

Hierarchical Meta-learning-based Adaptive Controller

Fengze Xie¹, Guanya Shi², Michael O’Connell¹, Yisong Yue¹, and Soon-Jo Chung¹

Abstract—We study how to design learning-based adaptive controllers that enable fast and accurate online adaptation in changing environments. In these settings, learning is typically done during an initial (offline) design phase, where the vehicle is exposed to different environmental conditions and disturbances (e.g., a drone exposed to different winds) to collect training data. Our work is motivated by the observation that real-world disturbances fall into two categories: 1) those that can be directly monitored or controlled during training, which we call “manageable”; and 2) those that cannot be directly measured or controlled (e.g., nominal model mismatch, air plate effects, and unpredictable wind), which we call “latent”. Imprecise modeling of these effects can result in degraded control performance, particularly when latent disturbances continuously vary. This paper presents the Hierarchical Meta-learning-based Adaptive Controller (HMAC) to learn and adapt to such multi-source disturbances. Within HMAC, we develop two techniques: 1) Hierarchical Iterative Learning, which jointly trains representations to caption the various sources of disturbances, and 2) Smoothed Streaming Meta-Learning, which learns to capture the evolving structure of latent disturbances over time (in addition to standard meta-learning on the manageable disturbances). Experimental results demonstrate that HMAC exhibits more precise and rapid adaptation to multi-source disturbances than other adaptive controllers.¹

I. INTRODUCTION

Real-world autonomous robots often require precise control achieved through accurate dynamic models. However, the interaction between the robot and its environment, such as wind fluctuations, thermal updrafts and downdrafts, propeller wash, and nominal model mismatch, introduces complexities that make it challenging to model specific parts of the system precisely. The complexity and instability of the resulting dynamics can greatly adversely affect the outcome of conventional robot control methods.

To address this problem, we must estimate these dynamics accurately and design a controller that stabilizes the system under these dynamics [1]–[3]. In recent years, the combination of meta-learning and adaptive control has shown promise in estimating unmodeled dynamics, addressing domain shift challenges and real-time adaptation to new environments [4]–[9]. However, real-world environments often contain composite disturbances, which cannot be completely controlled (see Fig. 2), e.g., when a robot moves amidst fluctuating wind conditions. Neglecting the composite structure of disturbances can lead to learning poor representations. We de-

^{*}This work was supported in part by DARPA Learning Introspective Control (LINC) and Supernal.

¹Department of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA 91125, USA. ²Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA.

¹Videos and demonstrations in <https://sites.google.com/view/hmacproject>.

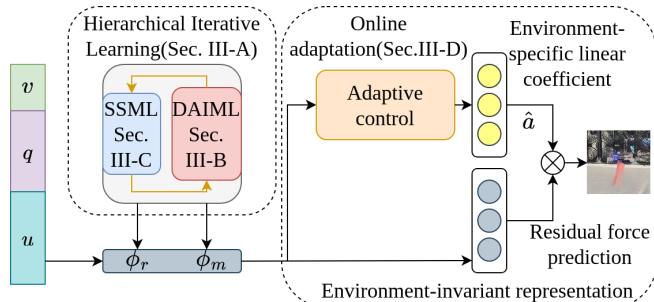


Fig. 1: Flowchart for Hierarchical Meta-learning-based Adaptive Controller (HMAC). HMAC explicitly accounts for multi-source disturbances (see Fig. 2) to learn two representations (ϕ_r and ϕ_m) that enable faster adaptation than standard meta-learning-based adaptive control (only ϕ_m). See Fig. 3 for details on Hierarchical Iterative Learning.

velop Hierarchical Meta-learning-based Adaptive Controller (HMAC), an approach that can learn from multiple types of disturbances (see Fig. 1). In our case, we consider two categories (see Fig. 2): 1) “manageable” disturbances that can be directly controlled and/or monitored during training; and 2) “latent” disturbances that cannot be directly measured or controlled (e.g., nominal model mismatch, air plate effects, and unpredictable wind).² We utilize existing meta-learning approaches to learn neural net representations that are invariant to the manageable disturbances and develop the Smoothed Streaming Meta-Learning (SSML) approach to learn neural net representations of the latent disturbances. As in prior work, we use composite adaptive control for online adaptation.

In our experiments, HMAC shows an average improvement of 26% in tracking performance over Neural-fly (NF) [5], a state-of-the-art work on deep-learning-based adaptive control, and more improvements over all other controllers. These results are accomplished using Crazyflies 2.1 with a thrust upgraded bundle and an air plate. Notably, even when shifting our trained configuration to Crazyflies with an alternate air plate configuration, HMAC still records a 21% performance boost over Neural-fly in tracking metrics.

The main contributions of this article are as follows. First, in Sec. III, HMAC is conceptually and theoretically motivated. Next, in Sec. IV, a series of experimental comparisons performed on an upgraded Crazyflie quadrotor demonstrate the improved performance of HMAC.

²Existing work in meta-learning-based adaptive control [5], [6], [10] only consider the manageable disturbance setting.

