raises several core questions regarding stability under learned dynamics, optimality for close-proximity flight, and the safe integration of learning within model-based control. Robotic control for close-proximity aerial maneuvers, particularly precision landing under downwash, has been widely explored [1]–[5]. Prior NMPC approaches address elements of the problem but lack formal optimality structure or robustness guarantees under strong, state-dependent disturbances. Neural Lander [2] models downwash using learned dynamics, but without a principled trajectory optimization framework. In contrast, this work develops a Pontryagin-augmented NMPC formulation that embeds analytic optimality conditions and integrates spectrally-normalized DNN-based residual downwash learning directly into the control law. The resulting controller provides stable, adaptive, and disturbance-aware performance for close-proximity quadrotor maneuvers. Simulation studies of precision landing demonstrate that the learned disturbance model significantly improves descent smoothness, input efficiency, and terminal accuracy relative to PMP-NMPC controller without using learned downwash dynamics as *disturbance rejection*. To our knowledge, this is the first formulation that (i) augments NMPC with PMP for close-proximity aerial maneuvers, (ii) provides stability guarantees with learned residual dynamics, and (iii) achieves real-time adaptation to strong downwash disturbances. The resulting framework advances robust, adaptive, and optimal control for autonomous aerial systems operating in disturbance-rich environments.

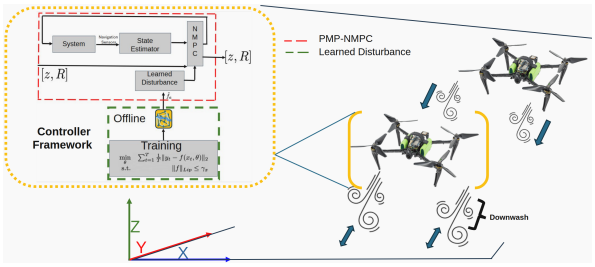

Fig. 1. Overview of the proposed PMP–NMPC framework with learned downwash compensation. The upper (red) block shows the closed-loop controller using real-time state estimation [z: altitude (vertical position), R: rotation matrix (orientation)] and the learned disturbance model, while the lower (green) block illustrates the offline neural-network training of residual downwash dynamics. The right panel depicts close-proximity flight scenarios where interaction-induced downwash (i.e. quadrotor landing or multi-robot interactions) is significant, and where the learned model is incorporated into the PMP–NMPC loop for stable, adaptive control.

II. PMP–NMPC Formulation

Problem Statement: This work addresses the problem of computing a stable and optimal quadrotor trajectory during close-proximity flight, particularly descent and landing, where strong downwash-induced disturbances significantly affect system dynamics. We formulate this task as a nonlinear *optimal control problem* (OCP) and solve it using a *Pontryagin-augmented Nonlinear Model Predictive Control* (PMP–NMPC) framework. The PMP structure provides analytic optimality

[1]P. Mukherjee is with Ocean and Mechanical Engineering Department, Florida Atlantic University, Boca Raton, FL USA, {pmukherjee}@fau.edu

conditions for real-time control under constraints, while a DNN learns the unknown residual downwash dynamics and injects them into the quadrotor dynamics as disturbance rejection. The resulting controller ensures stability and robustness to aerodynamic uncertainties during disturbance-dominated maneuvers.

Consider the nonlinear dynamics

$$\dot{x} = f(x) + g(x)u, \tag{1}$$

with state $x \in \mathbb{R}^n$ and control $u \in \mathbb{R}^m$. The NMPC cost over a horizon T_p is

$$J = \int_t^{t+T_p} (x - x_r)^T Q (x - x_r) + u^T R u \, d\tau, \tag{2}$$

with terminal penalty $(x - x_r)^T P (x - x_r)$.

PMP introduces a costate λ and Hamiltonian

$$H(x, \lambda, u) = \lambda^T (f(x) + g(x)u) + (x - x_r)^T Q (x - x_r) + u^T R u. \tag{3}$$

The optimality system becomes the *two-point boundary value problem* (TPBVP):

$$\dot{x} = f(x) + g(x)u, \tag{4}$$

$$\dot{\lambda} = -\nabla_x H, \tag{5}$$

$$u^* = \arg\min_{u \in U_C} H. \tag{6}$$

For quadratic R, the minimizer yields the explicit control law: $u^* = -\frac{1}{2} R^{-1} g(x)^T \lambda$, subject to saturation $u \in U_C$. Thus, PMP–NMPC converts the constrained OCP into a structured indirect optimality system with analytic gradients and together with the costate dynamics and boundary conditions forms a TPBVP solved over the prediction horizon to obtain the costate λ and the explicit control law u^*.

