Learned Incremental Nonlinear Dynamic Inversion for Quadrotors with and without Slung Payloads

Eckart Cobo-Briesewitz, Khaled Wahba, and Wolfgang Hönig

I. INTRODUCTION

The rise in complexity of tasks to be executed by multirotors has necessitated the development of more accurate flight controllers that are able to model all forces acting on these robots. Traditional flight controllers [1] model a large part of these forces but are unable to model all of them. The remaining forces, called residual forces, can arise from diverse sources such as blade flapping, drag, or ground forces [2]–[4]. Trying to compute residual forces directly can be computationally complex and intractable for real time control, especially on quadrotors with restricted hardware.

Recent work has shown that learning algorithms for residual force prediction can lead to significant improvements in flight performance [4], even modeling the interaction forces between multiple quadrotors [2]. By only learning the residual forces rather than the full dynamics of the quadrotor, the amount of training data needed is greatly reduced. In addition, learning models can help improve flight performance when sensor readings are noisy or lacking.

Incremental Nonlinear Dynamic Inversion (INDI) [5] derives the residual forces from the mismatch between the nominal dynamic model and real-time sensor feedback. Although INDI can significantly enhance flight performance, its effectiveness depends on the availability of specific sensor data and may still be affected by measurement noise.

We propose a neural-network-based approach that achieves comparable flight performance to the INDI controller without requiring rotor RPM measurements on both a single multirotor system and one carrying a slung type payload. By replacing the incremental model with a learned component integrated into two traditional flight controllers [1], [6], our approach removes the need for specialized sensors. Quantitative results show that it matches INDI's performance, and we further analyze a hybrid setup combining both INDI and neural network predictions. Figure 1 summarizes the three evaluated approaches.

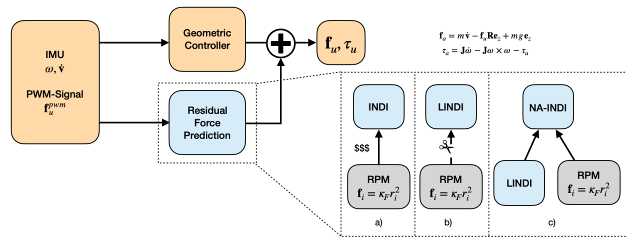

Fig. 1. Overview of the three approaches for computing residual forces compared in our paper: a) Standard INDI relies on additional sensor measurements, which can be expensive and prone to noise. b) LINDI learns the outputs of INDI, eliminating the need for specialized sensors and yielding smoother predictions. c) NA-INDI fuses LINDI predictions with the original sensor data, resulting in higher-quality residual force estimates.

II. APPROACH

A. Problem Definition

Given the nominal models and feasible reference trajectories $\mathbf{x}_r(t)$, we are aiming to find controllers that can minimize the mean tracking error:

$$\underset{\pi}{\operatorname{argmin}} \frac{1}{D} \int_{t=0}^{D} d(\mathbf{x}_\pi(t), \mathbf{x}_r(t))dt, \tag{1}$$

where $\pi(\mathbf{x}, \mathbf{x}_r) \mapsto \mathbf{u}$ is the control law that influences the state $\mathbf{x}_\pi(t)$, D the duration of the reference trajectory $\mathbf{x}_r(t)$, and d is a distance metric. For the multirotor case, we consider the positional tracking error, i.e., $d(\mathbf{x}_\pi, \mathbf{x}_r) = \|\mathbf{p}_\pi - \mathbf{p}_r\|_2$ and for the payload transport case the positional tracking error of the payload.

In this work, we augment classical controllers that ignore the residual force $\mathbf{f}_a$ and torque $\boldsymbol{\tau}_a$ with measured and/or predicted components to improve the tracking error (1).

B. Background: Incremental Nonlinear Dynamic Inversion (INDI)

The key idea of Incremental Nonlinear Dynamic Inversion (INDI) is to estimate the residual force $\mathbf{f}_a$ and torque $\boldsymbol{\tau}_a$ in real time from the discrepancy between IMU-based state measurements and the state evolution predicted by nominal dynamics and measured RPM inputs. INDI has been implemented on multirotors in [5], [7].

Within the INDI framework, the multirotor–payload nominal dynamics can be formulated in different ways. Some approaches explicitly include the dynamics of the payload, cable, and multirotors [6], [8], [9], while others capture the coupling through the cable tension, balancing the forces between the UAV and the payload [10], [11].