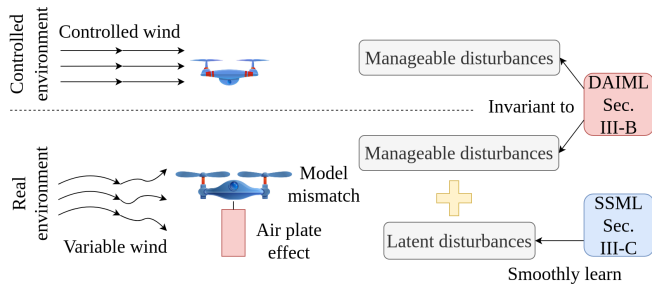


Fig. 2: Depiction of manageable and latent disturbances. A controlled environment can control and measure environment and system conditions; a real-world environment has disturbances that cannot be controlled and/or measured (e.g., variable wind conditions, drone system). Conventional meta-learning-based adaptive control (e.g., using only DAIML [6]) cannot accurately deal with the latent disturbances.

A. Related Work

1) *Multi-environment Learning-based Control*: Meta-learning is a technique to learn an efficient model from data across different tasks or environments [10]–[13], giving the capability to adapt to different domains or tasks with limited data rapidly. In robotics, meta-learning [11], [12] has found applications for robots to accurately adapt to highly dynamic environments [8], [9], [14], [15]. Multi-environment Learning-based Control usually includes two phases: offline learning and online adaptation. In the offline phase, the goal is to learn a model with shared features across all environments. Given limited online data, the online adaptation phase aims to use adaptive control methods [16] to adapt the offline-learned model to a new environment quickly. A subclass of Multi-environment Learning-based Control uses adaptive control methods to adapt a relatively small part of the offline-learned model, which enables a fast and stable online adaptation [4]–[7]. While our method belongs to this class, it distinguishes itself from previous approaches by not assuming a perfectly controlled training environment, which is more realistic in many practical settings.

2) *Hierarchical Iterative Learning*: Multi-model Learning refers to jointly training multiple DNNs together. A well-known subclass of Multi-model Learning is Ensemble Learning. This machine learning technique involves training multiple models and combining their predictions to make more accurate and robust predictions than individual models [17]–[19]. Recently, Ensemble Learning is widely used in DNNs [20]–[23] to improve the performance and robustness of predictions while quantifying the uncertainty in DNNs. Another subclass of Multi-model Learning is Hierarchical Learning, which involves organizing information or tasks into multiple levels of abstraction. The outputs of the models at one level can serve as inputs to the models at the next level, enabling information flow and interaction between the different levels. Hierarchical learning has many applications in different domains and tasks, including computer vision [24] and reinforcement learning [25], [26]. This article

employs an iterative and cascaded training approach to train our additive DNNs effectively, enabling us to disentangle multiple variables from the learning target.

3) *Smoothed Learning on Streaming Data*: Incorporating regularization techniques into the learning process can assist in stabilizing training, alleviating issues with forgetting, and facilitating convergence. Smoothed Online Convex Optimization (SOCO) [27], [28] in online learning incorporates a movement cost term into the optimization problem, which helps control the trade-off between exploration and exploitation. It addresses the challenges of online convex optimization in dynamic and uncertain environments [29], [30]. The Momentum Contrast (MoCo) [31] method is a technique that enables a slowly evolving key encoder that makes it possible to use a queue. It is commonly used in self-supervised learning, particularly in the field of computer vision [32], [33]. The primary allure is its capability to discern and learn data representations, all while circumventing the tedious process of manual labeling. In this article, Smoothed Streaming Meta-Learning utilizes a sliding window approach to sequentially process data over time. Incorporating a movement cost associated with parameter update, it aids in stabilizing the training process and mitigating the forgetting problem.

II. PROBLEM STATEMENT

We consider a general robot system

$$M(q)\ddot{q} + C(q, \dot{q})\dot{q} + g(q) = u + f(q, \dot{q}, w), \quad (1)$$

where q, \dot{q}, \ddot{q} are the n dimensional position, velocity and acceleration vectors; $M(q)$ is the symmetric, positive definite mass and inertia matrix; $C(q, \dot{q})$ is the Coriolis matrix; $g(q)$ is the gravitational force vector; and $u \in \mathbb{R}^n$ is the controller force. $f(q, \dot{q}, w)$ includes all unmodeled dynamics, with w representing underlying environmental conditions.

We can assume the residual dynamics $f(q, \dot{q}, w)$ can be decomposed in the form of

$$f(q, \dot{q}, w) \approx \phi_m(q, \dot{q})a_m(w_m) + \phi_r(q, \dot{q})a_r(w_r), \quad (2)$$

where w_m is a hidden state that represents manageable disturbances, and w_r is a hidden state that represents latent disturbances (see Fig. 2 and example below). ϕ_m and ϕ_r are learned representations that only depend on state; a_m is a set of linear coefficients that are associated with wind conditions; a_r is a set of linear coefficients that captures latent disturbances w_r of each condition.

Example 1: As depicted in Fig. 2, during data collection, we typically encounter both manageable disturbances w_m , and latent disturbances w_r . While we can control manageable disturbances (e.g., control wind profiles of a wind tunnel), latent disturbances persistently fluctuate in the background. While treating the total disturbances as a unified value [5] can lead to $\phi(q, \dot{q})$ memorizing information from w_r , resulting in incorrect values during transitions between the testing and training sets, our decomposition approach is capable of capturing both types of disturbances.