III. Disturbance-Aware Optimal Control

Quadrotor vertical dynamics include unknown downwash forces f_a. A neural network $\hat{f}_a(\zeta, \eta)$ (spectrally normalized) learns the residual dynamics. The disturbance-compensated control becomes $u = u^* - \hat{f}_a(\zeta, \eta)$. Since $\hat{f}_a$ is 1-Lipschitz, its Jacobian is uniformly bounded, enabling consistent integration in the PMP costate dynamics: $\dot{\lambda} = -\nabla_x H(x, \lambda, u, \hat{f}_a)$

IV. Stability via Lipschitz Contraction

Local finite-time practical stability is established using a contraction mapping formed at the motor-level input η:

$$\eta_k = B_o^{-1} \left[u^* - \hat{f}_a(\zeta, \eta_{k-1}) \right]. \tag{7}$$

where B_o is the quadrotor control allocation matrix that maps motor speeds to total thrust and body torques. If $\sigma(B_o^{-1})\,\Gamma < 1$ (with Γ the Lipschitz constant of the closed-loop dynamics), the mapping is contractive:

$$\|F(\eta_1) - F(\eta_2)\| \leq \alpha \|\eta_1 - \eta_2\|, \ \alpha < 1, \tag{8}$$

ensuring convergence of the discrete control update. Together with constraint tightening and Lyapunov arguments, this yields finite-time practical stability of the closed-loop PMP–NMPC system.

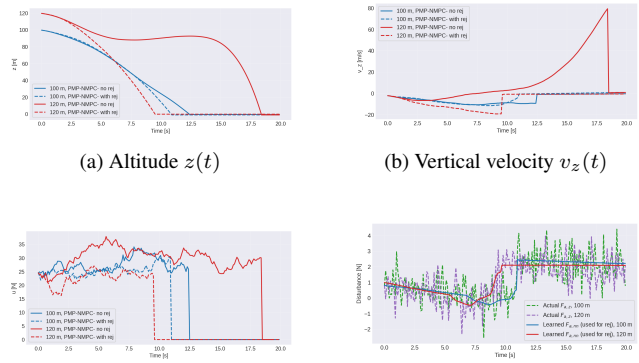

(a) Altitude $z(t)$ (b) Vertical velocity $v_z(t)$

(c) Control input $u(t)$ (d) Actual vs. learned downwash

Fig. 2. PMP–NMPC landing performance with and without learned downwash compensation for 100 m and 120 m descents.

V. Simulation Results

Simulation studies (Fig.2) on 100m and 120m descent scenarios confirm the effectiveness of the proposed approach. The PMP–NMPC controller with learned downwash (`with rej`) compensation achieves stable and accurate landing (Fig.2a) despite strong aerodynamic disturbances, reducing thrust variability (Fig.2c) and shortening descent time relative to PMP-NMPC without rejection (`no rej`). The spectrally-normalized DNN provides reliable residual force estimates (Fig.2d) throughout the maneuver, enabling smoother costate evolution and more consistent satisfaction of the PMP optimality conditions. Collectively, these results demonstrate the advantages of integrating learned dynamics within a principled PMP-based optimal control framework for disturbance-rich aerial operations.

VI. Conclusion

This work demonstrates that Pontryagin-augmented NMPC, combined with spectrally-normalized learned downwash compensation, yields a theoretically grounded and disturbance-aware control architecture for quadrotor landing. The indirect optimality structure enables analytic control laws and stability guarantees rarely available in learning-based aerial control.

References

[1] M. Jacquet, M. Kivits, H. Das, and A. Franchi, "Motor-Level N-MPC for Cooperative Active Perception With Multiple Heterogeneous UAVs," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2063–2070, Apr. 2022, multimedia file: 10.1109/LRA.2022.3143218/mm1. [Online]. Available: https://ieeexplore.ieee.org/document/9682606/

[2] G. Shi, X. Shi, M. O'Connell, R. Yu, K. Azizzadenesheli, A. Anandkumar, Y. Yue, and S.-J. Chung, "Neural Lander: Stable Drone Landing Control using Learned Dynamics," *2019 International Conference on Robotics and Automation (ICRA)*, pp. 9784–9790, May 2019, arXiv: 1811.08027. [Online]. Available: http://arxiv.org/abs/1811.08027

[3] M. Jacquet and A. Franchi, "Motor and perception constrained nmpc for torque-controlled generic aerial vehicles," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 518–525, 2020.

[4] M. Pagone, M. Boggio, C. Novara, A. Proskurnikov, and G. C. Calafiore, "A barrier function approach to constrained pontryagin-based nonlinear model predictive control," *arXiv preprint arXiv:2203.16377*, 2022.

[5] E. Pagone *et al.*, "Continuous-time nmpc using penalty methods for constrained systems," *Automatica*, 2024.