C. Learned Incremental Nonlinear Dynamic Inversion (LINDI)

Using a dataset containing example trajectories with IMU data, RPM data, and state estimates, one can also compute $\mathbf{f}_a$ and $\boldsymbol{\tau}_a$. In this case, noisy sensor data can be pre-processed

All authors are with Technical University of Berlin, Germany cobo-briesewitz@campus.tu-berlin.de, k.wahba@tu-berlin.de, hoenig@tu-berlin.de.
This work has been accepted at the 7th Annual Learning for Dynamics & Control Conference (L4DC).

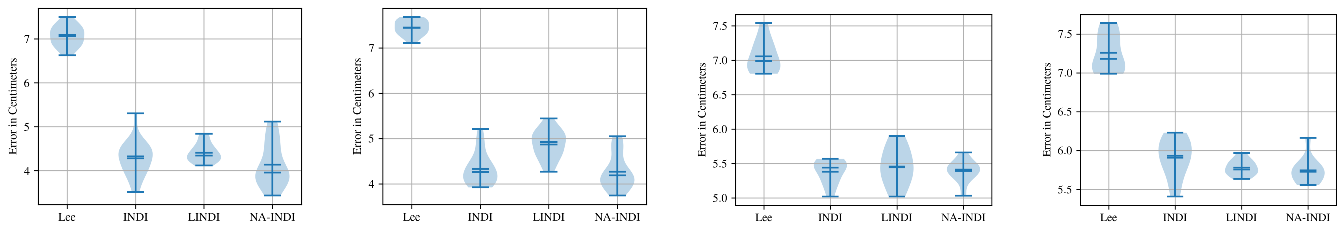

(a) Figure8 (no payload) (b) Circle (no payload) (c) Figure8 (payload) (d) Circle (payload)

Fig. 2. Performance comparison between all four methods. Left to right: no payload (Figure8, Circle) and payload (Figure8, Circle). Each violin plot shows mean tracking error (centimeters) (1) for ten trials.

using zero-delay filters, such as spline fitting. The resulting values of $\mathbf{f}_a$ and $\boldsymbol{\tau}_a$ are the labels of a supervised learning problem.

In the single quadrotor case, we use a multi-layer perceptron (MLP) with 19 inputs, 6 outputs, 2 hidden layers with 24 dimensions each, and Leaky-ReLU activation. The input includes $\mathbf{v}$, $\dot{\mathbf{v}}$, $\boldsymbol{\omega}$, the first two columns of the rotation matrix $\mathbf{R}$, and the motor PWM signal (which is different from r_i as it cannot observe motor delays). The output is the residual force and torque $(\mathbf{f}_a, \boldsymbol{\tau}_a)^\top \in \mathbb{R}^6$. The inputs and outputs of the network are scaled using min-max normalization to the range $[-1, 1]$ to mitigate disparities in value magnitudes and enhance training stability.

For the payload network the input is expanded to include the payload velocity, acceleration and cable direction, resulting in 28 inputs. The architecture remains the same except for the hidden layers which are reduced to a dimension of 16. This is due to the residual forces in the payload scenario being more pronounced and making the training process easier allowing for a smaller network size.

For pre-processing, we fit splines with cubic polynomial segments minimizing the L_2 error on the datapoints, rather than connecting points exactly. We fit splines on the collected INDI outputs and use this smooth signal as label for the training. The effects of training on smooth data can be seen in Figure 3.

D. *Neural-Augmented Incremental Nonlinear Dynamic Inversion (NA-INDI)*

We split the unmodeled dynamics in two parts $\mathbf{f}_a = \mathbf{f}_{a,NN} + \mathbf{f}_{a,INDI}$ and similar for $\boldsymbol{\tau}_a$, where $\mathbf{f}_{a,NN}$ and $\boldsymbol{\tau}_{a,NN}$ are learned functions that were trained using the steps described in II-C. Then the INDI control law only needs to reason about the remaining mismatch in the dynamics.

III. RESULTS

We compare our methods with INDI and the geometric controllers of [1] for the quadrotor system and [6] for the quadrotor–payload system, without residual force prediction.

Figure 2 shows that residual prediction can reduce the tracking error to the desired flight trajectory by approximately 40%. As seen in the results, LINDI achieves performance comparable to INDI without relying on RPM measurements. The results also indicate that combining sensor measurements with learned residuals (NA-INDI) yields the best performance, albeit by a small margin. This improvement arises because noise in the RPM measurements is easier to filter when the INDI component only needs to account for a smaller portion of the residuals, with the remaining effects handled by the LINDI component.