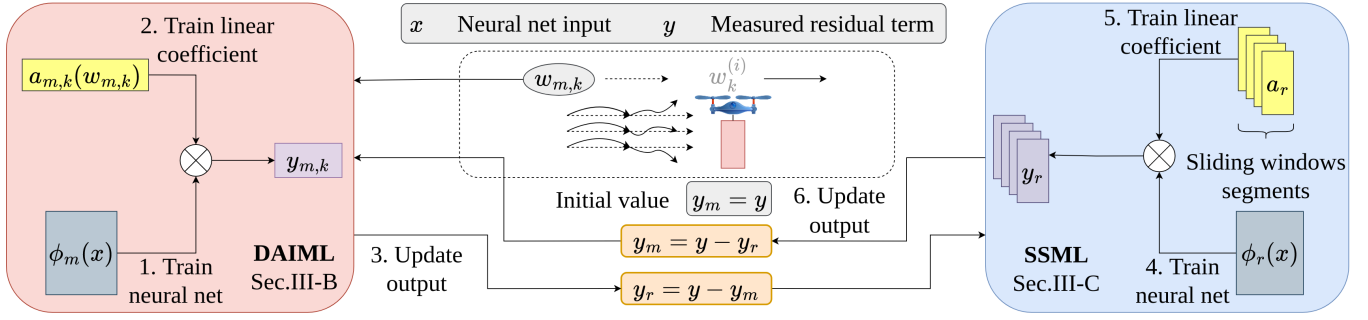


Fig. 3: Hierarchical Iterative Learning flowchart. As mentioned in Figure 2, conventional meta-learning for adaptive control (DAIML) learns a representation ϕ_m that is invariant to manageable disturbances, and Smoothed Streaming Meta-Learning (SSML) smoothly learns a representation ϕ_r of latent disturbances. The two modules run iteratively during offline training.

III. HIERARCHICAL LEARNING WITH ADAPTATION

A. Hierarchical Iterative Learning

We now introduce our Hierarchical Iterative Learning method (see Fig. 3). Hierarchical Iterative Learning integrates two submodules, Domain Adversarial Invariant Meta-Learning (DAIML, Sec. III-B) and Smoothed Streaming Meta-Learning (SSML, Sec. III-C). We assume the presence of K distinct environments with Manageable disturbances in different values from which we can gather data. Our dataset is denoted as $\mathcal{D} = \{D_1, \dots, D_K\}$, where $D_k = \{x_k^{(i)}, y_k^{(i)}\}_{i=1}^{N_k}$ consists of N_k pairs of input and output data. Our primary goal in Hierarchical Iterative Learning is to train two additive DNNs simultaneously. As illustrated in Fig. 3, we employ an iterative training strategy for both DNNs, capitalizing on carefully curated datasets to achieve this. Specifically, we define two new datasets: 1) $D_{m,k} = \{x_k^{(i)}, y_{m,k}^{(i)}\}_{i=1}^{N_k}$, where y_m represents the manageable disturbances and 2) $D_{r,k} = \{x_k^{(i)}, y_{r,k}^{(i)}\}_{i=1}^{N_k}$, where y_r corresponds to the latent disturbances. We commence training the DNN ϕ_m for n epochs, halting when it nears convergence. Subsequently, we initiate an iterative training process for ϕ_r and ϕ_m , each for a fixed number of epochs in every iteration. In this context, a_m is computed for each dataset $D_{m,k}$, while a_r pertains to each sliding window in our Smoothed Streaming Meta-Learning method, as elaborated in Sec. III-C.

B. Domain Adversarial Invariant Meta-Learning

This section briefly overviews the Domain Adversarial Invariant Meta-Learning (DAIML) [5] method we use to train ϕ_m . One challenge in the multi-environment training is the domain shift, for which $\phi_m(x)$ may memorize the distribution of $f(x, w_m)$ in different wind conditions in x rather than w_m . To solve this domain shift problem, the following adversarial optimization framework is applied:

$$\max_h \min_{\phi_m, a_{m,1}, \dots, a_{m,K}} \sum_{k=1}^K \sum_{i=1}^{N_k} (\|y_{m,k}^{(i)} - \phi_m(x_k^{(i)})a_{m,k}\|^2 - \alpha \cdot \text{loss}(h(\phi_m(x_k^{(i)})), k)) \quad (3)$$

where h is a discriminator to predict the environment index out of K environments, loss is the cross-entropy loss, and α is the hyperparameter to control the regulation rate. By applying DAIML, we can make ϕ_m domain-invariant and keep the environment-specific features in w_m .

C. Smoothed Streaming Meta-Learning

For learning ϕ_r , we propose Smoothed Streaming Meta-Learning (SSML), which enables gradual online adaptation of a_r for changing environmental conditions that cannot be directly controlled or measured. SSML is visualized in Fig. 4 and described in detail in Algorithm 1. We have a sliding window of length b on all sequential data across the entire data set. During each iteration, this algorithm executes two primary steps, described below.

The *adaptation step* (Lines 11-13) solves a least squares problem concerning ϕ_r for each adaptation set \mathbf{D}_k . We can perform this adaptation since we assume that w_r changes gradually. Specifically, we choose c data points in each sequential data to randomly serve as the adaptation set \mathbf{b}_k^c and the rest $b - c$ data points \mathbf{b}_k^a serve as the training set.

The *training step* (Line 16) updates ϕ_r on the training set. We address this objective by formulating an optimization problem incorporating regularization.

$$\min_{\phi_r, a_{r,1}, \dots, a_{r,K}} \sum_{k=1}^K \|y_{r,k}^{(i)} - \phi_r(x_k^{(i)})a_{r,k}\|^2 + c(\theta_{\phi_r}, \theta_{\phi_r}^{(t-1)}), \quad (4)$$

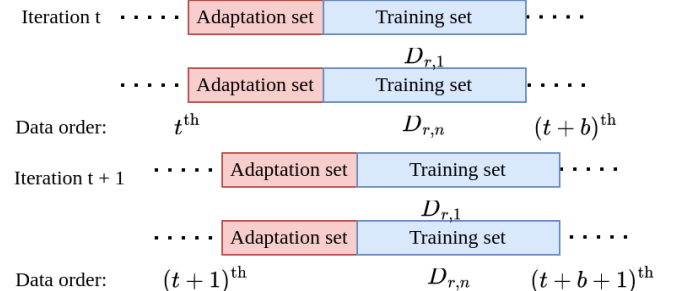


Fig. 4: Smoothed Streaming Meta-Learning visualization

Algorithm 1: Smoothed Streaming Meta-Learning

```
1 Input:  $\mathcal{D}_r = \{D_{r,1}, \dots, D_{r,K}\}$ 
2 Initialize: Neural network  $\phi_r$ 
3 Result: Trained neural network  $\phi_r$ 
4 repeat
5   for  $t \leftarrow 1$  to  $N_k - b$  do
6     Get Batch  $\mathbf{B}_t = \{D_{r,1[t:t+b]}, \dots, D_{r,K[t:t+b]}\}$ 
7     Loss  $\mathbf{L} = 0$ 
8     for  $k \leftarrow 1$  to  $K$  do
9        $\mathbf{d}_k \leftarrow \mathbf{B}_t[k]$ 
10      Split  $\mathbf{d}_k$  into training set  $\mathbf{b}_k^a$  and adaptive
11      set  $\mathbf{b}_k^c$ 
12      Solve the least square problem  $a^*(\phi_r) =$ 
13       $\operatorname{argmin}_a \sum_{i \in \mathbf{b}_k^c} \|y_{r,k}^{(i)} - \phi_r(x_k^{(i)})a\|^2$ 
14      if  $\|a^*\| > \gamma$  then
15         $a^* \leftarrow \gamma \frac{a^*}{\|a^*\|}$ 
16       $\mathbf{L} = \mathbf{L} + \sum_{i \in \mathbf{b}_k^a} \|y_{r,k}^{(i)} - \phi_r(x_k^{(i)})a^*\|^2 +$ 
17       $c(\theta_{\phi_r}, \theta_{\phi_r}^{(t-1)})$ 
18     $\mathbf{L} = \mathbf{L}/K$ 
19    Train DNN  $\phi_r$  using SGD and spectral
20    normalization with loss  $\mathbf{L}$ 
21 until convergence;
```

where c is the cost function that regularizes the change of the trainable parameters w_{ϕ_r} in ϕ_r , and t represents epochs. In this article, we use a quadratic regularization as

$$c(\theta_{\phi_r}, \theta_{\phi_r}^{(t-1)}) = (\theta_{\phi_r} - \theta_{\phi_r}^{(t-1)})^\top W (\theta_{\phi_r} - \theta_{\phi_r}^{(t-1)}). \quad (5)$$

We increment the value of W after each iteration. This regularization serves two primary purposes: 1) It enhances the stability and smoothness of the training curve, and 2) as W approaches a value of 1, the model ϕ_r that performs well on current data remains effective for previous data, aiding in mitigating the forgetting problem inherent in this training approach. During the training step, we aggregate the losses for the training set \mathbf{b}_k^t across different data sequences to facilitate convergence during training. After each iteration, we slide this window by a single data.

D. Online Adaptation

During online adaptation, the goal is to minimize the position tracking error, and we can use the update law from Neural-fly [5] modified for two additive DNNs. Our online adaptive control algorithm is summarized by the following control law, adaptation law, and parameter update equations:

$$u_{\text{HM}AC} = M(q)\ddot{q}_r + C(q, \dot{q})\dot{q}_r + g(q) - Ks - \phi(q, \dot{q})\hat{a} \quad (6)$$

$$\dot{\hat{a}} = -\Lambda\hat{a} - P\phi^\top R^{-1}(\phi\hat{a} - y) + P\phi^\top s \quad (7)$$

$$\dot{P} = -2\Lambda P + Q - P\phi^\top R^{-1}\phi P, \quad (8)$$

where

$$\phi = [\phi_m \quad \phi_r], \quad \hat{a} = [\hat{a}_m, \hat{a}_r]^\top, \quad (9)$$

$$\Lambda = \begin{bmatrix} \lambda_m & 0 \\ 0 & \lambda_r \end{bmatrix}, Q = \begin{bmatrix} Q_m & 0 \\ 0 & Q_r \end{bmatrix}, \quad (10)$$

$u_{\text{HM}AC}$ is the control law, $\dot{\hat{a}}$ is the online linear-parameter update; P is a covariance-like matrix used for automatic gain tuning; $\dot{q} + \Lambda\tilde{q}$ is the composite tracking error; $\tilde{q} = q - q_d$ is the position tracking error; $\dot{q}_r = \dot{q}_d - \Lambda(q_d - q)$ is the reference velocity. y is the measured aerodynamic residual force; and K, Λ, R, Q , and λ are gains. Based on our empirical observations and assumptions, w_m exerts a more substantial influence and more frequent changes on the robot's dynamics. This influence directly affects the ratio between Q_m and Q_r , as well as between λ_m and λ_r . With this adaptation law, $\begin{bmatrix} s \\ \tilde{a} \end{bmatrix}$ exponentially converges toward a bounded region. For detailed proof, we direct readers to Section S5 of [5].