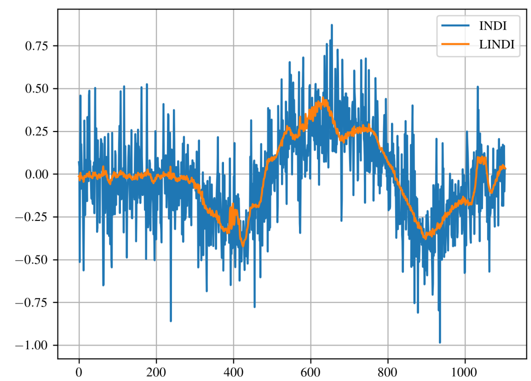

Fig. 3. The figure shows the residual force predictions of both INDI and LINDI as accelerations along the y-axis during a Figure8 flight trajectory for a quadrotor carrying a payload. Using splines to smooth the data collected with the RPM measurements helps the MLP learn smooth residuals and can reduce the amount of training data needed to produce accurate predictions.

IV. CONCLUSION

We introduce LINDI, a novel approach for learning residual predictions which trains a neural network on smooth data based on spline fitting. The results demonstrate that INDI can be effectively replaced by a neural network, which can output smoother residual force predictions without requiring special sensor measurements. While the proposed method requires a quadrotor equipped with RPM measurement sensors to collect the training data, once the neural network has been trained, it can be deployed on any number of quadrotors with similar specifications. This can significantly reduce the cost of maintaining a large fleet of drones with accurate flight controllers.

This work lays the groundwork for extending learning-based residual modeling to scenarios with external aerodynamic disturbances. Studying LINDI under such effects may enable robust, disturbance-aware residual prediction in real-world applications.

REFERENCES

[1] T. Lee, M. Leok, and N. H. McClamroch, "Geometric tracking control of a quadrotor uav on se(3)," in *IEEE Conference on Decision and Control (CDC)*, 2010.

[2] G. Shi, W. Hönig, Y. Yue, and S.-J. Chung, "Neural-swarm: Decentralized close-proximity multirotor control using learned interactions," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020.

[3] G. Shi, W. Hönig, X. Shi, Y. Yue, and S.-J. Chung, "Neural-swarm2: Planning and control of heterogeneous multirotor swarms using learned interactions," *IEEE Transactions on Robotics*, vol. 38, no. 2, pp. 1063–1079, 2022.

[4] L. Bauersfeld, E. Kaufmann, P. Foehn, S. Sun, and D. Scaramuzza, "Neurobem: Hybrid aerodynamic quadrotor model," in *Robotics: Science and Systems XVII*. Robotics: Science and Systems Foundation (RSS), 2021.

[5] E. J. J. Smeur, Q. Chu, and G. C. H. E. Croon, "Adaptive incremental nonlinear dynamic inversion for attitude control of micro air vehicles," *American Institute of Aeronautics and Astronautics*, vol. 39, no. 3, pp. 450–461, 2016.

[6] B. Yu, K. Gamagedara, S. Kim, T. Lee, and J. Suk, "Geometric control and experimental validation for a quadrotor UAV transporting a payload," in *Conference on Decision and Control (CDC)*, 2020, pp. 201–207.

[7] E. Tal and S. Karaman, "Accurate tracking of aggressive quadrotor trajectories using incremental nonlinear dynamic inversion and differential flatness," *IEEE Transactions on Control Systems Technology*, vol. 29, no. 3, pp. 1203–1218, 2021.

[8] K. Sreenath, T. Lee, and V. Kumar, "Geometric control and differential flatness of a quadrotor UAV with a cable-suspended load," in *Conference on Decision and Control*. IEEE, 2013, pp. 2269–2274.

[9] G. Li, X. Liu, and G. Loianno, "Rotortm: A flexible simulator for aerial transportation and manipulation," *IEEE Transactions on Robotics*, vol. 40, pp. 831–850, 2023.

[10] S. Sun, X. Wang, D. Sanalitro, A. Franchi, M. Tognon, and J. Alonso-Mora, "Agile and cooperative aerial manipulation of a cable-suspended load," *Science Robotics*, vol. 10, no. 107, p. eadu8015, 2025.

[11] K. Sreenath and V. R. Kumar, "Dynamics, control and planning for cooperative manipulation of payloads suspended by cables from multiple quadrotor robots," in *Robotics: Science and Systems (RSS)*, 2013.