IV. EXPERIMENTS

We employ quadrotors based on the Bitcraze Crazyflie 2.1 platform, enhanced with a thrust upgrade bundle. These quadrotors possess a rotor-to-rotor scale of 9 cm and weigh 41g (including the battery). To facilitate the execution of desired trajectories, we utilize the Crazyswarm framework [34] for communication with the Crazyflies. The quadrotor is outfitted with multiple reflective markers for tracking at 100Hz using the Vicon system. All tested controllers, extended Kalman filter, and neural network evaluation are executed onboard the STM32 microcontroller. We employ the micro SD card deck for data collection, saving binary encoded data at approximately 10 ms intervals.

A. Quadrotor Dynamics

We now introduce the dynamics of our quadrotor. Consider states given by global position, $\mathbf{p} \in \mathbb{R}^3$, velocity $\mathbf{v} \in \mathbb{R}^3$, attitude rotation matrix $R \in \text{SO}(3)$, and body angular velocity $\omega \in \mathbb{R}^3$. Then the dynamics are

$$\dot{p} = v, \quad m\dot{v} = mg + Rf_u + f, \quad (11)$$

$$\dot{R} = RS(\omega), \quad J\dot{\omega} = J\omega \times \omega + \tau_u, \quad (12)$$

where m is the mass, J is the inertia matrix of the quadrotor, $S(\cdot)$ is the skew-symmetric mapping, g is the gravity vector, $f_u = [0, 0, T]^\top$ and $\tau_u = [\tau_x, \tau_y, \tau_z]^\top$ are total thrust and body torques from four rotors predicted by the nominal model, and $f = [f_x, f_y, f_z]^\top$ are forces from unmodelled aerodynamic effects due to varying wind conditions and all other disturbances. Based on the data we collect, we observe that the unmodelled aerodynamics in our experiment mainly affects translational dynamics.

We transform the position dynamics (11) into the equivalent form of (1) by considering $M(q) = mI$, $C(q, \dot{q}) = 0$, and $u = Rf_u$. Our method is implemented within the position control loop, where we utilize it to compute a desired force u_d . Subsequently, this desired force is decomposed into the desired attitude R_d and desired thrust T_d

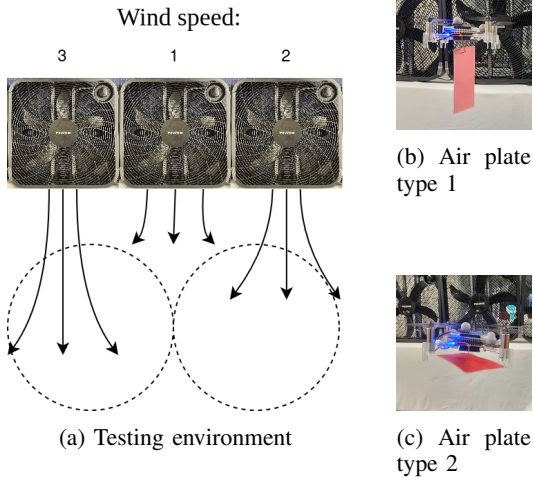


Fig. 5: Experiment Environment. (a) We place three box fans in one row horizontally as our training and testing environment. The dashed line is the figure-8 trajectory in Fig. 7. (b) In the training phase, we use the Crazyflie with air plate type 1. (c) In the testing phase, we use both Crazyflies with air plate type 1 and type 2.

using kinematics. Finally, the desired attitude and thrust are transmitted to our attitude and angular rate controller [35].

B. Data Collection

During the data collection, the drone with the configuration shown in Fig. 5b tracks a trajectory with the PID controller for 1 minute in three conditions generated by box fans. The collection of input-output pairs for each trajectory, denoted as the k^{th} sub-dataset, is referred to as D_{w_k} , where w_k includes a manageable wind condition, $w_{k,m}$, and latent disturbances, $w_{k,r}$. During each trajectory, we gather time-stamped data in $[q, \dot{q}, u]$. Subsequently, we compute the acceleration, \ddot{q} , using a Five-point stencil approach [36]. By combining this acceleration with (1), we obtain a noisy measurement of the unmodeled dynamics, denoted as $y = f(x, w) + \epsilon$, where ϵ encompasses the noise component, and $x = [q, \dot{q}] \in \mathbb{R}^2$ represents the system state. This framework allows us to define the dataset, $D = \{D_{w_1}, \dots, D_{w_K}\}$, where

$$D_{w_k} = \{x_k^{(i)}, y_k^{(i)} = f(x_k^{(i)}, w_k^{(i)}) + \epsilon_k^{(i)}\}_{i=1}^{N_k}. \quad (13)$$

C. Neural Network Architecture and Training Details

We implement our learning framework in Python using Pytorch [37]. We noticed that the aerodynamic effects depend on both velocity v (3-d) quaternion q (4-d) and rotor speed u (4-d) of the quadrotor, we can construct an 11-dimensional vector as our input state $x = [v, q, u]$ to both of our neural network ϕ_m and ϕ_r . ϕ_m consists of four fully-connected hidden layers: $11 \rightarrow 30 \rightarrow 40 \rightarrow 30 \rightarrow 3$, and ϕ_r consists of three fully-connected hidden layers: $11 \rightarrow 20 \rightarrow 10 \rightarrow 2$. The activation function used in each hidden layer is the Rectified Linear Unit (ReLU). The disturbances effect in different directions are highly correlated and share the same




Fig. 6: Residual: This histogram depicts the distribution of disturbances as represented by ϕ_r . We see that the distribution becomes more concentrated after ϕ_r is applied.

feature, so we use single ϕ_r and ϕ_m for all directions. We approximate the aerodynamic force caused by all external forces as follows:

$$f \approx \begin{bmatrix} \phi_m(x) & 0 & 0 & \phi_r(x) & 0 & 0 \\ 0 & \phi_m(x) & 0 & 0 & \phi_r(x) & 0 \\ 0 & 0 & \phi_m(x) & 0 & 0 & \phi_r(x) \end{bmatrix} \begin{bmatrix} a_{m,x} \\ a_{m,y} \\ a_{m,z} \\ a_{r,x} \\ a_{r,y} \\ a_{r,z} \end{bmatrix} \quad (14)$$

where $a_{m,j} \in \mathbb{R}^3$, $a_{r,j} \in \mathbb{R}^2$, $j = \{x, y, z\}$ are linear coefficients for each direction of the aerodynamic force caused by external force. We use adversarial learning mentioned in [5] to train ϕ_m and Algorithm 1 to train ϕ_r . We add spectral normalization to guarantee the stability and convergence of the training of both neural networks [38], [39].

We use data collected in no wind, wind 1, and wind 2 conditions as training data and test on wind 3 conditions. Fig.6 highlights the advantages garnered from employing Hierarchical Iterative Learning, particularly along the z-axis, where DAIML does not fully capture the structure of the residual term. By applying Hierarchical Iterative Learning, the remaining disturbances in the residual term are effectively captured by ϕ_r .

D. Control Performance in Flight Tests

We compare the performance of HMAC with other adaptive controllers, including Neural-fly [5] and INDI [40], as well as PID controller to follow two different trajectories. The first trajectory is a 1.0m long and 0.5m wide figure-8 trajectory; the second is a 0.8m long and 0.8m wide wave trajectory. Both trajectories are under the same wind conditions shown in Fig. 5a, where three different commonly used box fans generate wind with different speeds [41]. We further assess the generalization capabilities of our HMAC and NF models by applying their trained representations to Crazyflies equipped with a distinct air plate configuration (as depicted in Fig. 5c) – referred to as HMAC-transfer and NF-transfer. The flight trajectory for each experiment is shown in Fig. 7. We employ Crazyflies' integrated PID and INDI controllers, fine-tuning their parameters to enhance performance. We see that the PID controller's performance diminishes in the test environment. Meanwhile, the INDI controller relies on IMU-derived acceleration as a reference, although this

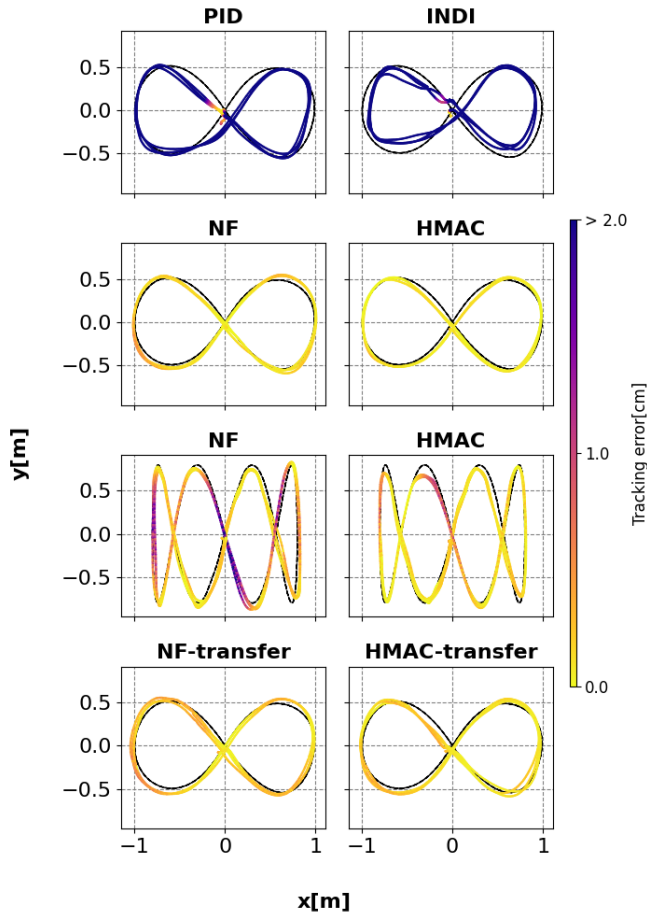


Fig. 7: Trajectory tracking performance of each controller. The PID controller is susceptible to inaccuracies in the dynamics model. INDI [40] controllers are affected by noise in acceleration data from IMUs. NF and HMAC exhibit similar performance characteristics, whereas HMAC demonstrates enhanced performance through more accurate prediction.

value exhibits considerable noise. To ensure system stability, INDI applies a low-pass filter to the IMU acceleration data, resulting in a notably conservative acceleration.

In contrast, NF and HMAC perform considerably well in our testing environment. Analyzing the wave trajectory presented in Fig. 7, we observe that NF encounters challenges in tracking performance along the y-axis. This performance degradation is attributed to the fact that in the proposed DAIML approach, the wind condition remains constant along the y-axis. On the contrary, the wind generated by the box fan exhibits non-linearity and instability, rendering adaptation difficult for NF in this specific scenario. Table I tabulates the

TABLE I: Trajectory tracking error in m .

	Figure 8	Wave	Figure8-transfer	Wave-transfer
PID	0.251	crash	-	-
INDI	0.312	crash	-	-
NF	0.036	0.059	0.047	crash
HMAC	0.027	0.042	0.037	0.051




Fig. 8: The top figure displays the changes of \hat{a} in the time domain, showing differences in magnitude between \hat{a}_m and \hat{a}_r , while the bottom figure presents its power spectral density in the frequency domain, revealing \hat{a}_r has a lower frequency compared to \hat{a}_m .

mean-square error values over all tested trajectories and air plate configurations. In conclusion, HMAC outperforms NF, including figure-8, wave, and figure-8-transfer trajectories.

Furthermore, our analysis examines the variations in a_r and a_m during the execution of a figure-8 trajectory, as illustrated in Fig. 8. Notably, both a_r and a_m exhibit a similar pattern; however, the evolution of a_r is characterized by smoother changes, illustrated in both the time-domain and power spectral density plots. This observation provides empirical support for our SSML assumption.

V. CONCLUSION

This paper presents the Hierarchical Meta-learning-based Adaptive Controller that enables agile and stable flight. Compared to previous work, we avoid the strict restriction on the label of disturbance and achieve better performance under multi-source disturbances since we design a composite adaptive controller that can adapt to multiple disturbances changing in different frequencies. For this purpose, we develop a multi-model training structure with a Smoothed Streaming Meta-Learning method to extract and learn latent disturbances that cannot be easily controlled and labeled. We apply our algorithm on Crazyflies with a thrust upgrade bundle, and our controller provides outstanding performance in trajectory tracking problems under wind disturbances.

We illustrate that the learned latent disturbances are crucial by comparing tracking performance between our and NF controllers. The resulting trajectory demonstrates that our controller exhibits a distinct advantage in accurately tracking trajectories under the environment with latent disturbances. This superiority is achieved through rapid and precise estimation and adaptation to these disturbances.

In future research, we aim to investigate the performance of our algorithm on robot platforms encountering more structured, latent disturbances.

REFERENCES

- [1] D. Hanover, P. Foehn, S. Sun, E. Kaufmann, and D. Scaramuzza, "Performance, precision, and payloads: Adaptive nonlinear MPC for quadrotors," *CoRR*, vol. abs/2109.04210, 2021.
- [2] G. Torrente, E. Kaufmann, P. Foehn, and D. Scaramuzza, "Data-driven mpc for quadrotors," 2021.
- [3] 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.
- [4] S. M. Richards, N. Azizan, J.-J. Slotine, and M. Pavone, "Adaptive-control-oriented meta-learning for nonlinear systems," 2021.
- [5] M. O'Connell, G. Shi, X. Shi, K. Azizzadenesheli, A. Anandkumar, Y. Yue, and S.-J. Chung, "Neural-fly enables rapid learning for agile flight in strong winds," *Science Robotics*, vol. 7, no. 66, p. eabm6597, 2022.
- [6] M. O'Connell, G. Shi, X. Shi, and S.-J. Chung, "Meta-learning-based robust adaptive flight control under uncertain wind conditions," 2022.
- [7] R. Sinha, J. Harrison, S. M. Richards, and M. Pavone, "Adaptive robust model predictive control via uncertainty cancellation," 2022.
- [8] A. Nagabandi, I. Clavera, S. Liu, R. S. Fearing, P. Abbeel, S. Levine, and C. Finn, "Learning to adapt in dynamic, real-world environments through meta-reinforcement learning," 2019.
- [9] S. Belkhal, R. Li, G. Kahn, R. McAllister, R. Calandra, and S. Levine, "Model-based meta-reinforcement learning for flight with suspended payloads," *IEEE Robotics and Automation Letters*, vol. 6, pp. 1471–1478, apr 2021.
- [10] G. Shi, K. Azizzadenesheli, M. O'Connell, S.-J. Chung, and Y. Yue, "Meta-adaptive nonlinear control: Theory and algorithms," 2021.
- [11] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," 2017.
- [12] T. Hospedales, A. Antoniou, P. Micaelli, and A. Storkey, "Meta-learning in neural networks: A survey," 2020.
- [13] W. Xiao, T. He, J. Dolan, and G. Shi, "Safe deep policy adaptation," *arXiv preprint arXiv:2310.08602*, 2023.
- [14] X. Song, Y. Yang, K. Choromanski, K. Caluwaerts, W. Gao, C. Finn, and J. Tan, "Rapidly adaptable legged robots via evolutionary meta-learning," 2020.
- [15] C. D. McKinnon and A. P. Schoellig, "Meta learning with paired forward and inverse models for efficient receding horizon control," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3240–3247, 2021.
- [16] J.-J. E. Slotine and W. Li, *Applied nonlinear control*. Englewood Cliffs, NJ: Prentice Hall, 1991.
- [17] L. Breiman, "Bagging predictors," *Machine learning*, vol. 24, pp. 123–140, 1996.
- [18] Y. Freund and R. E. Schapire, "A decision-theoretic generalization of on-line learning and an application to boosting," *Journal of Computer and System Sciences*, vol. 55, no. 1, pp. 119–139, 1997.
- [19] L. Breiman, "Random forests," *Machine Learning*, vol. 45, pp. 5–32, 10 2001.
- [20] B. Lakshminarayanan, A. Pritzel, and C. Blundell, "Simple and scalable predictive uncertainty estimation using deep ensembles," 2017.
- [21] R. Chandra, R. Chen, and J. Simmons, "Bayesian neural networks via mcmc: a python-based tutorial," 2023.
- [22] A. D. Cobb and B. Jalaian, "Scaling hamiltonian monte carlo inference for bayesian neural networks with symmetric splitting," 2020.
- [23] M. Welling and Y. W. Teh, "Bayesian learning via stochastic gradient langevin dynamics," in *Proceedings of the 28th international conference on machine learning (ICML-11)*, pp. 681–688, 2011.
- [24] T. Hoyoux, A. J. Rodríguez-Sánchez, and J. H. Piater, "Can computer vision problems benefit from structured hierarchical classification?," *Machine Vision and Applications*, vol. 27, no. 8, pp. 1299–1312, 2016.
- [25] T. D. Kulkarni, K. Narasimhan, A. Saedi, and J. Tenenbaum, "Hierarchical deep reinforcement learning: Integrating temporal abstraction and intrinsic motivation," in *Advances in Neural Information Processing Systems* (D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, eds.), vol. 29, Curran Associates, Inc., 2016.
- [26] Y. Yang, G. Shi, X. Meng, W. Yu, T. Zhang, J. Tan, and B. Boots, "Cajun: Continuous adaptive jumping using a learned centroidal controller," 2023.
- [27] M. Lin, Z. Liu, A. Wierman, and L. L. H. Andrew, "Online algorithms for geographical load balancing," in *2012 International Green Computing Conference (IGCC)*, pp. 1–10, 2012.
- [28] G. Goel, Y. Lin, H. Sun, and A. Wierman, "Beyond online balanced descent: An optimal algorithm for smoothed online optimization," in *Advances in Neural Information Processing Systems* (H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, eds.), vol. 32, Curran Associates, Inc., 2019.
- [29] G. Shi, Y. Lin, S.-J. Chung, Y. Yue, and A. Wierman, "Online optimization with memory and competitive control," 2021.
- [30] W. Pan, G. Shi, Y. Lin, and A. Wierman, "Online optimization with feedback delay and nonlinear switching cost," *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, vol. 6, no. 1, pp. 1–34, 2022.
- [31] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, "Momentum contrast for unsupervised visual representation learning," 2020.
- [32] J.-B. Grill, F. Strub, F. Altché, C. Tallec, P. H. Richemond, E. Buchatskaya, C. Doersch, B. A. Pires, Z. D. Guo, M. G. Azar, B. Piot, K. Kavukcuoglu, R. Munos, and M. Valko, "Bootstrap your own latent: A new approach to self-supervised learning," 2020.
- [33] M. Caron, H. Touvron, I. Misra, H. Jégou, J. Mairal, P. Bojanowski, and A. Joulin, "Emerging properties in self-supervised vision transformers," 2021.
- [34] J. A. Preiss, W. Honig, G. S. Sukhatme, and N. Ayanian, "Crazyswarm: A large nano-quadcopter swarm," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3299–3304, 2017.
- [35] D. Mellinger and V. Kumar, "Minimum snap trajectory generation and control for quadrotors," in *2011 IEEE International Conference on Robotics and Automation*, pp. 2520–2525, 2011.
- [36] M. Abramowitz, *Handbook of Mathematical Functions, With Formulas, Graphs, and Mathematical Tables*. USA: Dover Publications, Inc., 1974.
- [37] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, high-performance deep learning library," 2019.
- [38] 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," in *2019 International Conference on Robotics and Automation (ICRA)*, IEEE, may 2019.
- [39] G. Shi, W. Hönig, X. Shi, Y. Yue, and S. Chung, "Neural-swarm2: Planning and control of heterogeneous multirotor swarms using learned interactions," *CoRR*, vol. abs/2012.05457, 2020.
- [40] E. Smeur, Q. Chu, and G. Croon, "Adaptive incremental nonlinear dynamic inversion for attitude control of micro air vehicles," *Journal of Guidance, Control, and Dynamics*, vol. 39, pp. 1–12, 12 2015.
- [41] K. Huang, R. Rana, A. Spitzer, G. Shi, and B. Boots, "Datt: Deep adaptive trajectory tracking for quadrotor control," in *7th Annual Conference on Robot Learning*, 2023.